

Server +

Mega Guide

Prepare With Confidence

This PrepLogic Mega Guide was written by certified subject matter experts and published authors to provide you accurate, in-depth exam coverage. All exam objectives are covered in detail, giving you the knowledge and confidence you need to pass your exam.



PrepLogic

Be Prepared. Be Confident. Get Certified.



Server+ (SK0-003) Mega Guide

Copyright © 2010 by PrepLogic, LLC.

Product ID: 12061

Production Date: April 15, 2010

All rights reserved. No part of this document shall be stored in a retrieval system or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. No patent liability is assumed with respect to the use of the information contained herein.

Warning and Disclaimer

Every effort has been made to make this document as complete and as accurate as possible, but no warranty or fitness is implied. The publisher and authors assume no responsibility for errors or omissions. The information provided is on an "as is" basis. The authors and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this document.

PrepLogic, LearnSmart Video Training, Printables, Lecture Series, Quiz Me Series, Awdeeo, and other PrepLogic logos are trademarks or registered trademarks of PrepLogic, LLC. All other trademarks not owned by PrepLogic that appear in the Software or on Web Site(s) are the property of their respective owners.

Volume, Corporate, and Educational Sales

PrepLogic offers favorable discounts on all products when ordered in quantity. For more information, please contact PrepLogic directly:

1-800-418-6789

solutions@preplogic.com

International Contact Information

International: +1 (813) 769-0920

Australia: (02) 8003 3878

South Africa: (0) 11 083 9973

United Kingdom: (0) 20 8816 8036

Abstract..... 5

What to Know 5

Tips 6

Changes in the Server+ 2009 Exam:..... 7

1.0 Server Background.....8

 1.1 Kinds of Servers 8

 1.2 Server Architectures 11

 1.3 Monitoring Protocols..... 11

2.0 System Hardware12

 2.1 Motherboards..... 12

 2.2 Processors..... 15

 2.3 Memory..... 17

 2.4 BIOS, CMOS Memory, and POST 21

 2.4.1 POST..... 22

 2.5 System Resources 23

 2.6 Multiprocessing 24

 2.7 Upgrades 26

3.0 Software28

 3.1 Operating Systems 28

 3.2 Disk Partitioning 29

 3.3 File Systems 30

 3.4 Baselines for Monitoring Servers 32

 3.5 Software for Monitoring Servers 33

4.0 Storage34

 4.1 IDE/ATA Drives..... 35

 4.2 SATA Drives..... 37

 4.3 SCSI Drives 38

 4.4 RAID 41

 4.5 Fibre Channel 42

 4.6 SAN and NAS..... 44

5.0 IT Environment44

 5.1 The Server Room 44

 5.2 Physical Security of the Server Room 46

 5.3 Installing Servers 46

5.3.1 Rack Mounting	46
5.3.2 Power	47
5.3.3 Cabling.....	48
6.0 Disaster Recovery	50
6.1 Backups	52
6.3 Redundancy.....	55
7.0 Troubleshooting	56
7.1 How to Troubleshoot	56
7.2 Hardware Tools.....	57
7.3 Software Tools.....	57
7.4 Performance Bottlenecks	59
Practice Questions.....	62
Answers & Explanations	70

Abstract

This Mega Guide prepares you for CompTIA's Server+ exam. Passing this single exam awards you CompTIA's Server+ certification.

The Server+ exam is a vendor-neutral certification that covers Industry Standard Server Architecture (or ISSA) technology. It is aimed at IT professionals having at least 18 to 24 months of experience.

This Mega Guide is your "Cliffs Notes" key to the test material. It is a short, highly-targeted summary of what's on the exam. It can't cover everything, but if you understand and memorize its contents, you will be very far along towards passing the exam.

What to Know

The major test topics or "knowledge domains" in the Server+ exam – 2009 version -- and their distribution in the test questions are as follows:

- System Hardware 21%
- Software 19%
- Storage 14%
- IT Environment 11%
- Disaster Recovery 11%
- Troubleshooting 24%

From this distribution you can see that System Hardware, Software, and Troubleshooting represent nearly two-thirds of the test questions. Be sure to focus on these areas.

System Hardware includes motherboard and bus architectures, processors, memory, the Basic Input/Output System or BIOS, and the Power-On Self-Test or POST. Be familiar with the different kinds of servers and server topologies. You must also know the basics of fault tolerance, clustering, and multiprocessor systems.

Software topics you should study include how to install and configure Network Operating Systems or NOS's. These include the server versions of Windows, Linux, and Unix. Among the topics are system requirements and Hardware Compatibility Lists or HCL's, file systems, and how to configure new disks. You must know about performance monitoring and optimization and how to establish baselines in order to monitor servers. Finally, you should understand how hardware upgrades impact software configuration.

Storage topics include hard disk types and their interfaces (IDE/ATA, SATA, SCSI, and their variants). Study disk configuration, fibre channel, RAID, and RAID levels. Understand network storage systems and the differences between Storage Area Networks or SAN's and Network Attached Storage or NAS.

IT Environment topics cover the physical aspects of the server room. Among these are issues relating to security, such as limiting access to the server room and the fundamentals of physical security. Also included are environmental issues such as temperature and humidity, protection from power surges, backup generators, and water and flood control. Hardware set-up and configuration includes rack mounting, power supplies, network cabling, adapters, and network interface cards or NIC's.

For **Disaster Recovery**, know how to develop a disaster recovery plan. Know what planning is necessary to successfully implement systems that can be recovered. Understand backup hardware, software and strategies, and server redundancy and fault tolerance.

For **Troubleshooting**, know what diagnostic tools are commonly available and how to use them. These include both hardware and software tools. Understand how to perform problem determination and be able to troubleshoot and resolve problems remotely. Know how to identify and remedy performance bottlenecks.

Tips

- CompTIA recommends 18 to 24 months experience before taking this exam. It is certainly possible for candidates to pass this exam without the recommended experience -- *but only if they study very rigorously and thoroughly.*
- Conversely, one might ask if an experienced person could pass this test without study. Some might, but even well-qualified candidates find that some test topics lie outside their experience. Reading this Mega Guide ensures that you know what topics the exam covers and allows you to gauge if you're really ready to take it.
- All candidates benefit from studying multiple different resources: books, study guides, flash cards, forum threads, and practice exams. Diverse resources are absolutely essential if you lack the recommended 18 to 24 months hands-on experience.
- CompTIA recommends their A+ certification as relevant experience for the Server+ test. While the degree of overlap between the two certifications is debatable, prior testing experience with *any* CompTIA test can only be to your benefit when taking the Server+ exam.
- CompTIA offers a handful of free sample questions at their certification web site. These are only intended to show you what the test questions look like. The sample is far too small to help you learn what you need to know for the exam.
- Taking practice tests is critical to passing. PrepLogic has gained an excellent reputation for practice tests that prepare you well for the real exam. If PrepLogic does not offer Server+ practice tests at the time you will take the exam, search the web for other offerings.
- The test consists of multiple-choice questions with one correct answer. It also contains multiple-choice questions where you select 2 or 3 correct items out of a list of possible answers. For multiple-answer questions, you are told how many choices you should select.
- Be sure to answer *every* question since there is no penalty for guessing. *Never leave a question unanswered!*
- Read each question carefully. You don't want to select an incorrect answer because you did not accurately understand the question.
- Read and consider all answers before picking one. Often the differences between answers are slight and you don't want to pick a wrong answer because you did not read them all. Some questions may have more than one answer that could be considered correct. You must pick the *best* answer, not just an answer that could be correct under the right circumstances.
- When the answer to a question is not immediately apparent, it helps to use a process of elimination. Eliminate answers you know are wrong, and you'll often then be able to take your best guess from the remaining couple alternatives.

- “Word association” is a useful technique in the Server+ exam. From your study you’ll recognize that the words offered in some answers have nothing at all to do with the question. They are trying to confuse you if you don’t know the proper terminology. You can eliminate these errant answers before selecting your final answer.
- Many Server+ questions are factual. You are asked which answer is factually correct. Other questions involve judgment based on your evaluation of a situation presented in the question. Since the judgment questions can be difficult, you greatly increase your chances of a “pass” by correctly answering all the factual questions. Memorization pays off.
- Compared to certification exams from other vendors, the Server+ exam presents you with a lot of questions with little time to answer them. You’ll face 80 questions in 90 minutes. You need to be able to identify answers quickly. Nailing factual questions immediately gives you more time to consider the judgment questions.
- You can always mark a question you’re not sure about for later review.
- It is not unusual to find information in other questions that will help you answer a question on the test.
- The Server+ exam is graded using a *scaling system* that gives you a resultant score between 100 and 900. This precludes tracing your answers back to your score and provides test security. A secret (unpublished) mathematical formula converts your correct questions into the scaled grade. It’s best not to worry about the scaling process as it is not transparent. Instead, concentrate on learning the material and answering the questions correctly.

Changes in the Server+ 2009 Exam:

The Server+ exam was originally introduced in 2001, then revised in 2005. It was revised again in the summer of 2009.

There are two differences between the three exam versions:

- Newer exams test more up-to-date material
- Different exam versions shift the emphasis among the test topics

This Mega Guide covers the 2009 exam. It reflects the 2009 question distribution listed in the above section “What You Should Know To Pass This Test.”

Be sure to visit CompTIA’s [web page on the Server+ exam](#) as they reserve the right to make adjustments or updates to the exam at any time.

While at CompTIA’s web site you can verify:

- The exam objectives
- Available study materials
- Find available courses or classes
- Learn the mechanics of signing up for and taking the test

Be sure that the book you study for the Server+ exam addresses the version of the test you’ll take!

Search on “Server+” at Amazon and you’ll see a dozen books copyrighted prior to 2005 – covering the 2001 Server+ exam – and books copyrighted about 2005 or 2006 – which cover the 2005 version Server+ exam.

Books for the 2009 version exam are appearing in late 2009 and 2010.

1.0 Server Background

A **server** is a computer device that responds to or “services” requests from clients. A **client** is a machine that makes requests of a server.

Software programs can also be considered servers. For example, an Oracle database server could be considered both as a hardware device acting in the role of server, and also as a database software application that acts as a server. In either case, the device or application services queries coming in to it from clients.

1.1 Kinds of Servers

There are many kinds of servers classified by their function with the organization:

General purpose: General-purpose servers typically perform multiple functions and so can not be classified as any specific type of server. An example would be a server in a small company that handles file, print, email, and web proxy functions.

Appliance: Appliance servers are single-function servers built with the idea of being “plug ready” with only very minimal configuration. Appliance servers are **Field-Replaceable Units** or **FRU's** in that if one fails you just plug in another to replace it.

Mail: Mail servers send, receive, and store email. They usually support industry-standard email protocols like:

- **Post Office Protocol** or **POP3** for incoming mail
- **Internet Message Access Protocol** or **IMAP** for incoming mail
- **Simple Mail Transfer Protocol** or **SMTP** for outgoing mail
- **Multipurpose Internet Mail Extensions** or **MIME** for multimedia data

The protocol called **Multipurpose Internet Mail Extensions** or **MIME** adds the ability to send and receive multimedia files as attachments. Both sender and receiver must have MIME capability for the multimedia data transfer to work.

Firewall: Firewall servers can be hardware or software based. They protect an organization’s internal network and servers from outside penetration. They secure the organization from the internet and typically scan incoming and outgoing data traffic. Firewall security can be configured in any number of ways including **packet filtering** (scanning data packets as they come in), and permitting or denying access based on the originating **Internet Protocol** address or **IP** address.

Web: Web servers host web sites. They provide security through such techniques as **user authentication**, **data encryption**, and **web permissions** for access to particular data or folders. **Digital certificates** validate the web server and its supporting organization to internet users. The most popular two software products for setting up web servers are Apache (open source) and **Internet Information Server** or **IIS** (Microsoft).

Terminal: Terminal servers provide remote dial-in access to other servers or a network. To support this access they use serial protocols like:

- **Point-to-Point Protocol** or **PPP**
- **Serial Line IP** or **SLIP**

PPP is more popular than SLIP today because it supports more protocols (TCP/IP, IPX/SPX, etc) and it adds features like data encryption and compression. Adequate security is a big concern with any dial-in facility. There must be one or more ways to **authenticate** users – to determine that they really are who they say they are.

Database: Database servers store data. They usually run relational database products like Oracle, SQL Server, DB2, MySQL or competing software databases to perform this function.

File: File servers provide shared storage and shared access to files. They tie together clients within work groups through file sharing.

Print: Print servers control printers for shared access, thereby avoiding the need for each individual to have their own computer-attached printer.

FAX: Fax servers provide shared telephone fax access, analogous to the same way Print Servers share printers. They are often considered superior to fax machines because they can service remote users, add security, transmit faxes faster, and they don't jam.

News: News servers run software to store and distribute news articles to and from newsgroups either on a specific network or on the internet. The internet's USENET consists of many news servers running **Network News Transfer Protocol** or **NNTP**.

FTP: File Transfer Protocol or **FTP** servers transfer files across the internet. They support **user authentication** for security purposes. Or you can allow users to log in with **anonymous FTP** if you want files to be accessible to the public without user authentication.

SNA: Systems Network Architecture or **SNA** servers enable clients to access IBM mainframes and other IBM-proprietary computers using IBM's SNA networking protocol. SNA protocol is multi-layered and hierarchical. Microsoft's **SNA Server** and its successor product **Host Integration Server** allow personal computer clients to access IBM mainframes and iSeries™ servers in data centers.

NAS: Network Attached Storage or **NAS** servers are an alternative to local file servers. NAS servers use **Network Interface Cards** or **NIC's** to connect to the network and are assigned internet addresses. NAS is very flexible since all kinds of storage devices can use this principle to connect to the network. It's also very easy to take NAS devices on- or off- line. NAS servers can be used by any client or server on the network regardless of operating system.

SAN: A Storage Area Network or **SAN** attaches remote storage devices like disk arrays and tape libraries to servers such that they appear to the server operating systems to be locally attached. **SAN servers** provide the connections to the storage devices used in SAN's.

DNS: Domain Name Servers or **DNS** servers translate high-level English-like names–TCP/IP names–into **Internet Protocol** or **IP addresses**. For example, a server name like www.mycompany.com might translate into a numeric IP address that looks like this: 108.45.93.17. DNS servers replace the old **HOSTS** file that lists static IP address translations in operating systems like Windows server, Unix, and Linux. DNS servers centralize this information and make it easier to update and manage.

- **DNS zones** or boundaries categorize and help organize the DNS information.
- **Primary domain servers** store read/write copies of the IP database while **secondary domain servers** store read-only copies to improve IP query performance. Clients that request DNS address translation are called **resolvers**.
- A **dot server** (or root server) is a DNS server that maintains a database of IP addresses for .com, .net, .edu, etc. Any network that uses Microsoft's Active Directory must use DNS.
- **Dynamic DNS** or **DDNS** keeps a standard domain name like www.mycompany.com linked to dynamically changed IP addresses. It is useful for web servers that need to keep an unchanging domain name while themselves receiving dynamically assigned IP address.

WINS: Windows Internet Naming Service or **WINS** is a Microsoft protocol that translates older NetBIOS names for computers and other network resources into IP addresses. **WINS servers** support this function. Usually there are multiple WINS servers that **replicate** the translation database (in whole or in part) amongst themselves. Replication can occur at **replication intervals** or based on **replication triggers**. It can either be **pull replication** or **push replication**. A **WINS proxy agent** is a server configured to listen to WINS broadcast messages and then communicate with the WINS server to resolve or register NetBIOS names.

RAS: RAS Servers support **Remote Access Service** or **RAS** for remote clients. This allows clients to remotely connect to a local area network. RAS is a term coined by Microsoft and implemented in Windows as a feature or service and so RAS servers often refer to this. Remote clients can access RAS servers by dial-up or by internet connections (such as in a **Virtual Private Network** or **VPN**). VPN's are established using a **tunneling protocol** such as **Point-to-Point Tunneling Protocol** or **PPTP** or **Level 2 Tunneling Protocol** or **L2TP**. When setting up a RAS Server, security is paramount, since anybody on the internet could try to gain access.

Key security protocols you can configure on a RAS Server include:

- **Password Authentication Protocol** or **PAP** (unencrypted or **clear text**)
- **Shiva Password Authentication Protocol** or **SPAP** (encrypted)
- **Challenge Handshake Authentication Protocol** or **CHAP** (encrypted)

DHCP: DHCP servers support the **Dynamic Host Configuration Protocol** or **DHCP**. They dynamically assign temporary IP addresses to clients. For client computers that will get dynamically assigned IP addresses from the DHCP server, you dictate this when configuring their network connections. (The alternative to a dynamically assigned IP address is a permanent **static IP address** assignment.) Clients request a dynamic IP address by a **discovery broadcast** during their boot processing and a DHCP server responds by **leasing** an IP address to the client for a limited time. In Windows server you can administer DHCP by accessing the **DHCP panel** through Administrative Tools.

Proxy: Proxy servers act as intermediaries between clients in the organization and the internet. They often cache or store web pages for faster client access. **Anonymous Proxy Servers** strip out IP addresses to provide client anonymity.

Gateways, Routers, and Bridges: Servers can be used in various network-support roles, for example, as network gateways, routers, and bridges. Dedicated devices are often viewed more favorably in these roles than servers, since they are typically more "plug-and-go" and also designed to be secure in their roles. However for smaller organizations properly configured servers with appropriate software can provide lower-cost solutions.

For setting up servers as routers, there are two common router protocols:

- **Router Information Protocol** or **RIP** – simple but serviceable, a distance vector-based protocol that calculates the best route to a destination based on the least number of hops.
- **Open Shortest Path First** or **OSPF** – more sophisticated w/ more features and better scalability, flexibility, and accuracy than RIP. It calculates routes based on routing tables reflecting the network topology and their current status (rather than simply the number of hops like RIP).

1.2 Server Architectures

Servers can be arranged in various **architectures** or **models** or **topologies**. For example, large companies may organize servers into three levels or **tiers**. **Front-end servers** interface with the world beyond the organization. Examples include mail servers, web servers, proxy servers, and firewalls. **Mid-tier servers** interact with clients located within the organization and include file servers, print servers, and email servers. **Back-end servers** are typically responsible for voluminous data and archival storage. Database servers and data warehouse servers are examples.

Dedicated servers are allocated to a single function or use. For example, if a database server runs only a relational database management system and performs no other function it would be considered dedicated. On the other hand, if that server also performs other tasks like print or file serving, it is a **shared server**.

Distributed servers are geographically dispersed. In a **peer-to-peer server architecture**, all distributed servers are equal. This is easy to set up but may lead to security issues or administrative difficulty. In a **hierarchical server architecture**, some servers are designated as control points for the other servers. This may mean easier administration and better security through a single control point.

When increasing server power and capability you can either **scale up** or **scale out** (sometimes referred to as **building up** or **building out**). Scaling up adds hardware to an existing server to increase its power. Scaling out adds more servers to the network to increase capabilities.

Scaling out means more servers to manage and administer, and could possibly lead to data sharing and security issues. On the positive side, scaling out means greater redundancy among systems and probably increased availability of at least some systems. It is also very easy to scale out if using a “cookie-cutter” approach to setting up new servers.

1.3 Monitoring Protocols

To manage a large group or network of servers, you need software that helps manage the servers based on one of the two predominant protocols.

Simple Network Management Protocol or **SNMP** raises alerts when issues occur and informs server administrators about problems and changes. Notification can be by email, pager alert, log message, or network message to the client the administrator is logged on to.

SNMP components include:

- **SNMP Agents** – software that runs on any server or system being monitored. When an Agent notices an event that causes it to issue a notification it is called a **trap**.
- **SNMP Management system or console** – central system(s) that manage the SNMP software, receive notifications of traps, etc. Typically the **Network Management System** or **NMS** software controls the SNMP Agents by running on the management server.
- **Management Information Base or MIB** – a database that sets out the parameters for monitoring the different network devices and servers.
- **SNMP Community** – a logical collection of systems and servers tied together in a single SNMP monitoring and security structure.

SNMP typically uses TCP port number 162 for the remote Agents to report back to the central NMS. TCP port 161 is used for event communications or traps. Version 3 of SNMP, current as of 2004, uses packet encryption for security. (Earlier versions sent data in unencrypted **clear text**, a potential security risk.)

Many network management products from vendors are based on SNMP underneath the covers. Popular examples include **Tivoli** from IBM, **Unicenter TNG** from Computer Associates, Inc., and **OpenView** from Hewlett-Packard.

Desktop Management Interface or **DMI** is a less popular alternative to SNMP. Since it complements SNMP by carrying more specific device-level information, it is sometimes used in conjunction with SNMP.

DMI stores the device information in its **Management Information File** or **MIF** (which is analogous to SNMP's MIB).

2.0 System Hardware

Along with the Troubleshooting domain, System Hardware is the biggest part of the exam. Even many questions in other topics rely on your knowledge of system hardware. Thus this section is the longest in this Mega Guide.

2.1 Motherboards

A **motherboard** is the “mother of all circuit boards,” the primary circuit card to which all others in the computer connect.

Key motherboard characteristics:

- **Bus** – the data path the motherboard provides for communications across the system.
 - 32 or 64 bits wide. Wider means more bits transferred simultaneously and is generally faster.
 - The bus is extended through the **I/O expansion bus**. Common I/O expansion bus standards for card-based interfaces and external (hardware port) interfaces are:
 - **PCI** or **Peripheral Component Interconnect**. Benefits:
 - **Plug-and-Play** by default for auto-configuration

- ▶ Bus mastering eliminates the **Central Processor Unit** or **CPU** from communications. **Bus mastering** allows devices to communicate across the bus with little to no processor involvement, so it can be faster and also conserve processor resources.
 - 32-bit data path
 - **PCI interrupts** and **PCI steering** mean more addresses are available for **Interrupt Requests** or **IRQ's**
 - **PCI hot swap** means devices can be replaced while the computer remains up
 - **PCI hot plug** means boards can be powered on/off independently, so adapters can be added/removed without powering down the entire server.
 - **Note:** Not all PCI cards are hot capable.
 - **Peer PCI bus** increases expansion slots, offers flexible bus width and speed, and facilitates load balancing
- **PCI-X or PCI Extended.** Benefits:
 - ▶ PCI-X is an extension standard to PCI and is generally physically backward-compatible with cards based on PCI 2.x
 - ▶ Maximum bandwidth varies, ranging from 1024 **MBps** or megabytes per second for version 1, up to 2.15 GB/s or 4.3 GB/s throughput for PCI-X version 2.
 - ▶ Parallel interface
 - ▶ 64-bit data path at faster speeds than PCI
 - ▶ Ten-fold performance increase over PCI
- PCI-Express. Benefits:
 - ▶ Intended to replace PCI, PCI-X, AGP, etc
 - ▶ Uses point-to-point **serial** connections, called **lanes**, between devices & slots
 - ▶ Faster, smaller cables and connectors, many other benefits
 - ▶ Hot swappable and hot plug capable
 - ▶ Data rate speeds up to 250 MB/s for v1.0, 500 MB/s for v2.0, and 1 GB/s for v3.0 per lane. The data transfer rates are 2.5 GT/s for v1.0, 5 GT/s for v2.0, and 8 GT/s for v3.0 **GT/s** stands for billions of transfers per second.
- **AGP** (Accelerated Graphics Port)
 - ▶ For graphics cards for video displays
 - ▶ Not commonly encountered in servers

- ISA and EISA
 - Obsolete standards that preceded PCI

- USB or Universal Serial Bus.
 - An external bus rather than an interface card bus
 - Fast serial port communication designed to replace traditional serial and parallel ports
 - Hot swappable and plug-and-play
 - Provides power to low-power devices
 - A USB hub expands one USB port to several devices
 - USB 1.0 transfers data at 1.5 Mbits/sec, 1.1 at 12 Mbits/sec, 2.0 at 480 Mbits/sec
 - USB devices can present a security risk because they are highly portable, hot swappable, and plug-and-play

- **Firewire**
 - An **external bus** rather than an interface card bus
 - A competing standard to USB
 - Also known as the **IEEE 1394 interface**
 - Hot swappable and plug-and-play
 - Several serial **data transfer rates** or **DTR's** with Firewire 400 topping out at 400 Mbits/sec on the high end for half-duplex transfers, and Firewire 800 topping out around 800 Mbits/sec on the high end for full duplex data transfers. Full duplex data transfers double the DTR of half duplex transfers because **full duplex** means data can be transferred in both directions simultaneously.

- **Clock Frequency** – the number of times per second that a quartz crystal vibrates or “oscillates.” Measured in millions or billions of times per second – **megahertz** or **gigahertz**. Provides for synchronous system operation and helps determine system speed or performance. Processor or CPU **instructions** are executed on the basis of the clock cycle. Higher frequency means better performance.

- **Chipsets** – these subdivide the bus into logical components that run at different clock frequency speeds, thereby avoiding the bottleneck that a single system-wide clock speed would create.

Chipsets create a **hierarchical bus** that places the slower buses beneath the faster ones for maximum performance.

- ▶ Front side bus or FSB -- path to communicate with main memory and graphics card running at motherboard clock speed.
- ▶ PCI bus – 32-bit path I/O for adapter cards, USB, and IDE ports. Connects to system clock and CMOS memory chip.
- ▶ North Bridge chipset – divides FSB (or the “processor bus”) from the PCI bus and manages data traffic in that area. It sets the speed for the FSB and determines how many CPUs and how much memory the machine can have. Sometimes called the system controller chip.
- ▶ South Bridge chipset – divides PCI bus from the ISA bus and Super I/O chip and manages data traffic in that area for slower devices like IDE ports, USB ports, ISA bus, etc. Sometimes called the peripheral bus controller.
- ▶ Accelerated Hub Architecture, now known as Intel Hub Architecture or IHA, was designed to replace the traditional North Bridge/South Bridge design. IHA offers higher-speed channels between sections and it optimizes data transfer based on data type.
- **I2O – Intelligent Input/Output** has an I2O processor on the device itself (such as an adapter card) that communicates with the I2O driver. I2O may still be on the exam but is it largely defunct as an active standard.
 - ▶ Benefits include efficient interrupt handling, hot-plugging, and direct memory access or DMA instead of depending on the processor for memory access
 - ▶ Operating System Module or OSM interfaces to the host OS
 - ▶ Hardware Device Module or HDM handles hardware controller to device I/O

2.2 Processors

Three key factors direct the effective speed and performance of the **Central Processing Unit** or **CPU** – **clock speed**, **data bus width**, and **cache** memory. We just discussed the clock speed above. Let's talk about the other two components here:

- **Data bus width** refers to how many bits can pass into or out of the processor in a single cycle. It is increasingly 64-bit (versus slower 32-bit) for servers. 64-bit also yields a larger memory address space.
- There are two or three levels of **cache** (fast processor memory). Cache speeds performance by holding most recently-used data. From fastest to slowest, and from smallest sized to largest, they are –
 - ▶ **L1 cache** – proximate to the processor and runs at processor's speed rather than the motherboard's speed
 - ▶ **L2 cache** – used to be **discrete** (separate) from the processor, running through the **back side bus**. Since **Advanced Transfer Cache** or **ATC** it is on the processor die like the L1 cache
 - ▶ **L3 cache** – a third level of cache on some systems.
- **Cache memory** also applies to places other than processors in servers. For example, disk drives and CDs or DVDs also have their own cache memory. In this section we are only talking about processor cache memory.

Intel designs and manufacturers both **Pentium** and **Xeon** CPU's. They are instruction-set compatible but Xeon costs more and is targeted at servers due to advantages in:

- Cache size
- Cache speed
- Larger size of addressable memory
- Support for **Advanced Programmable Interrupt Controller** or **APIC**, which enables various devices to communicate with different CPU's via the IRQ's
- **Symmetric Multiprocessing** or **SMP** designs and support (SMP is a way to package more than one processor in a server and is described in section 2.6 below)

Xeon produces more heat than Pentiums and thus is physically distinguished by its type of enclosure because of its cooling needs.

Celeron processors are Pentiums with lesser cache sizes and cheaper supporting chipsets, so Celerons are not used for servers. Though some low-end servers do use Celeron processors they are, by and large, consumer-oriented chips.

Dual core and **multi-core processors** mean more than one execution core on a single die or chip. Their strength is in multi-threaded applications or when several programs run simultaneously. Where tasks can not be separated into multiple threads, or where only a single monolithic program runs at a time, they offer less benefit.

Dual core processors share memory through a single memory controller and appear as one to the outside world via a single system request interface. Designs differ about cache sharing. L1 may not to be shared while lower level caches usually are shared.

Processors can be either 32-bit or 64-bit in respect to two different measurements:

- Data bus width
- Internal operations and instructions

For servers the overall trend is towards 64-bit in both areas. Intel's 64-bit processors are:

- **Itanium**
 - Supports **Explicitly Parallel Instruction Computing** or **EPIC** to simultaneously process many operations
 - Includes **Machine Check Architecture** or **MCA** to identify and catch internal errors
 - Includes L3 cache
 - The instructions it executes -- its instruction set -- is completely different and incompatible with the traditional x86 instruction set
- **Intel 64** (formerly EM64T)
 - This is a fully 64-bit version of the Pentium 4
 - In contrast to the Itanium's new and unique instruction set, the EM64T's instruction set is a superset of the traditional x86 instruction set

Advanced Micro Devices or **AMD** is Intel's main competitor in processor chip design and manufacture. Their **Opteron** series of server CPUs has included dual-core designs since 2003 and quad-core starting in 2007. The distinguishing feature of AMD's multi-core design is its compatibility with the traditional 32-bit x86 architecture. So existing 32-bit operating systems and applications run on AMD 64-bit processors without emulator overhead. Intel's Itanium completely departs from the 32-bit x86 design and is incompatible with it.

It's important to know the advanced features processors have gained over the years:

- **Protected Mode** – processor-level memory protection between programs (programs or program subroutines are usually referred to as **processes** in Unix/Linux or **threads** in Windows)
- **Instruction Pipelining** – overlaps processing of instructions if CPU registers are available
- **Superscalar Architecture** – multiple pipelines can be in progress at once
- **Out of Order Execution** – processor can execute some instructions out of order
- **Branch Prediction** – prefetch loads code on a predictive basis
- **Speculative Execution** – processor predictively executes code, discarding the results if they are not needed
- **Hyper-threading** – Intel's term for simultaneous threading where a single processor appears as two virtual processors to the operating system
- **Explicitly Parallel Instruction Computing** or **EPIC** – as supported by the Itanium and others, allows the processor to carry out up to 20 instructions per clock cycle.
- **On-chip cache** – various levels of cache memory (1st, 2nd, and even 3rd) on die with the processor(s)
- **Dual and quad-core processors** – more than one processor on the same chip die
- **Reduced Instruction Set Computer** – a processor with a small set of instructions or **instruction set** that tries to gain speed through reducing instruction set complexity. The alternative to RISC is **CISC – Complex Instruction Set Computer**. RISC benefits have diminished as chips have become more dense, and while it has many success stories, most computers are based on CISC architectures today (eg: x86 family and others). Hybrids have also blurred the once-clear distinction between RISC and CISC.

2.3 Memory

You need to be familiar with the different kinds of computer memory. Memory is called **Randomly Accessible Memory** or **RAM** because any memory location can be read directly (as opposed to media in which data can only be read sequentially, such as tape). **DRAM** or **Dynamic RAM** is called dynamic because the chips need periodic electrical refresh in order to hold the data. This is as opposed to **Static RAM** or **SRAM** which does not require the refresh and is not as volatile. SRAM is traditionally higher cost and so it is used in specialty memory subsystems. DRAM is the primary form of main memory used in servers and consumer computers.

Most modern memory is DRAM based on **Dual Inline Memory Modules** or **DIMM**'s. DIMMs have separate contacts on both sides of the module that make electrical connection to the memory socket or bank into which it is inserted. This is faster than the earlier **Single Inline Memory Modules** or **SIMMs** that preceded DIMM's. Let's look at the types of DIMM's available. In the order in which they were introduced, from earliest to most current:

- **EDO or Extended Data Out**
 - Early type of DIMM
 - Fast for multiple sequential memory accesses because it eliminates an extra look-up step when accessing multiple data items within the same row

- **SDRAM or Synchronous Dynamic RAM**
 - Runs at clock speed of the system bus
 - Often referred to by the associated bus speed, for example **PC-100** is designed for motherboards operating at **100** mhz

- **RDRAM or Rambus RAM**
 - Developed by Rambus Technology which charges a royalty fee for its use
 - Immediately identifiable because you can't see the memory chips, they are enclosed in a aluminum **heat shield** or heat dissipator that covers the module
 - Data transfer is 16 bits wide from the **Rambus Inline Memory Module** or **RIMM**
 - Transfers data twice per clock cycle. Accomplishes this by data transfer on both the leading and trailing edge of the clock cycle.
 - *All memory slots must be filled for RDRAM to work*, so slots not filled with memory modules are filled with dummy slabs called **continuity modules**. This requirement is unique to RDRAM.
 - DDR SDRAM is often used instead of RDRAM due to the royalty fee adding to RDRAM costs. RDRAM never achieved above 10% market penetration.

- **DDR SDRAM or Double Data Rate SDRAM** (also called **DDR** or **DDR-1**)
 - Improves over and replaces SDRAM
 - Transfers data twice per clock cycle. Accomplishes this by data transfer on both the leading and trailing edge of the clock cycle.
 - Introduces the **prefetch buffer**, a memory cache on the RAM module that stores data before it is actually needed. Initially only 2 bits wide, DDR2 and DDR3 standards expand this prefetch buffer to 4 and then 8 bits..

- **DDR2 (or DDR-2)**
 - Enhanced form of DDR SDRAM with increased pre-fetch, enhanced registers, and on-die termination
 - Newer and faster than DDR and works at higher bus speeds

- **DDR3 (or DDR-3)**
 - Data transfer rate is twice that of DDR-2
 - Newer and faster than DDR-2

Memory stick characteristics (from oldest to newest):

Memory Type:	Connector:	Volts:	Prefetch Buffer Width:	Comments:
SIMM DRAM	30 or 72 pin	5 or 3.3	-	Popular in the P-I and P-II eras
SDRAM	168 pin	3.3	-	Popular in the P-III era
Rambus RAM	184 pin	2.5	varies	Never gained big market share
DDR	184 pin	2.5	2 bit	Popular in the P-IV era
DDR-2	240 pin	1.8	4 bit	Introduced in 2003
DDR-3	240 pin	1.5	8 bit	Introduced in 2007, currently popular

The different “generations” of memory sticks have one or two notches in different locations along their connecting edge so that you can’t put the wrong kind of memory into a memory slot designed for some other form of memory. These edge notches physically prevent you from inserting incorrect memory into a motherboard.

Here are performance characteristics for DDR, DDR-2, and DDR-3:

Name:	Aka:	Memory Clock:	I/O Bus Clock:	Data Transfers / Second
DDR-200	PC-1600	100 mhz	100 mhz	200 million
DDR-266	PC-2100	133 mhz	133 mhz	266 million
DDR-333	PC-2700	166 mhz	166 mhz	333 million
DDR-400	PC-3200	200 mhz	200 mhz	400 million
DDR-500	PC-4000	250 mhz	250 mhz	500 million
DDR2-400	PC2-3200	100 mhz	200 mhz	400 million
DDR2-533	PC2-4200	133 mhz	266 mhz	533 million
DDR2-667	PC2-5300	166 mhz	333 mhz	667 million
DDR2-800	PC2-6400	200 mhz	400 mhz	800 million
DDR2-1066	PC2-8500	266 mhz	533 mhz	1066 million

DDR3-800	PC3-6400	100 mhz	400 mhz	800 million
DDR3-1066	PC3-8500	133 mhz	533 mhz	1066 million
DDR3-1333	PC3-10600	166 mhz	667 mhz	1333 million
DDR3-1600	PC3-12800	200 mhz	800 mhz	1600 million

Registered memory (more current) and **buffered memory** (older technology) re-drive or amplify the signal entering the memory module. They handle heavily loaded server memory and allow modules to have more memory chips with higher reliability. They cost more than the unregistered, unbuffered memory used in consumer computers.

Memory interleaving allows memory access between two or more memory banks or boards to occur alternately. This means faster access because it eliminates wait states. All modules involved must be of the same kind (density and speed).

Interleaved memory must be configured identically across the **banks** (or **boards**) involved. Interleaving is described in terms of the number of banks. For example, with two banks on each of two boards, you have 2 banks times 2 boards or **4-way interleaving**.

Error correcting code or **ECC** memory is a specialized form of SDRAM. ECC recognizes when memory errors occur.

To determine if a memory stick is ECC, add up the number of chips on the stick. If it is evenly divisible by 3, you have ECC memory. Then check the part numbers on the chips. If all are the same, the extra error-detection **check bits** reside on each chip. If one chip has a different part number than the others, then the **parity bits** reside in that **parity chip**.

ECC recognizes memory errors by appending extra bits during memory writes, and then spotting errors using the check bits during read access. It employs a **checksum** to spot errors.

Some forms of ECC can only recognize single-bit errors but can not correct them, while **extended ECC** can both recognize and correct single-bit errors. All kinds of ECC identify but can not correct multiple-bit errors. When unfixable errors are found the response may be to generate a **Non-Maskable Interrupt** or **NMI** and shut down the computer.

Many motherboards will support either ECC or non-ECC memory depending on the motherboard setting. The motherboard configuration (described in the next section below) must be set to **ECC-enabled mode** in most BIOS's to turn on the parity or check bits and the error detection or correction feature. Not many motherboards will support both ECC and non-ECC memory simultaneously, so you typically choose to use either all ECC or all non-ECC memory inside a single server.

ECC requires more bits and is therefore more expensive and very slightly slower than equivalent non-ECC memory. ECC is standard for servers that require high reliability but is usually not worth the cost for consumer computers, where you'll rarely see it.

Memory mirroring is where memory banks mirror each other (duplicate each other's contents). Many servers support a **hot add** of memory where you can add memory while the system remains up and available.

2.4 BIOS, CMOS Memory, and POST

The **Basic Input/Output System** or **BIOS** is software that provides the lowest-level interface between the system hardware and the operating system. The vast majority of BIOS's are from either of two vendors: **Phoenix Software** or **American Megatrends Inc** or **AMI**.

The BIOS provides callable services to programs through its **Application Programming Interface** or **API**. **Device drivers**, operating system programs that provide for use of external devices, use the BIOS API to invoke its services. Usually these are referred to as **system calls**. From this viewpoint the BIOS can be considered a set of small programs that offer low-level services to other programs. Any device you can connect to a server requires a device driver that ultimately invokes BIOS services.

The BIOS is stored on a **flash memory chip**, meaning that you can change or update the BIOS through the flash-memory procedure provided by the computer manufacturer. Flash memory is formally known as **EEPROM** or **Electrically Erasable Programmable Read-Only Memory**. Follow any procedure to "flash the BIOS" carefully and to the letter. Failure might mean destruction of the existing BIOS programs without validly replacing them. This could render your system unbootable.

Along with the BIOS are the CMOS configuration settings. **CMOS** stands for **Complimentary Metal Oxide Semiconductor**. Settings on the CMOS chip are maintained by a small battery that looks like a watch battery when the server is powered down and not receiving electricity. Replace the battery if the system clock is slowing down, as this is a sure sign that the battery is nearing end of life.

You access the CMOS configuration settings for the motherboard by pressing a manufacturer-specific key when the system boots. You then enter a series of panels allowing you to view and change certain configuration settings. Common settings you might see include:

CMOS / BIOS Setting:	Use:
Exit saving changes	Allows you to update CMOS memory settings
Exit without saving changes	Allows you to exit the configuration panels without changing anything
Hardware	View lots of hardware specifics. Change some of them.
Processor	Processor info like stepping (or version) and other characteristics
Memory Test	Allows you to test memory and indicate what level of testing should occur upon boot
Ultra DMA Settings	Configures the UDMA disk controller
PCI Bus Mastering	Enables/disables bus mastering, and sets its characteristics
Memory Scrubbing	Enables/disables memory error correction (ECC)
Management Port	Indicates a serial port that will be used for remote diagnostics
Logging	Determines the extent of system logging. Allows you to view hardware logs.
Security	Often allows you to set passwords to the configuration panels, as well as a server boot password
Boot Order	Determines the order in which devices are selected for booting an operating system
Power Management	Controls power management settings for the system

For servers it is important to protect the existing CMOS/BIOS configuration. If someone has physical access to the server they could easily alter these if they are not protected. CMOS settings are complicated and a well-intentioned but under-qualified person could really mess them up.

Most CMOS configuration systems have a password to prevent unauthorized access. There may be a system boot password as well. Set these and keep them in a secure location so that you don't forget them. If the CMOS is deprived of its battery it loses the configuration information. Thus if you ever forget the passwords you can eliminate them by physically removing the battery from the system for 5 minutes. Or you can set motherboard jumpers for this purpose, often marked **CPW** for **Clear Password** or **RPW** for **Reset Password**.

Unfortunately this process clears more than just the password. It also clears all configuration settings. Most CMOS/BIOS panels have an option to **RESET DEFAULT CONFIGURATION** to help you get back to the standard manufacturer's default configuration if you're forced to.

2.4.1 POST

After briefly giving you the option to enter the CMOS configuration panels, servers enter their **Power-On Self-Test** or **POST** procedure. What is checked depends on the system and its CMOS configuration options, but typically all hardware is verified:

CPU characteristics, BIOS version, memory, video memory and card, keyboard, mouse, ports, disk drives, CD or DVD, sound, and the operating system

Successfully POST leads directly to **booting** -- the loading the operating system. Unsuccessful POST usually stops the boot process with a manufacturer- or machine- specific error code. You'll have to look up this code either in the system documentation or on the web at the manufacturer's web site.

Sometimes you'll see minor warnings or informational messages fly by as booting continues. In this case you can often press the **PAUSE** key to read the message. Or view the CMOS configuration panels' hardware log file.

Here's a simplified view of the boot process:

1. POST runs. It displays an error message and stops if a fatal error occurs, otherwise the boot process continues.
2. CMOS configuration dictates the **boot order** (the order in which devices are selected for booting an operating system).
3. The BIOS **boot loader** program tries to boot from the specified device. For this it switches control to a 512 byte **Master Boot Record** or **MBR** on that device. The MBR is the first 512 bytes of the disk – it exists outside of any partition. It contains code to switch to or **bootstrap** an operating system pointed to by the MBR. The MBR also contains the **partition table** that keeps track of all the partitions on the disk. (Details on disk partitioning are in section 3.2 below.)

2.5 System Resources

Every server system and its motherboard support these limited internal resources:

- **Interrupt Request or IRQ** – devices use IRQ's to interrupt the processor to ask for resources or to tell it they have completed a task. There are 16 IRQ's, numbered 0 through 15. Many are pre-assigned to common devices.

Multiple devices can use the same IRQ number through PCI-based **IRQ steering**. When you have more devices than available IRQ's you can sometimes have a problem of **IRQ contention**. Operating systems sometimes assign **virtual IRQ's** to the real IRQ's to better manage them and avoid IRQ contention.

You can view and change IRQ assignments through the system CMOS panels available upon boot-up. You'd do this when IRQ contention occurs and you have to resolve it by manually re-assigning IRQ's.

If you have to manually re-assign IRQ's, note that the first 8 (IRQ 0 through 7) are serviced by the **master Peripheral Interface Controller** or **PIC** chip, while the next 8 (IRQ 8 through 15) are serviced by a **slave** PIC chip. Only the master directly signals the CPU. The slave always signals the master, which then signals the CPU on its behalf.

Advanced Programmable Interrupt Controller or **APIC** is a more advanced form of peripheral interface controller or PIC. It contains more outputs, a more complex priority scheme, and more advanced IRQ management. Intel's Xeon processors (described in section 2.2 above) are associated with APIC advances.

- **Direct Memory Access** or **DMA** – this resource allows devices to directly access memory without involving the processor. This conserves the processor resource plus provides faster memory access. It has often been used, for example, for video display control. There are eight DMA channels numbered 0 through 7.

The alternative to DMA is **Programmed I/O** mode or **PIO**, which occupies the CPU during the entire I/O operation and makes it unavailable for other work. There are many different **PIO modes**, reflecting improved and faster ways of performing I/O.

- **Memory Address** – main memory consists of equal-sized units referred to as bytes, each of which is given a unique **memory address**. It is the operating system's job to ensure that memory is allocated to only one process or program at a time. This is its **memory protection** feature.
- **I/O Port** – a memory location used for communication amongst devices and processes. Specific I/O ports are usually assigned particular communications roles.

In Windows server, you can view IRQ assignments, IRQ sharing and conflicts, DMA channel use, the I/O memory map, and dedicated memory use through the **Device Manager** or the command line program **msinfo32.exe**. The **System Information** panels provide this data too.

In Linux, the command **cat /proc/interrupts** shows the IRQ assignments and the **procinfo** program gives that plus additional information about memory use, devices, etc. Issue the command **cat /proc/ioports** to see I/O port assignments. The **/proc** directory contains much Linux operational and configuration data.

This chart shows the mapping between IRQ's and their common assignments:

IRQ	Device	IRQ	Device
0	System timer	8	Real-time clock or RTC
1	Keyboard controller	9	Free (cascaded from IRQ 2)
2	Points to IRQ 9	10	Free (often used by sound cards)
3	serial ports 2 or 4	11	Free
4	serial ports 1 or 3	12	Free (often used by mouse)
5	Printer/parallel port LPT2	13	Math co-processor (if any)
6	Floppy disk controller	14	Primary IDE controller (if any)
7	Printer port/parallel LPT1	15	Secondary IDE controller (if any)

2.6 Multiprocessing

Multitasking computers can run more than a single program at one time. Multitasking is an operating system feature. Windows server, Unix, Linux, and NetWare are all multitasking operating systems.

On a single-processor computer, the operating system directs the processor to switch between different tasks or programs at different times, very quickly, thereby giving the impression it is working on several tasks or programs at once. This is multitasking.

Besides fast task-switching, computers can work on multiple tasks or programs simultaneously when they contain **dual-** or **quad-core processors**. We discussed this in section 2.2 above on multi-core processors.

Another possibility is to put more than one processor in one computer. This is **multiprocessing**, which extends and enhances multitasking by placing multiple processors within the same computer. A multiprocessing computer has more than one processor, each of which can multitask as directed by the operating system.

Symmetric Multiprocessing or **SMP** is a multiprocessor design that ties together multiple processors on one server, as supported by SMP motherboard and server design. The multiple processors appear as one computer to the user. Internally all the processors share the same main memory and their use is coordinated by the operating system. Performance benefits depend on...

- Minimized communication and coordination overhead between the processors
- Whether the application programs benefit from **multiple threads**. Many requests coming into a server – Yes. An individual on a standalone computer – Not so much.

SMP advantages are:

- Greater **processor density** and work potential in one machine: without the extra hardware overhead that would be required to do the same work with multiple separate computers
- **Single-system image**: a single operating system runs all the processors and operates in a transparent fashion to users and server support personnel
- Greater **scalability** in a single box
- Effective at handling programming problems that can be divided into many discrete tasks

Due to coordination overhead between the processors, SMP systems typically contain from 2 to 16 processors (although a few manufacturers make systems that extend up to 64 processors).

SMP processors are typically purchased as a group because they must:

1. Be of identical release and have the same **stepping code**. (Using identical processors is a good general principle to follow when you add processors to *any* existing system).
2. Run at same internal clock speed
3. Run at the same front side bus or FSB speed

An alternative to SMP design is **MPP** or **Massively Parallel Processing**. MPP loosely couples the multiple processors within the server in comparison to SMP. In MPP designs, each processor has its own memory (versus the shared processor memory of SMP). MPP designs may either be a single computer with many processors or a coordinated set of computers that together act as an MPP system. The effectiveness of MPP depends on whether a programming problem can be separated into identical parallel tasks. It also depends on the coordination overhead involved in combining results from those parallel tasks.

Since MPP systems share a few hardware resources among their many processors, they are sometimes called **shared nothing systems**. An alternative way to tie together multiple processors and harness them to common tasks is **shared disk systems**, also known as **clustering**. Clustering ties together entirely separate computers, each with its own processor(s), memory, and operating system, but the computers share disks. They have a coordination mechanism to ensure data integrity when writing to the disks. Clusters may extend from two nodes up to some vendor-defined limit.

An example would be an Oracle database cluster. This involves two or more separate computers, each running its own copy of the Oracle relational **Database Management System** or **DBMS** software, and each reading and writing to a single copy of the Oracle database that resides on the shared disks. The Oracle DBMS software is “cluster-aware” so it can handle the problem of coordinating disk updates to ensure data integrity.

Shared disk systems enable you to direct the power of multiple computers at a shared database, but like SMP they suffer the overhead of coordination and communications – in this case, between the multiple computers in the cluster. Their benefits include redundancy and fault tolerance, scalability, high availability and the ability to improve performance incrementally by adding more nodes to the cluster. **Fault tolerance** refers to the ability of a system to remain available after a component fails. It is one of the techniques used to achieve a high percentage of uptime—**high availability**.

Another example clustering system is **Microsoft Cluster Server** software or **MSCS**. This operating system product can tie together separate computers into a high-availability cluster to support essential services to the organization.

Finally we mention **Non-Uniform Memory Access** or **NUMA** computers. NUMA is basically a variation of SMP—there are multiple processors, all of the same type and speed, and they share memory. The difference is that NUMA recognizes different speeds of memory and views memory as hierarchically arranged. NUMA tries to capitalize on this for extra performance – hence the terminology, “non-uniform memory access.” NUMA tries to ensure that no processor is ever waiting on data, and that all are fully utilized all the time, by alleviating possible memory-access bottlenecks.

This chart summarizes the common processor configurations for servers:

Name:	Abbreviation:	Definition:
Uniprocessor	--	A server with a single processor
Dual-core and quad-core processors	--	Either two or four processors on a single CPU chip
Symmetric Multiprocessing	SMP	Usually 2 to 16 processors that share memory and present a single system image (a single operating system image)
Non-Uniform Memory Access	NUMA	An SMP variant that hierarchically arranges memory for non-uniform memory access – leading to greater performance and processor scalability
Massively Parallel Processing	MPP	A shared-nothing architecture that ties together many CPU's or computer systems, yielding the ability to run a very large number of parallel processes or tasks
Clustering	--	A shared-disk architecture through which separate computers share common disk (file systems or database)

2.7 Upgrades

Most servers require upgrades over their life spans. The Server+ exam requires you to know two broad upgrade topics: the proper procedures to follow for upgrades, and some of the specifics of hardware upgrades. This section cover those two topics.

The advance steps to planning an upgrade are:

- Plan on a time to upgrade. This will vary by whether the upgrade will inconvenience users and how long it will take. Target the time of lowest server utilization for the upgrade.
- Notify the users about the upgrade and the chosen **upgrade window** (time when the upgrade will occur)
- Confirm you have necessary upgrade components, check for pre- and co- requisites, verify you have the system resources for the upgrade. Creating an **inventory list** ensures you haven't overlooked anything.
- Test the upgrade via a **pilot program** first to ensure it will work and you know how to do it.
- Perform a **full backup** prior to the upgrade and have a **Backout Plan** in case the upgrade fails. The Backout Plan lists the steps and procedures to follow to get the server back to the state it was in prior to the upgrade attempt, if necessary.

The specific steps you take in replacing hardware components are:

1. Confirm OS and **Hardware Compatibility List** or **HCL** compatibility
2. Prepare yourself for the upgrading by reading any vendor documentation, researching on the internet, reading FAQ's, etc.
3. Document all prior settings (label connectors and cables, record CMOS settings, etc).
4. Make a full backup
5. Inventory supplied parts to ensure everything required is in the upgrade package. Print out any README files and ensure any software is virus-free
6. Perform a **pilot** to test the upgrade on a test server
7. Schedule an **upgrade window** at a time of lowest server utilization
8. Ensure that the UPS is in place and connectivity is good
9. Verify that all required resources are available (examples: IRQ's, DMA, I/O)
10. Once the upgrade is over, verify BIOS detection and new CMOS settings
11. Depending on the kind of upgrade you did, you may have to notify the operating system (example: adding disks). Or you may not have to, if the OS can pick up the changes from the BIOS / CMOS (example: most memory upgrades)
12. Create a new server performance baseline
13. Document what you did and how it turned out

Specifics you should know about hardware upgrades:

- Avoid **Electrostatic Discharge** or **ESD** and other electrical problems by:
 - Touching the chassis
 - Using a grounding kit
 - Unplugging the power
- For processor upgrades be sure you researched compatibility between the motherboard and the **stepping** (version) of the processor you intend to install.
For SMP and MPP upgrades, all processors should be the same stepping and specifically compatible for your equipment as per vendor spec.
- Processor upgrades require pin and socket compatibility, as well as proper cooling. CPU fans should be adhered by grease for good heat transfer.
- Upgrading memory requires proper compatibility, good seating in the bank, BIOS recognition with proper CMOS setting, and thorough testing after the upgrade to ensure the memory is working accurately.

- BIOS upgrades must be done very carefully. If they fail in mid-stream the computer may not boot. If the computer becomes unbootable, many motherboards have a "recovery mode" board jumper. Then you can get the system to try your BIOS boot floppy again to obtain a usable BIOS.
- You can use a **digital multimeter** or **DMM** to test power supplies and their connectors. Hot swappable **power supply units** or **PSU's** can be removed and inserted while the system is still up. Non-hot-swappable units will require powering down and ensuring you've cabled everything back together properly with the new PSU.
- Inserting card adapters properly is simply a matter of matching the slot to your card format. Then ensure the BIOS / CMOS and the operating system both pick up the new card. Sometimes the OS step will require a driver for the card or any device(s) you may have attached to it.

3.0 Software

3.1 Operating Systems

Operating systems control system hardware and resources and provide services to application programs. The functions or services of an **OS** include:

1. Processor control
2. Memory management (including main memory, video memory, caches and buffers, and virtual memory)
3. Device control and software interfaces (or **drivers**) to devices
4. Security
5. The file system

Before trying to install any operating system, check the vendor's **Hardware Compatibility List** or **HCL** to ensure the OS supports all your hardware. If a device is not on the list, that does not mean that it will not work, it just means that it has not been tested and verified as working with that OS.

Here are minimal **system requirements** for various OS's, as published by the vendors themselves. You should always review these on the vendor web sites before installing (or upgrading) any server operating system. Many vendors set these requirements artificially low because they're only talking about the OS and not the applications that run on it. (It's up to you to figure out what resources your applications require.) Marketing reasons are another factor. In the real world you often need to double vendor's minimum recommendations for CPU and memory for reasonable performance. The chart includes "recommended minimum requirements" as well where the vendors provide them:

Operating System:	Minimum / Recommended Requirements:
Windows 2000 Server and Advanced Server Edition	Pentium 133 128 M memory / 256 M 1+ G disk
Windows 2003 Server Standard, Enterprise and Web Editions	Pentium 133 mhz / 550 mhz 128 M memory / 256 M 1.25 to 2 G disk
Windows 2003 Server Data Center Edition	Pentium 400 mhz / 733 mhz 512 M memory / 1 G 2 G disk
Windows 2008 Server	1 ghz for x86, 1.4 ghz for x64 512 M memory 20 G disk for 32-bit systems, 32 G disk for 64-bit systems
Netware 6.5	Pentium II / 2 -way Pentium III, Pentium III Xeon, Pentium 4, or Intel Xeon 700 MHz or higher processors 512 M memory / 1 G 1 G disk / 4 G disk
Novell Open Enterprise Server 2 (this is NetWare's successor product)	Pentium IV or AMD Athlon at 1.5 ghz 768 M memory / 1 G 10 G disk
Unix	Varies by Unix variant and vendor
Linux	386DX / Pentium III for most server distributions No official minimum ram / 512 M for most server versions or distributions

Windows Server 2003 R2 requirements are the same as those listed above for Windows Server 2003.

Operating systems that are designed to run on 64-bit hardware are called **64-bit**. Other operating systems are normally **32-bit**. Operating systems like Windows server and Linux distributions like Red Hat are available either in 32- or 64- bit versions. Check your server hardware to see whether it requires a 32- or 64- bit operating systems before installing one.

3.2 Disk Partitioning

Software helps you to configure hard disks for use by the operating system. **FDISK** is the utility traditionally used to define **disk partitions**, logical portions of a disk. Disk partitions are then formatted with a file system for the operating system's use by the **FORMAT** command. **FDISK** and **FORMAT** were originally MS-DOS commands and there are today several slightly different versions of them under different DOS, Windows and Linux releases.

Key **FDISK** command options allow you to create and work with:

Term:	Meaning:
Primary Partition	A bootable partition on which you install an operating system or OS . Up to 4 Primary Partitions are allowed per disk. Or you can have up to 3 Primary Partitions and an Extended Partition on one disk.
Extended Partition	Contains Logical Partition(s), also known as Logical Drive(s).
Logical Partition or Logical Drive	You can create any number of Logical Partitions – also called Logical Drives – within the Extended Partition.
Bootable Partition	The Bootable flag makes the OS on the partition bootable
Active Partition	The partition that is booted from when the computer starts

Windows Server 2000/2003/2008 comes with the **Disk Management Utility** as part of the **Computer Management Console**. This gives you a graphical interface to view and manage disk partitions and formatting. Most Linux distributions come with graphical tools like **GParted** or **QtParted** that provide full partition management and formatting (even for Windows servers and file systems). These newer Windows and Linux graphical interfaces have largely superseded the old command-line **FDISK** and **FORMAT** commands for most server administrators. Windows Server 2008 and Windows 7 also use **diskpart**.

While Windows requires rebooting after creating logical drives for the OS to recognize them before you can format them, Linux does not require rebooting for immediate use of the new partitions. Just format them and mount them at will. Linux supports Windows and Unix partitions and file formats but Windows does not recognize Linux or Unix partitions.

3.3 File Systems

Operating systems use **file systems** to name, store and organize files on disk. Different OS's use different file systems. Most OS's support a short list of recognized file systems with a particular one usually considered the de facto "standard" or predominant among users.

The important file systems to be familiar with for the exam are:

- **FAT** – originally dates from DOS, has undergone several iterations to increase the size of logical partitions or disks under management and the file size limits. **FAT32** was the last major, popular version. FAT systems lack security, encryption, and native compression in most versions. The file system wastes space at the end of the file (**slack space**) due to large allocation units called **clusters** (set at 32K in earlier versions and down to 4K in FAT32). The largest file size in FAT32 is 4 G and the maximum volume size ranges between 512 M and 8 T depending on the cluster size selected.

- **NTFS – NT File System** is the major, default file system for Windows servers. It has undergone several iterations but the differences between NTFS versions are generally very minor (especially when compared to the big differences among different FAT versions). NTFS has mechanisms to ensure data reliability, built-in security, encryption options, quotas or limits on space allocations to individual users, optional compression, and even a **Distributed File System** or **DFS** to create a single logical directory structure across multiple physical servers. Cluster size can range from 512 bytes to 64K for very efficient use of space. The size of an individual file is limited only by available disk space, while the maximum volume size is huge (in the 16 exabyte range) and is suggested as two terabytes by Microsoft from a practical standpoint.
- **Unix File Systems** – Unix uses a number of similar file systems as favored by particular brands of Unix (Solaris vs AIX vs HP/UX, etc). Examples include AIX's Journaled File System or JFS, Solaris's ZFS and UFS, and many others. These file systems generally support good security through permission bits, large file and volume sizes, and reliability and data integrity features.
- **Linux File Systems** – Most Linux distributions can use any of several popular Linux file systems, which include ext2, ext3, ext4, and rfs. These are similar to the many Unix file systems and include similar features. Most Linux distributions can read and write to both FAT and NTFS file systems. Linux and Unix file systems use **I-nodes** to contain the **directory structure** and file information.
- **Berkeley System Distribution or BSD Unixes** – These systems often use some form of **Unix File System** or **UFS**, as well as other Linux or Unix file systems. Past exams have sometimes referred to UFS as if it were used on all Unix systems so be aware the exam might present UFS as representing all Unix file systems.
- **NetWare File System** – This competes with NTFS and has similar, competitive features. You can add **Novell Storage Service** or **NSS** to an existing NetWare File System install to increase performance and maximum storage capacity.

This chart summarizes the different file systems:

File System:	Native To:	Usable By:	Maximum File Size:	Maximum Volume Size:
FAT32	DOS	DOS, Windows, Linux	4 G	512 M to 8 T *
NTFS	Windows	Windows, Linux	16 EiB	16 EiB
UFS, JFS, ZFS, etc	Various Unix	Various Unix	big and varies	big and varies
UFS	BSD-based Unix	Various Unix	big and varies	big and varies
ext2, ext3, ext4, rfs	Linux	Linux (some Unixes)	big and varies	big and varies

* Maximum volume size varies depending on the cluster size used

Windows, Linux, and Unix use similar file system naming conventions. But Windows uses backward slashes within names while Linux and Unix use forward slashes. Fully-qualified Windows file names include a disk drive letter which Linux and Unix do not use. The root directory for Windows is therefore per logical disk volume – eg: **C:** -- while Linux and Unix does not include a drive letter – so their file system root is simply: **/**. Windows, Linux, and Unix all employ hierarchical file systems with Windows using the term **folder** for what Linux and Unix term a **directory**.

Most Windows, Linux, and Unix file systems include **attribute bits** for each file. This information often includes permission bits for ownership and security, archive bit, read-only bit, system or hidden bit, and sometimes more.

With Windows, Linux, and Unix you should schedule periodic runs of the file verification utility. Under Windows this is the error-checking utility called **ScanDisk**, under Linux and Unix systems it is the file system check or **fsck** command.

Windows file checking utility finds:

- **Lost file fragments** – disk file allocation units that have data but no reference in the file allocation structure (usually caused by disorderly or unexpected system shutdown). ScanDisk fixes this by converting the lost data sectors into files with names like **file0000.chk** and **file0001.chk**.
- **Cross-linked files** – two or more files point to the same disk allocation unit. ScanDisk gives these sectors to the newer file.

Windows' NTFS file system will give better performance if periodically **defragmented** by running the Windows Defrag utility. All Windows disk utilities can be found under the **Disk Properties | Tools** tab or through the **Control Panel | Administrative Tools | Computer Management** (access varies slightly by Windows version). Linux and Unix file systems do not require defragmentation.

3.4 Baselines for Monitoring Servers

To measure whether a server is performing adequately, you first need to establish a baseline. A **baseline** is a point-in-time snapshot of relevant performance information. Baselines typically include measurements for these critical server resources:

- Processor(s)
- Memory
- Disks or disk subsystem
- Network segment

Take a baseline when the server is experiencing a period of maximum sustained activity and behaving normally. Then you can compare this baseline to one you take later, when you believe the system has slowed down or not working up-to-par. Look for changes in resource utilization between the baseline and your "problem-state" system performance snapshot. The differences will point you to the **bottleneck**, the constrained resource that is causing the issue. Now you know what needs to be fixed or adjusted to improve server performance.

To use baselines in problem-solving, you must have created a baseline *prior* to any problems, when the server is performing normally. Otherwise you won't have a normal baseline to compare to when performance problems occur.

A series of baselines or system performance snapshots can also be useful in **capacity planning**. They allow you to view server performance over time and see if any critical resources are stretched to their limits. If so, you know which resources you need more of. And the trend line will indicate the future server capacity you must meet.

3.5 Software for Monitoring Servers

While performance monitoring concepts are applicable across operating systems, the tools you use for this purpose are unique to each operating system. Generally you want to monitor these resources:

Resource:	Look at:
CPU	Percent of CPU usage, called CPU utilization
Memory	How much memory is in use and how much remains free
Caching	Cache size and performance
Paging	How much virtual memory is used and what is the paging rate
Disk Activity	Percent of disk I/O, called disk utilization
Processes	How many processes, and whether any are hung or stopped
Users	How many users, and whether any are consuming inordinate resources and need to have limits or quotas imposed

Unix and Linux offer various command line commands for monitoring resources:

Unix/Linux Line Command:	What It Tells You:
vmstat	Info on memory use, CPU, and interrupts
iostat	Info on disk I/O activity
netstat	Info on network I/O activity
top	Info on processes consuming the most resources
ps	Process listing, with memory and CPU usage
df	Disk usage, space used and available
sar	CPU activity

Unix and Linux also offer GUI performance monitoring tools. The Unix tools are unique to each version of Unix (Solaris, AIX, and HP/UX all have their own, for example.) Linux systems using the GNOME desktop graphical interface use the **GNOME System Monitor**. Linux systems running the KDE graphical desktop interface use the performance monitor called **KDE System Guard** or **KsysGuard**.

System messages are written to different log directories depending on the variety of Unix or Linux, but most write messages to the **/var/log/** directory. Some also write system messages to the **/var/adm/** directory. In most systems the message and system logs are simple text files.

Novell NetWare includes the **Monitor Utility** to track server performance. The **Traffic Manager** can monitor network traffic. The Novell **Internet Caching System** or **ICS** utility is useful for performance monitoring web and FTP services under NetWare. **ConsoleOne** is the basic NetWare graphical interface into a variety of useful tools.

Windows servers have a graphical **Performance Monitor** or **System Monitor** available under **Administrative Tools** or as part of the **Microsoft Management Console** or **MMC**. The key to using the tool is that you need to add **Objects** and **Counters** you want to track. **Objects** are the entity you want to monitor. **Counters** represent the statistics you want to collect about that object. For example, the object might be the processor and the counter its percent utilization. Or the object might be memory and the counter the percent of real memory in use. With objects and counters established, the **Chart View** and tracks them graphically over time.


Windows' **Performance Logs** and **Alerts** are another part of how you track server performance. View these with the **Event Viewer** or through the snap-in to the MMC.

Finally, remember that the Windows **Task Manager** is always accessible. This shows the real-time operation of **Applications** and **Processes**. The **Performance tab** shows CPU and page file usage, as well as other memory statistics. The **Network tab** shows real-time network traffic levels. You can always access

the Task Manager by right-clicking on the bar at the bottom of the display screen.

4.0 Storage

Hard disks consist of a number of **disk platters** that store data, **drive heads** to read and write to the platters, a **spindle motor** to rotate the disk platters, an **actuator mechanism** to move the drive heads over the proper portion of the disk platters. Drives also have **connectors**, **jumpers** and a **disk cache** to store recently-retrieved data from the disk.



On a very basic level disk drive performance is determined by these factors:

- **Seek Time** – how long it takes the drive to move its arm out to the data position
- **Rotational Delay** – how long it takes for the platters to rotate until the data is under the read/write head
- **Throughput** – how many megabytes the drive transfers per second
- **Cache** – cache is fast disk semiconductor memory that can speed performance depending on its size and how frequently data is re-used

Since disk drives involve physical, mechanical movement, reading data from a disk is orders of magnitude slower than reading it from memory. This principle underlies **caching** or **buffering** in operating systems and database management systems. The fastest disk I/O is I/O that is not performed at all... instead the data is read from memory, where the most-recently used disk data has been stored or cached. The secret to speeding disk I/O is to ensure caches or buffers are as large as practical for the application.

4.1 IDE/ATA Drives

One predominant standard for disk drives is called **Integrated Drive Electronics** or **IDE**. It is also known as **AT Attachment** or **ATA**. IDE or ATA drives distinguish themselves from earlier drives by integration of the controller circuitry on the drive. (If an adapter card is used that mainly provides the physical interface for the drive(s) to the motherboard... the controller is actually on the drive itself.)

Key facts about IDE/ATA:

- The **disk controller** is on the drive itself (the name "IDE" derives from this innovation)
- Very high degree of backward compatibility across the IDE/ATA spectrum. You can plug a very old IDE/ATA drive into a new motherboard's IDE/ATA connector and it should work (albeit at slow speeds and with meager storage capacity compared to current drives).
- While the technology interface has been highly standardized, IDE/ATA naming has not. Various generations of drives have multiple names since vendors have managed to popularize their own proprietary naming (as shown in the table below).
- A maximum of two drives can be placed on one IDE/ATA cable, interfacing to one IDE/ATA motherboard or adapter card plug. Only one of the two drives uses the cable for data transfer at any one time.
- CD devices, DVD devices, and tape drives use a variation on the IDE/ATA standard called **ATA Packet Interface** or **ATAPI**.
- IDE/ATA drives are inexpensive but their throughput tops out at 133 **MBps** (or **megabytes per second**)

- Key IDE/ATA terms:
 - **Programmed I/O or PIO** – standards for processor-directed I/O
 - **Direct Memory Access or DMA** – standards for data transfer without processor involvement leading to less processor work and faster data transfer
 - **Self-Monitoring and Reporting Technology or SMART** – drive self-diagnostic capability introduced in ATA-3. Reports drive degradation to software with messages appearing in the system log or a vendor utility log.

This chart summarizes IDE/ATA standards and their performance:

IDE/ATA Standard:	Also Known As:	Performance (in MBps or megabytes per second):
ATA-1		3.3 – 8.3
ATA-2	Enhanced IDE or EIDE, Fast-ATA, Fast-ATA-2, AT Attachment Interface with Extensions	11.1 (PIO 3) – 16.7
ATA-3	Enhanced IDE or EIDE	11.1 – 16.7
ATA-4	ATA-33, Ultra-DMA or UDMA, Ultra-ATA, UDMA/33, Ultra-ATA/33	16.7 - 33
ATA-5	ATA-66, Ultra-DMA/100	44.4 – 66.7
ATA-6	ATA-100, ATA-133	100 – 133

As the chart shows, the names of the standards often tell you their data transfer rates. For example, ATA-33 operates at 33 MBps, while Ultra-DMA/100 runs at 100 MBps. Note also the advances in technology. Newer standards implement faster **Programmed I/O or PIO** modes, and also faster forms of **Direct Memory Access or DMA**. These improvements mean increased data transfer rates across the standards. DMA is generally better and faster than PIO because it eliminates the need for the processor to be involved in data transfers from the disk.

IDE/ATA cables are flat ribbon-style cables that are unshielded, therefore limited to a length of 18 inches or 457.2 mm. Since the interface supports only up to two drives per cable, cables have either one or two connector plugs for drives.

The standard **40-pin connector** on the cable has position 20 blocked, so that the drives are connected properly to the cable. The cable also has a notch to ensure it is plugged into the motherboard (or adapter card) properly as well. Pin 1 on the cable is traditionally denoted by a red stripe which also ensures proper orientation. Pin 1 is on the inner side of the drive, so the red stripe goes towards the drive center when attached.

If you have one drive to connect, place it on the outer-most cable connector. If you have two drives traditionally the master goes on the outer-most cable connector. Physically set the **drive jumpers** on each drive to indicate which is master and which is slave when two drives are connected to the same cable. (With most operating systems the master is assigned the first drive letter and the slave the second drive letter. Some operating systems will not boot off a slave drive while others will).

Cable select allows the master and slave determination to be based on the drives' relative cable positions. It is supposed to be an easier alternative than setting master and slave drive jumpers. *But cable select does not work with all 40-pin cables!*

The ATA-5 standard introduced an **80-pin cable** to replace the older 40-pin standard. The 40-pin cable works with newer standard drives but the newer 80-pin cable is required to achieve the maximum data transfer rate. **Cable select** always works with 80-pin cables whereas it does not work with all 40-pin cables.

4.2 SATA Drives

The IDE/ATA drive standard is now often referred to as **Parallel ATA** or **PATA** to distinguish it from its successor, **Serial ATA** or **SATA**. SATA drives had all but replaced IDE/ATA drives in both consumer and server computers as of 2009.

SATA drives are physically incompatible with PATA drives. SATA uses a 7-in connector cable to support its point-to-point serial communications. The narrower SATA cable makes the physical aspects of connecting drives much easier in small computer enclosures and takes less room. It also makes for better airflow inside server cabinets or racks. The cables are only 8mm wide and can be up to a meter long versus PATA's 18 inches because the SATA cable includes grounding pins that reduce interference. Like a USB port, the SATA interface can be used as an external port for outside-the-box peripherals like external drives.

While SATA drives, cables, power plugs, and connectors are physically incompatible with PATA, SATA provides backwards software compatibility with the older systems. This protects older software and device drivers at the cost of performance where compatibility is used. Here are SATA performance rates:

SATA Standard	Also Known As	Performance (in MBps)
SATA		150
SATA II	SATA-2, SATA2	300

SATA II's advantages over the original SATA standards include:

- **Port Multipliers** – allow up to 15 drives to connect to one interface
- **Port Selectors** – two hosts can connect to one drive
- **Native Command Queuing** – drive commands are executed in the most efficient order rather than simply by first-come-first-serve

The old IDE/ATA standards provided for inexpensive disk drives with reasonable performance, simple configuration, and excellent backwards compatibility. The major advantages of SATA over PATA include:

- Smaller, more flexible cabling, with longer cable lengths due to grounding pins and reduced interference
- Much faster data transfer rates due to fast serial point-to-point communications
- **Fully duplexed** – can send and receive data simultaneously
- **Hot swappable** – can add or remove devices while operating
- SATA II adds advantages like more disk drives per interface and the connection of drives to two hosts (**port multipliers** and **port selectors**)

4.3 SCSI Drives

Traditionally **Small Computer Systems Interface** or **SCSI** (pronounced “scuzzy”) provided the main alternative to the IDE/ATA standards. Generally SCSI systems performed better but at the price of greater cost and complexity. So SCSI is often found in servers while consumer systems almost exclusively used IDE/ATA (and now SATA).

The Server+ exam traditionally expects you to know some of the characteristics of the different SCSI standards. This is difficult because almost every source you can consult will offer a different chart with inconsistent naming. The chart below therefore includes all alternate names for each entry as well as the standards specification on which it is based:

Popular Name	Also Known As	Specification Document	Performance (Clock / Width / Bandwidth)
Synchronous or just “SCSI”	SCSI-1, Narrow SCSI, SCSI-1 Fast-5	SCSI-1 (1986)	5 mhz / 8 bits / 5 MBps
Fast	SCSI-2 Fast-10	SCSI-2 (1994)	10 mhz / 8 bits / 10 MBps
Wide	Wide SCSI	SCSI-2	5 mhz / 16 bits / 10 MBps
Fast-Wide	SCSI-2 Fast-10/Wide	SCSI-2; SCSI-3 SPI (1996)	10 mhz / 16 bits / 20 MBps
Ultra	Fast-20	SCSI-3 SPI	20 mhz / 8 bits / 20 MBps
Ultra Wide	Fast-20 Wide	SCSI-3 SPI	20 mhz / 16 bits / 40 MBps
Ultra2	Fast-40	SCSI-3 SPI-2 (1997)	40 mhz / 8 bits / 40 MBps
Ultra2 Wide	Fast-40/Wide	SCSI-3 SPI-2	40 mhz / 16 bits / 80 MBps
Ultra3	Ultra-160, Fast-80 Wide, Fast-80DT	SCSI-3 SPI-3 (1999)	40 mhz / 16 bits / 160 MBps
Ultra-320	Ultra-4 SCSI, Fast-160 SCSI, Fast-160DT	SCSI-3 (2002)	80 mhz / 16 bits / 320 MBps
Ultra-640	Ultra-5	SCSI-3 (2003)	160 mhz / 16 bits / 640 MBps

SCSI cables and connectors have varied over the years with many different varieties. Especially common are:

- 50-pin or **A Cable**: for 8 bit SCSI
- 68-pin or **P Cable**: for 16 bit SCSI
- 80-pin cable: for 16 bit SCSI

SCSI employs several different methods of **signaling**—transmitting data through electrical impulses:

- **Single-ended** or **SE** – original signaling method of SCSI-1 based on a binary system whereby positive voltage represents 1 and zero voltage represents 0
- **Disconnect-reconnect** – allows devices to dynamically disconnect and reconnect to the drive string when delays are experienced
- **High Voltage Differential** or **HVD** – based on the same signaling principle as SE but implemented with higher voltage for a longer disk string. HVD cable ranges up to 82 feet or 25 meters. You can **not** put SE devices on an HVD string or they could burn out.
- **Low Voltage Differential** or **LVD** – like HVD but based on low voltage (3 volts instead of the older 5 volt levels). This allows cables up to 39 feet or 12 meters long. LVD is a popular signaling method and is the only method used for Ultra3 onwards. SE and **multimode SE/LVD disks** can be integrated on LVD chains with some restrictions.

SCSI connectors do **not** tell you which signal method is used. Therefore these equipment symbols describe the signaling method:



SE



HVD



LVD



LVD/SE multimode

Strings of SCSI drives or cables must be terminated or else they become unreliable and may not even work at all. A **terminator** ends the SCSI chain and prevents any **signal bounce** from the end of the string back down the chain. You must match the kind of SCSI terminator you use with the signal mode of the string. Here are the kinds of terminators:

- **Passive termination** – uses resistors and is least reliable and not widely used
- **Active termination** – adds voltage regulators for more reliable termination
- **Forced Perfect Termination** or **FPT** – ensures precise voltage through diode clamps. Very reliable and the most popular terminator for SE buses.
- **HVD termination** – required for HVD chains
- **LVD termination** – required for LVD strings. Often SE/LVD multimode terminators can handle chains of SE devices by acting as an active terminator.

Passthrough termination allows you to terminate the chain at the point at which the last device is connected.

SCSI drive cables with devices are often called **chains** because the drives are connected or “**daisy chained**” one after another on the cable. 8-bit SCSI standards normally allow 7 connected devices beyond the adapter while 16-bit SCSI usually allows 15 devices plus the adapter.

Each device in the chain must have a unique **SCSI ID** assigned within that chain. The adapter is usually preset at 7 (which is highest priority of all SCSI ID assignments). You assign SCSI ID's through jumpers, wheels, or buttons on the devices. Device ID's determine the order in which devices on the chain are seen and prioritized (physical location within the chain is irrelevant).

Sometimes **logical unit numbers** or **LUN's** are used to identify sub-units within a device. Not all host adapters support LUN's, so if you need this feature you need to verify whether the host adapter supports it.

SCSI expanders extend SCSI capabilities. For example, they enable more devices per chain, allow mixing of certain interfaces, allow longer cable lengths and the like. Their exact capabilities depend on the manufacturer and the product.

To summarize, SCSI is similar to IDE/ATA drives in that both:

- Have undergone several generations of standards but still feature good backwards compatibility for older drives
- Achieve backwards compatibility by "stepping down" performance to that of the older device
- Have the disk controller circuitry residing with the disk itself

The traditional advantages of SCSI over IDE/ATA disks are:

- Better performance with higher data transfer rates including higher sustained data transfer rates
- Ability to string more drives off of a single adapter (versus the IDE/ATA limit of two drives per interface adapter).
- Maximum SCSI cable length is longer.
- Many SCSI drives are manufactured to higher reliability and performance specifications since they are targeted for the server market rather than the consumer-oriented IDE/ATA drives. SCSI drives generally have higher **Mean Time Between Failure** or **MTBF**.
- The SCSI parallel bus allows simultaneous I/O
- SCSI's **Single Connector Attachment** or **SCA-2** connectors allow hot-swapping of disks.

The downside to SCSI versus IDE/ATA is its greater cost and complexity of configuration. Different vendors and products have different configuration rules you have to be aware of when installing SCSI disks.

Internet SCSI or **iSCSI** is a communications protocol that enables a computer to communicate with SCSI storage devices over an IP network. SCSI commands are put into IP packets and sent over the internet or Ethernet to the SCSI controller, which then executes the commands and sends the data back into packets to the originating computer.

4.4 RAID

Redundant Array of Inexpensive Disks or **RAID** transforms individual disk drives into a single disk subsystem to provide greater reliability and/or performance.

RAID can be implemented in hardware, software, or a combination of the two in an independent RAID subsystem. **Software RAID** is supported by operating systems such as Windows server (through its Disk Manager). Software RAID is easy to set up and monitor and is also inexpensive, since it is a built-in feature of the operating system. But software RAID performs very poorly compared to RAID that includes specialized hardware and software. Dedicated RAID systems are also typically more reliable and secure than operating system-based RAID. Also, software RAID does not support all RAID options or levels.

The best-performing RAID is often based on a hardware/software combination shipped by vendors as an intelligent data storage system. These subsystems usually include a large **RAID cache**, semiconductor memory that retains recently-used data for quick reuse of that data without additional disk I/O. The size of the RAID cache is critical to determining the performance of a RAID subsystem. Often performance increases directly with the price of intelligent RAID subsystems due to the value-added of large, expensive caches.

RAID capabilities are determined in part by the type or RAID or **RAID level**. RAID 0+1 and RAID 5 are probably the most common. Here are the RAID levels and how they operate:

Level	What It Is
0	Disk striping. Data is transparently striped across two or more physical drives. Advantage: performance. Disadvantage: one failing drive loses all the data due to the striping, which effectively decreases the overall Mean Time Between Failure or MTBF . Result: few sites use RAID 0. RAID 0+1 is a very popular alternative.
1	Disk mirroring. Controllers simultaneous write any data to two disks. All drives should have the same storage capacity. Advantage: redundancy. Disadvantage: every disk is mirrored so you need double the number of disks.
0+1	Disk striping and mirroring. Data is striped across two or more disks by the first channel and data is mirrored to a second disk via the second channel. Advantage: the performance of RAID 0 with the redundancy of RAID 1. Disadvantage: just like RAID 1, requires double the number of disks (and controllers). RAID 0+1 is very popular.
1+0	Creates a mirror of a RAID 0 array. Most sites use RAID 0+1 instead as it is more flexible with mirror smaller units mirrored than the entire RAID 0 disk array.
3	Striping with parity disk. Stripes data like RAID 0 and adds a parity disk . Advantage: can recover data if one disk fails. Disadvantage: requires an extra disk, can not recover data if two drives fail, also vulnerable to the failure of the single parity disk.
5	Striping with parity. Writes parity data like RAID 3 for data reliability but the parity data is written across multiple disks in a round-robin fashion. Advantage: can recover data like RAID 3 after a disk failure yet does not have a single vulnerable parity disk. Disadvantage: can be slower due to parity-writing overhead. Also takes a long time to rebuild data to a new replacement drive. Requires a minimum of at least 3 drives.

5+1	Mirrors the entire RAID 5 striped sets. This is a high cost solution that may be suitable for absolutely critical data requiring high availability.
0+5	A RAID 5 array composed of striped RAID 0 sets
5+0	A RAID 0 array striped across RAID 5 elements.
Others	You'll sometimes hear of other RAID levels but will rarely encounter them. Some are "standards" invented by individual storage vendors.

Some sources distinguish between disk mirroring and disk duplexing. **Disk mirroring** writes identical data on more than one disk drive, while **disk duplexing** adds a separate controller or adapter for each of the disks. Duplexing is just like mirroring except that it eliminates the interface as a single point of failure (and increases costs by duplicating it).

Note that the **dual** or **nested RAID levels** like 5+1 or 0+1 are sometimes referred to without the plus sign. For example you might see **RAID 5+1** written as **RAID 51**.

Intel developed the **RADIOS** or **RAID I/O Steering** as a specification for a RAID I/O controller embedded on a motherboard.

Zero Channel RAID or **ZCR** is a PCI RAID adapter card that uses the internal SCSI channels of the motherboard to implement low-cost RAID. The card must be installed in a PCI slot and the SCSI scanning option in the BIOS must be disabled for ZCR to work. ZCR is also known as **Modular RAID on the Motherboard** or **MROMB**.

4.5 Fibre Channel

Fibre Channel or **FC** is a set of standards designed for the use of fiber optic cable (although they can also be used for twisted pair copper wire media). **Fiber optic cable** communicates through pulsating light rather than by electrical impulses.

Fibre Channel Protocol or **FCP** is a communications protocol that transports SCSI commands over fibre channel networks.

Fibre channel **ports** are any entity that communicates over the network. There are three major fibre channel topologies that describe how ports are connected:

- **Point-to-Point** or **FC/P2P** – direct connection between two devices. Only two ports are involved.
- **Arbitrated Loop** or **FC/AL** – all devices are connected in a loop or ring. Failure of one device breaks the ring. Fibre channel **hubs** can be used to connect multiple devices together and by-pass failed ports. FC/AL supports a maximum of 127 ports. With its greater number of devices and much longer cable runs, many view FC/AL as the fibre channel replacement for traditional SCSI.
- **Switched Fabric** or **FC/SW** – all devices or looped groups of devices are connected to fibre channel **switches**. The switches dynamically manage communications paths, optimizing them and by-passing any failed devices. FC/SW supports a huge number of ports and yields very high reliability because it by-passes failed components.

Fibre channel is a **layered protocol** consisting of these five layers:

Layer Name:	Layer Function:	Explanation:
F0	Physical	Cables, fiber optics, pin-outs, connectors, etc
F1	Data Link	Coding and decoding of signals – 8b/10b encoding
F2	Network	FC-PI-2 standard for network protocol
F3	Common Services	For extra services (example: encryption)
F4	Protocol Mapping	Maps to supported higher-level protocols, for example, to SCSI commands

This table shows the line-speed and throughput for various FC standards:

Standard:	Line Speed (GBaud):	Throughput (MBps):	Available in:
1GFC	1.0625	200	1997
2GFC	2.125	400	2001
4GFC	4.25	800	2005
8GFC	8.5	1600	2008
10GFC Serial	10.52	2400	2004
16GFC	17	3200	2011 est.
20GFC	21.04	4800	2008

The connect distances supported by FC depend on whether **single-mode** or **multi-mode** fiber optic cable is used. Single-mode tops out at either 10 or 50 kilometers (depending on the variant), while multi-mode extends up to 500 meters for the longest-range variant.

Fibre Channel over IP or **FCIP** is analogous to iSCSI. It allows FC to operate over IP packet networks. The advantage is that this allows for geographically dispersed networks.

FC can be used to support many different storage technologies. Examples include:

- **Just a Bunch of Disks (JBOD)** is a term coined to contrast regular disk drives to RAID
- RAID
- Storage networks of various kinds
- (SAN, NAS, **Serial Storage Architecture** or **SSA**. SSA is a serial transport protocol for attaching disk drives to servers).

Fibre channel is usually considered ideal for such applications as:

- Organization-wide communications networks called **backbones**
- Very large networks
- Multimedia networks supporting digital audio & video
- Server clusters
- Large network-based storage systems

4.6 SAN and NAS

A **Storage Area Network** or **SAN** is an architecture to attach remote storage devices like disk arrays and tape libraries to servers such that they appear to the server operating systems as locally attached. The SAN is connected through a server or cluster of servers that act as client access points.

Network Attached Storage or **NAS** is a storage architecture that gives storage devices their own IP addresses. Clients then communicate with the devices (or their controlling servers) through these IP addresses.

NAS has several advantages. It is usually very easy to set up and attach the NAS storage devices to the network. This means it is easy to load-balance and provide redundancy by adding more NAS devices. Since communication is through universal IP protocols, NAS storage easily serves heterogeneous operating systems and servers on the network.

Both SAN and NAS are usually considered enterprise-wide storage solutions. While costs are coming down one typically thinks of them primarily in large organizations.

5.0 IT Environment

5.1 The Server Room

Planning for servers requires the proper physical environment. Key factors are:

- Temperature
 - Processors and disk drives especially throw off a lot of heat
 - Separate thermostats, proper airflow, and air conditioners are essential
 - Temperature is controlled thru **HVAC** systems (**Heating, Ventilation, and Air Conditioning**)
- Air quality
 - Dust dirties sensitive equipment and a layer of it also increases temperature
 - **Positive pressure** keeps dust outside of equipment thru proper air flow
 - You can clean sensitive electronics and circuit boards using compressed air canisters

- Humidity
 - Should be between 20 and 80 percent non-condensing relative humidity or ideally between 40 and 60 percent
 - **Electrostatic Discharge** or **ESD** damages equipment and can be aggravated by improper humidity control. Controlling ESD is so important that you may place anti-static mats in various locations around the server room or have anti-static covering for the entire area. Even minor electrical “shocks” can reduce reliability for sensitive electrical equipment.

- Flooring and Ceilings
 - A **plenum** provides space for running cable. It is a cavity that can be created by a dropped fake ceiling or in a raised panel floor
 - Plenum floor panels are usually 2 by 2 feet square 11 inches above the concrete. Use a **floor puller** to remove tiles when needed. Tiles are usually coated to minimize ESD and grounding is designed into the plenum.
 - Plenums are designed with both cabling and air flow (HVAC ducts) in mind

- Power
 - **Main power, Uninterruptible Power Supplies** or **UPS**, and **backup generators** provide server room power. Server room power is best dedicated and de-coupled from normal office power systems.
 - UPS is battery-based and provides short-term power if main power is interrupted
 - Backup generators are expensive and run off diesel fuel or natural gas for a longer-term solution when main power goes out
 - **Clean power** means power that is even – not subject to surges, spikes, dips or poor grounding. Clean power prevents short circuiting and tripping circuit breakers.
 - Power cables should be shielded to reduce **Electromagnetic Interference** or **EMI** or else this could affect computer and electronic equipment. Labeling and organized running of cables eases maintenance.
 - Circuit breakers should be labeled. **Circuit breaker finders** can help assure labels are accurate or help identify breakers that are not labeled.

- Labeling
 - Everything in the server room or data center should be labeled, preferably with an easily-understood numbering system. This includes servers, racks, cables, network boxes, HVAC units, you-name-it. Rapid and accurate identification of server room components is required for quick action to fix or replace failing components, for moving items to make space, for logging into the right server to make a change, and for many other purposes. The small amount of time it takes to label items when installed or changed is well worth it given the potential downsides of squandered time and server room mistakes when labels are not applied or applied inaccurately.

5.2 Physical Security of the Server Room

Good physical security is part of server room design. Physical access to the server room can be restricted by several mechanisms:

- Keypads
- Card Scanners
- Biometrics
- Security guards

The most effective systems combine two of the above. For example, you might require a scan of the employee's ID card along with entering a key code to enter the server room.

Server room access should also be monitored and tracked. Ways to accomplish this include:

- Sign-in sheets
- Video surveillance with saved tapes
- Logs kept by access-restriction tools like card scanners and keypads
- Logs kept by servers themselves when they are accessed or reconfigured

5.3 Installing Servers

This section covers the physical aspects of servers and installing them in the server room. The topics -- rack mounting, power, backup power, and cabling -- could also be considered part of the test objective on System Hardware.

5.3.1 Rack Mounting

Rack-mounted servers use vertical space in the server room effectively to reduce the horizontal space required and the overall size of the server room. Racks increase the **density** of equipment in the room to make better use of available space. Servers that are not rack-mounted are usually referred to as **free-standing**.

Blade servers are specifically designed with external dimensions to fit into server racks. This maximizes space utilization through vertical stacking of the servers. Many blade servers are built based on a single integrated circuit motherboard.

Racks are measured in rack units or **U's**, a unit defined by the **Electronic Industries Alliance** or **EIA**. Here is the size of a U as well as some common rack sizes:

Units:	Vertical Inches:	Centimeters:
1U	1.75	4.45
22U	38.5	97.8
24U	42.0	106.7
36U	63.0	160.0
42U	73.5	186.7

U opening width is 19 inches or 48.26 centimeters and exterior rack width is therefore about 23 to 24 inches or 58.5 centimeters. A full rack is usually defined as 42 U in height.

Racks often have **stabilizing feet** because of the weight they can contain. Fill racks from the bottom up to reduce the chance of tipping and ensure the stabilizing feet are in position. Locate the heaviest items in the bottom positions. Racks also may have **ballast** to stabilize them. Pull out only one piece of equipment from a rack at a time to avoid destabilizing racks. Weight distribution within racks may also be part of the planning for the server room floor plan.

Device grouping is a whole topic in itself that considers weight distribution one of several important factors. Cable lengths, clustering needs, location of power outlets and power supplies, cooling and air conditioning all fit into the floor plan or server room layout.

Cable management arms or CMA's can help reduce and manage cable-clutter among the server racks. Cable trays ensure cable plugs don't disconnect. Keyboard-Video-Mouse or KVM switches also help reduce clutter by enabling many servers to use a single keyboard, monitor and mouse. KVM consoles attach to up to 8 servers in typical configurations using the special KVM cabling for connection with the KVM switches.

Proper cooling is essential to rack-mounted equipment. Rack doors are often perforated for ventilation. In addition many racks have mounted fans to draw in cool air and push warm air out. Be sure to plan sufficient space around racks for proper air movement. Don't place boxes or items in front of rack air intakes or outflow vents, or else you've defeated the purpose of properly spacing the racks for airflow.

5.3.2 Power

Power Distribution Units or PDU's are the high-reliability, heavy-duty server equivalent of household power strips. They can be located vertically along racks in the rack bars. Usually they're at the rear of the rack as close to the bottom as possible. Don't overload PDUs by over 80% of their power ratings. Avoid any sudden in-rush of power by switching power on (and off) incrementally.

The server power supplies themselves are rack mounted in U's the same as other rack-mounted server components. You can usually tell which power supplies are hot-swappable because they have handles on them so you can pull them out of the rack.

Estimating how many **Power Supply Units** or PSU's and their wattage is important to designing a server room. If you don't know the overall power requirement of a server you can list its components and derive the power requirement by adding up those individual wattage requirements.

Many servers have redundant hot-swappable power supplies for PSU fail-over. You may hear of "**N + 1**" PSU design, where **N** is the number of PSU's and **1** is the spare for backup. This design allows for replacing a unit while others continue to power the system. Usually you want to keep PSU utilization under 100% for each power supply. At least 6% in reserve is considered reasonable.

While household power outlets are 110 volts or V, server rooms usually have 208/220 V outlets to accommodate their heavier loading. In power calculations remember the standard formula:

$$\text{watts} = \text{volts} * \text{amps}$$

The **load equipment** refers to all the equipment in the server room that requires power. When installing **Uninterruptible power supplies** or UPS, be sure the UPS can handle all the load equipment. The purpose of battery-based UPS is short-term power to protect the server room equipment and its data and software. For a brief main power outage a UPS might be all you need to continue operations. For a longer outage, UPS will at least allow graceful and orderly shutdown of equipment – which can prevent data loss and software corruption.

An easy way to determine UPS power requirements is to take the **volts-amps** or **VA** rating of the UPS and multiply it by **.60** to determine how many watts the UPS will support. Thus for a UPS with a rating of 1000 VA, the watts rating would be calculated as 1000 times **.60** or 600 watts.

Under normal conditions, a UPS operates like this:

1. AC power comes into the UPS
2. The UPS's **rectifier** converts incoming AC power to DC
3. The UPS uses some of the DC power to charge its batteries if necessary
4. The UPS's **inverter** converts remaining DC power to outgoing clean AC power
5. The AC power passes through the transformer to the load equipment

Under a power failure condition, a UPS changes its operation to this:

1. UPS batteries provide power to the inverter
2. The inverter converts the DC power to outgoing clean AC power
3. The AC power passes through the transformer to the load equipment

The power-protection system described above is called an **online/double conversion system**, because a rectifier converts incoming AC power to DC power, then back again to clean outgoing AC power. This is the kind of system used in most server rooms. Alternative systems are:

- **Standby/Offline** – A transistor switches a large transformer, which has a small amount of available power, before transferring power to the UPS battery. This is mainly used for consumer systems.
- **Line Interactive** – The same as Standby/Offline but includes an automatic voltage regulator for power stabilization

5.3.3 Cabling

The most common kinds of cabling in use and cabling standards are:

- **Thicknet** – based on the 10Base5 standard to transmit data at 10Mbps over a maximum distance of 500 meters or 1,640.4 feet. (**Mbps** stands for **megabits per second**. Note that this is a different abbreviation than **MBps**, which stands for **megabytes per second**.) Thicknet is 1 centimeter thick. It's popular for network backbones but widely being replaced by fiber optic media.
- **Thinnet** – based on the 10Base2 standard to transmit data at 10Mbps with maximum segment lengths of 185 meters or 606.9 feet. 0.5 centimeters thick and more flexible and therefore easier to work with than Thicknet. Also known as RG-58 cable. Like Thicknet, it uses **BNC connectors** and its use is fading out.

- **Shielded Twisted Pair or STP** -- this is the long-time standard for carrying telephone signals and is widely used for that purpose. STP transmits data at speeds up to 500 Mbps up to 100 meters or 328.08 feet.
The “twisted pair” is two copper wires, each inside of color-coded insulation, twisted around each other. Enough twists per inch ensures against **crosstalk** or **Electromagnetic Interference** or **EMI**. The STP cable encases a number of twisted pairs which are contained in a shielding to thwart EMI from external sources.

- **Unshielded Twisted Pair or UTP** – UTP transmits data at speeds up to 100Mbps and at lengths up to 100 meters or 328.08 feet. UTP uses filtering and balancing techniques to reduce EMI and crosstalk rather than shielding and it doesn’t offer the same level of protection from those problems as alternatives. But it is lightweight, flexible, and easy to work with, so it remains popular. There are several UTP standards, here are the most common:
 - **Category 3 or Cat 3** – Data transmission up to 10Mbps. Cheaper but not quite as reliable as the much more popular Category 5.
 - **Category 4 or Cat 4** – Data transmission up to 16Mbps. Once common in 16Mbps **Token Ring** networks.
 - **Category 5 or Cat 5** -- Data transmission up to 100Mbps. Popular and used for 10Base-T, 100Base-T, 1000Base-T, and Token Ring networks.
 - **Category 5e or Cat 5e** – Enhanced version of Category 5 cable supporting all the same uses but with better internal and external insulation to reduce EMI at marginal extra cost.
 - **Category 6 or Cat 6** – Data transmission at up to 600Mbps. The heaviest insulation and shielding among the generally-used UTP cable standards.

- **Fiber Optic Cabling – Fiber optics** transmits data by using light pulses within glass or plastic threads. Fiber optics is very different from all the above cabling methods in that it transmits data further and faster. Speeds are up to 2GBps and lengths are up to 25 kilometers or 15.5 miles. Fiber optic does not suffer from EMI issues and is very secure. The downside is that fiber optic cabling is expensive to install, it is fragile compared to wire cabling, and it is difficult to split.
As result of these unique characteristics fiber optic cabling is often used as the **backbone** for large campus or corporate networks.

The most common kind of cable connector for non-optical cable is the **RJ-45** plug. It looks just like a common household RJ-11 telephone connector but is a bit wider. RJ-45 typically has a **straight-through connection** in that the pins and their colors on each RJ-45 plug at both ends of the cable match. RJ-45 plugs have 8 color-coded wires.

Cross-over cables are used to directly connect computers. They are **not** straight-through connections and use a 568A cable end on one end and a 568B cable end on the other end.

Ethernet is a set of standards for connecting computers in **Local Area Networks** or **LAN**'s. Its relevance here is that it is the most common way to connect servers together. Ethernet is standardized as IEEE 802.3.

The network protocol Ethernet uses is called **Carrier Sense Multiple Access with Collision Detection** or **CSMA/CD**. Using this protocol network nodes try to transmit data on the LAN when they believe it not to be in use. If a data transmission conflict or **collision** occurs, a node will back off and wait to re-transmit at a later time.

As a successful standard for local area networks, Ethernet has undergone constant improvement and has dozens of standards. On a broad level these can be classified by their data transmission speeds as:

- **Ethernet** -- 10 Mbps
- **Fast Ethernet**-- 100 Mbps
- **Gigabit Ethernet** or **GbE** -- 1000 Mbps

The speeds above are in **megabits per second**, abbreviated **Mbps**. This is the traditional unit of performance measurement for Ethernet and cabling discussions. It differs from the data transfer rates for disks given in section 4, which were provided in **megabytes per second** or **MBps**. Since it is easy to mistake **Mbps** and **MBps**, some sources will use **Mbits/sec** or **Mbytes/sec** to be more clear, instead of the shorter abbreviations. Here I've gone by the convention you'll likely see on the exam.

This table shows some of the major cable standards for each category of Ethernet:

Category/Speed:	Standard:	Media:
Ethernet:		
	10BASE2	Dominant for many years for 10 Mbps Ethernet
	10BASE5	The original standard, uses coaxial cable
	10BASE-T	Runs over 4 wires (two twisted pairs), often with Cat 3 or Cat 5 cable
	10BASE-FL	Fiber optic based Ethernet
Fast Ethernet:		
	100BASE-T	Generic term for Fast Ethernet over twisted pair cable. Includes 100BASE-TX and 100BASE-T4.
	100BASE-TX	Uses two twisted pair and requires Cat 5 cable.
	100BASE-T4	Uses four pairs in the cable and requires Cat 3 cable
	100BASE-FX	Based on fiber optic cable
Gigabit Ethernet:		
	1000BASE-T	Uses Cat 5e cable
	1000BASE-X	A collection of fiber optic cable standards

6.0 Disaster Recovery

Disaster recovery requires *disaster planning*. If a **Disaster Recovery Plan** or **DRP** has not been created in advance, disaster recovery will be difficult and perhaps even impossible. The DRP needs to be comprehensive and cover all possible contingencies. Otherwise when disaster strikes it may not address the problem you have. You must *test the plan* because inevitably testing will improve the plan by uncovering overlooked issues.

The DRP should be printed and stored in a safe place. Just like data backups themselves, a copy of DRP should also be stored off-site at a secure location. It must also be kept updated as the organization's servers change. The goal is to have a Disaster Recovery Plan ready-to-go on a moment's notice that will address any possible situation the organization could face.

Some of the threats the DRP must address include:

- **Fire** – Server rooms should be designed to prevent, detect, and suppress fire. Federal **Occupational Health and Safety Agency** or **OSHA** regulations must by law be followed. Smoke alarms are key and since halon has been banned from production by the **Environmental Protection Agency** or **EPA** another fire suppressant good with electrical fires must be used.
- **Flood and Water Damage** – Water can come in various forms ... as a flood, as seepage, backed up by improper drainage, and as a result of fire from a wet sprinkler system. Good server room design minimizes these threats.
- **Other forms of physical damage** – Fire and flood always top the list but other disasters can happen, whether caused by nature (tornado, earthquake, hurricane) or by man (vandalism, sabotage, or terrorism).
- **Software or data damage** – While DRP's once only covered all physical contingencies, organizations today realize it must also cover software and data damage or loss. What if a virus destroys the operating systems on the servers? What if an innocent mistake or a malicious employee causes mass data destruction?

For organizations with consolidated server rooms or data centers, it's essential to have an alternate site available in case the primary becomes unavailable or unusable. The alternate site may be a secondary location the organization owns or it might be rented from a **service bureau**. Service bureaus rent the right to use their facilities in case of disaster. They make money by offering this form of insurance to many organizations.

Hot sites are alternate server rooms or data centers that are ready for immediate use. The DRP should specify exactly how the transition to the hot site will occur if it becomes necessary. Obviously equipment and software will have to have some level of compatibility with the stricken data center. The DRP thus covers hardware and software compatibility issues. It also must address personnel – how will your staff get to the hot site? Who is critical to the effort and how will they be contacted? Will you use outside personnel to augment your efforts or just your own?

Cold sites are ready to receive computer equipment. They have the proper infrastructure in place, with the plenum, power, HVAC, etc. But you provide the equipment. The DRP must address how the cold site is populated with the proper hardware and set up for operations. Hot sites typically make for a smoother transition under disaster conditions because they already have hardware and software in place, but cold sites cost less.

While the hot or cold site is in use, the DRP should include problem identification and resolution procedures for the stricken server room or data center. Getting it back into operating condition saves the company money and allows ordinary operations to resume as quickly as possible.

Don't forget **insurance**. DRP's and **business continuity planning** manage very expensive, exceptional occurrences... exactly what insurance is designed to address.

6.1 Backups

A **backup strategy** includes hardware, software, and procedures that ensure the organization's data is securely copied such that it could be restored if something happens to the live data on disk.

Let's start with the hardware. Key factors to consider when buying backup hardware are:

- How much data do you need to back up? What is the storage capacity of the backup device relative to this amount of data and how fast can it back up the data?
- How long would it take to restore your data using this hardware?
- What is the cost of the unit, and what is the cost of the media it uses?
- How long do the media hold information (without rewinding or other intervention)?
- How reliable is the backup media and writing to it? How often do errors occur and are you notified when they do?
- What is the **online retention period** of the device (for how long can the data be restored from a backup without manual intervention)?
- What software drives the unit? It might be bundled with the hardware or perhaps the unit only works with specific operating systems.
- Does the device support features like data compression, encryption, and write verification?
- Can the device run unattended, and if not, how much human intervention and oversight will be needed?

Tapes are the traditional backup media. They are inexpensive, portable, storable, reliable, and reasonably fast for write and read operations. Tapes are also **dense** -- each tape holds a lot of data which means a low cost per stored megabyte. Finally, tapes are a backup standard. Even if you don't have available the drives to read a particular tape or data format other sites or service bureaus do.

Alternatives to tape include optical media like writable CD's and DVD's. These are inexpensive but aren't usually a good solution to backing up large amounts of data quickly. They tend to work better for consumer personal computers than for servers.

Tape drives run the gamut in terms of their interfaces. Many use various forms of SCSI. Some of the largest have proprietary interfaces with vendor-supplied software while very low-end devices use IDE/ATA interfaces.

Tape libraries are self-contained units that include more than one tape in place. Most have **auto-loaders**, which mount and swap tapes automatically. **Tape arrays** use special controllers that stripe data across multiple tape drives in parallel.

Tape robots run the whole tape cartridge show without human intervention, automatically moving, mounting, rewinding, and removing tapes as needed. Some backup systems offer complete **Hierarchical Storage Management** or **HSM**, the ability to intelligently store the most-frequently accessed data on more expensive, quickly-accessible media, while storing less-used data on cheaper, more "offline" media. Some HSM managers automatically move data between various devices and media to most effectively and cost-efficiently use the different media based on data usage.

The larger, more expensive tape units include features to reduce the time human intervention consumes in loading and physically handling the tapes. This also reduces error and the tapes are automatically labeled as well. **Tape operators** used to physically handle tapes – today with these automated units they sit at a console, oversee tape robot logs, and issue the occasional command to the unit through its GUI software. Much of their role is troubleshooting problems that pop up.

There have been many different tape drive technologies and tape standards over time. Here are some of the better-known:

Technology:	Full Name:	Characteristics:
QIC	Quarter Inch Cartridge	A series of common formats popular years ago and obsolete today. Nonetheless, many sites keep one or two QIC tape units on hand to read old tapes if needed.
QIC Wide	Quarter Inch Cartridge Wide	Improved format QIC, holds more data. Also obsolete.
Travan		A family of standards popular for small office systems, backwards compatible with the QIC formats.
DAT	Digital Audio Tape	Tape formats adapted from consumer audiocassettes for server use. Can use helical scanning to increase data density, where writes are performed by rotating angular heads
DDS	Digital Data Storage	A set of standards for encoding data on DAT tapes.
Standard 8MM	Standard 8 Millimeter Tape	8 MM tape was adapted from video technology for servers.
Mammoth 8MM		A set of standards for more storage capacity beyond the original 8MM standards
AIT	Advanced Intelligent Tape	Another 8MM based helical scan technology.
DLT	Digital Linear Tape	½" wide versus ¼" wide for DAT tapes. DLT drives can simultaneously read and write (like some Travan devices), leading to higher data transfer rates. Compared to DAT, in general the DLT formats are faster, higher capacity, more reliable, but also more expensive.
SuperDLT	Super Digital Linear Tape	Successor formats to DLT with greater capacity and data transfer rates
LTO Ultrium	Linear Tape Open Ultrium	LTO is a tape technology from a vendor consortium (HP, IBM, and others). Ultrium is a formatting standard from the group that provides higher reliability, as well as greater data capacities and transfer rates, than DLT technologies.

All tape devices require periodic maintenance, such as cleaning the heads and other components. The more expensive units include self-cleaning capability, as well as error checking and reporting of problems to the controlling console when they occur.

Additionally, your backup strategy should recognize that all digital backup media deteriorate over time. If a government investigation required reading backups from ten years ago, would you be able to? Long-term storage raises additional questions beyond DRP requirements.

6.2 Backup Strategies

You must consider many factors in deciding what data to back up when. Among them:

- What data needs to be backed up, when, and how frequently?
- How disruptive will the backups be to server use? Must servers be down during backups or can backups be "live" while the servers are online? If the latter, do backups impact server performance?
- Do you have a **backup window** when backups will run and in which they must complete?
- What type of backups will be involved?
 - **Full or Normal** – backs up all files
 - **Incremental** – Backs up only files changed since the last full or incremental backup.
 - **Differential** – Backs up only files changed since the last full or incremental backup but does not mark the files as having been backed up (does not clear the **archive bit**)
- How many copies will you make of each backup? It's common to back up in triplicate. One tape goes in the on-site tape repository or **tape vault**, one goes to an off-site tape vault (for disaster recovery), and one resides wherever is most convenient for immediate recovery.
- Which of the backup strategies listed below will you use?

Here are the three most common backup strategies:

Grandfather-Father-Son or GFS –

1. The Son backups -- Label four tapes as "Monday" through "Thursday." Use these tapes for daily incremental backups during the week. Reuse these tapes each week.
2. The Father backups -- Label five tapes "Week 1" through "Week 5." These tapes are used for weekly full backups every Friday, the day you do not perform a Son backup.
3. The Grandfather backups – Label three tapes as "Month 1" through "Month 3." Perform a full backup on the last business day of the month.

At a minimum, GFS requires 12 tapes (assuming each backup fits on a single tape). The GFS strategy provides for tape backups up to a quarter of the year. The GFS system can be modified to fit local needs. For example, add duplicates of each backup tape for off-site storage.

The Six-Cartridge Backup –

This strategy reduces the GFS minimum of 12 tapes down to 6, at the cost of a shorter archival or backup period of only two weeks.

1. Label 6 tapes "Friday 1," "Friday 2," "Monday," "Tuesday," "Wednesday," and "Thursday."
2. Perform the first normal backup on the "Friday 1" tape, and store it off-site.
3. On Monday through Thursday, perform incremental backups on the tapes with these labels and store them on-site.
4. Complete the cycle by performing the second normal backup onto the "Friday 2" tape. Store
5. it off-site. On each successive Friday you will alternate between the tapes labeled "Friday 1" and "Friday 2."

Towers of Hanoi –

This strategy ensures you always have an archive with daily tapes going back 32 days. It only uses five tapes (assuming each backup fits on one tape). On the downside, it's named after the famous Towers of Hanoi puzzle, so it's complicated.

Use five tapes for this strategy, and mark them "A" through "E." The "A" tape will be reused every other day, the "B" tape will be reused every 4 days, the "C" tape will be reused every 8 days, the "D" tape will be reused every 16 days, and the "E" tape will be reused every 32 days. With this in mind, follow this procedure with all normal back ups:

1. On day one, back up to "A." (Reuse the "A" tape every other day.)
2. On day two, back up to "B." (Reuse the "B" tape every 4 days.)
3. On day four, back up to "C." (Reuse the "C" tape every 8 days.)
4. On day eight, back up to "D." (Reuse the "D" tape every 16 days.)
5. On day sixteen, back up to "E." (Reuse the "E" tape every 32 days.)

Whichever backup strategy you use, be sure to label the tapes! This includes both manually labeling them (if they pass through human hands), and labeling them through software headers. And be sure to practice recoveries. You don't want the first recovery you do to be the real thing. If you do your chances for error will be much higher and even if everything goes fine your mean time to recover will be far higher than if you had practiced recovery beforehand.

6.3 Redundancy

Hardware and software redundancy address many availability problems and reduce the chances of having to retrieve backup tapes from an off-site location to perform a recovery. They might even save you from having to put a DRP into action. Certainly redundancy increases availability and uptime. The trade-off is greater cost, as you duplicate various hardware and software resources in order to eliminate single points of failure. The most redundant solutions cost the most and yield the greatest degree of availability.

Hardware redundancy can be **in-chassis redundancy**, duplication of components within a single server, or **system redundancy**, where duplicate servers ensure redundancy. Redundant designs that work best avoid any single point of failure.

Redundancy must be considered for every server hardware resource:

- **Processors – SMP and MPP** systems provide processor redundancy. **Clustering** is another option, where a failing cluster node server can be taken offline while other node(s) remain up and keep the cluster functioning. Dual-node clusters can be **active-active**, in which both servers are always operational in the cluster, or they can be **active-passive**, in which the secondary or passive server is offline and simply waits to take over if the primary fails. This principle extends to multiple node clusters too. Cluster **failover** can be **automatic**, where a secondary machine takes over as primary when the primary fails to respond to a periodic message or when its **heart-beat** (periodic message) goes missing. Or failover can be **manual**, under direction of human operators.
- **Memory** – With memory so closely tied to processors in most computer systems, the same technologies discussed for processors apply here as well. Additionally, some servers have hot-swappable memory cards and cages. You can add in memory cards while the system is up and mark failing memory off-line through software control.

- **Disks – Hot-swappable disks** are considered standard at many data centers. **Hot spares** are drives that are already in place ready to be used. In section 4 on Storage we discussed a number of techniques to manage redundant disks, including various RAID levels:
 - ▶ RAID 1 - Mirroring
 - ▶ RAID 0+1 - Striping plus mirroring
 - ▶ RAID 3 - Striping with parity disk
 - ▶ RAID 5 - Striping plus striped parity
 - ▶ RAID 5+1 - Mirrors entire RAID striped set
- **Networks** – Ensuring servers have multiple **Network Interface Cards** or **NIC's** is key to ensuring that server won't go offline due a to a failed network adapter.
Local area networks should be designed to eliminate possible points of failure. This could duplicate whatever components are possible failure points—cabling, switches, hubs, routers—ultimately culminating in completely redundant LAN's.

As well as hardware resource redundancy, it's essential to have spare parts on hand for quick fixes to hardware problems. Some vendors sell **spare part kits**. A data center of any size keeps many extra parts on hand, everything from disks, to memory, to odd spare parts.

An **inventory list** of what is on hand and where it is kept is essential to effectively using extra hardware to avoid or minimize downtime. Accurately labeling servers and server room components ensures you can quickly identify items needing attention or replacement.

7.0 Troubleshooting

7.1 How to Troubleshoot

This is CompTIA's troubleshooting procedure:

1. **Identify the problem.** Do this by checking documentation and looking for error messages in the logs of the server, its operating system, and any other programs. Ask open-ended questions. You may have to try to re-create the problem to get an error message.
2. **Isolate the cause.** Determine whether the problem is hardware or software related. As when, what, how questions. Identify and question contacts who are involved. Use your senses to observe the problem. If necessary remove one component at a time to isolate causes.
3. **Identify possible solutions.** Create a list of approaches to solving the problem. Try to determine which might work best.
4. **Research the best solution:** Read documentation to see if this problem has occurred before. Look to outside sources (vendor documentation, vendor support, internet forums, internet search) to see if you can verify the proposed solution.
5. **Apply the fix:** Implement and verify the proposed solution. Test it thoroughly, and ensure it has no harmful side effects.
6. **Document results:** Document your findings in an easy-to-search manner.

7.2 Hardware Tools

Troubleshooting tools fall into two categories: hardware and software. Here are key hardware tools:

- **Time Domain Reflectometer** or **TDR** – tests for cable breaks and tells the approximate distance to the break by using the signal's **Velocity of Propagation** or **VOP**. This is one of several kinds of **advanced cable testers**.
- **Fox and Hound** or **tone generator and locator** – tests cable to help you identify it, then you can label it correctly so you don't have to do this again. The "fox" sends a tone down the cable and the "hound" is an amplifier probe that identifies the conductor within a bundle through which the tone was sent.
- **Motherboard diagnostics** – the **POST** tests the motherboard and BIOS perform at each boot contain valuable server hardware diagnostics. The tests include processors, chipsets, memory, interface cards, video, peripherals, etc. Fatal error messages usually appear immediately on the console, or you can view the POST log through the boot configuration panels.
- **POST card** - you can also obtain and insert a special PCI card that will monitor booting and output results on an LED display.
- **Multimeter** - tests the power on any wire, can be used with power supplies.
- **Logic probe** - tests for voltage, can help identify motherboard problems
- **Loopback adapter** - hook the appropriate loopback adapter onto a serial or parallel port to test that the port works
- **Memory testers** - identifies memory by speed, density, and type, and verifies it works. Software solutions are cheaper if they accomplish your goal.
- **Wake-on-LAN** - this hardware/software technology allows remote booting and rebooting of servers. Receiving a **magic packet** is what wakes up or boots the remote server. **WOL** requires that the remote network interface adapter and motherboard support the technology.

7.3 Software Tools

Section 3.5 above listed many software tools for monitoring and managing server performance. These can also be used to identify and diagnose software problems. This section lists additional tools for troubleshooting.

Here are some line commands useful for diagnosing network problems:

Windows Command:	Linux / Unix Command:	Use:
ping	ping	Sends packets to network hosts. You can see if the target is up, time how long packets take to propagate, and test if different packet sizes affect timings.
netstat	netstat	Displays incoming and outgoing network connections, routing tables, and various network statistics
nslookup	nslookup	Queries DNS servers to find DNS details like IP addresses for specific computers

tracert	tracert (or tracpath)	Describes the routing of IP packets over the network
ipconfig	ifconfig	Displays network IP address and other TCP/IP information for your computer
arp	arp	Displays network card and Ethernet connectivity, based on Address Resolution Protocol or ARP . This is the mapping between the IP address and the physical address (often called the Medium Access Control or MAC address).
net, netsh	services file (not a command)	Control network services.
hostname	hostname	Displays host name of the computer.

Section 3.5 listed monitoring tools in various operating systems. Here are some of their diagnostic tools.

For Windows server:

- **Computer Management** – umbrella access to a wide range of tools including the Event Viewer for system, application, and security logs, System Information, Device Management, Services Manager, etc.
- **Task Manager** – real-time performance monitoring and statistics
- **Network Monitor** – packet analyzer and network sniffer
- **Dr. Watson** – program debugging tool
- **Server Resource Kit** – additional OS utilities for performance measuring, tweaking, and diagnostics
- **/SOS switch** – add this to the **boot.ini** file to view drivers as they load
- **Performance Monitor** – add **Objects** and **Counters** to watch and gather statistics
- **System Information** – comprehensive hardware and software information in one place
- **Active Directory** – Domains and Trusts, Users and Computers, Sites and Services

Linux utilities vary a bit based on the distribution. The key determining factor is whether the distribution or **distro** you use is based on the GNOME or KDE graphical user interface:

- GNOME **System Monitor** and KDE System Guard or **KsysGuard** – real-time monitoring and statistics
- GNOME's **Control Center** or KDE **Control Center** – information on applications, device drivers, system settings, configuration, etc
- GNOME's **System Monitor** or KDE **Task Manager** – view and control processes
- **/proc pseudo-file system** – a set of files automatically created from system information about devices, ports, performance statistics, memory information, interrupts, swapping, disk statistics, etc.
- **Tripwire** – this tool detects changed files and directories for all Linux versions.
- **sysctl** -- configuration tool for the kernel, networking, file systems, memory use, etc.

Finally, there's a platform-independent standard called **Intelligent Platform Management Interface** or **IPMI** that defines common interfaces used to monitor system health and generally manage the system. IPMI forms the basis for several products from different vendors.

7.4 Performance Bottlenecks

Often troubleshooting performance issues will uncover a system bottleneck. A **bottleneck** is a key resource which is in short supply, thereby slowing down system-wide performance. As described in section 3.4 above, you can uncover bottlenecks by using your normal performance **baseline** and comparing it to a current performance snapshot. A shortfall indicates a possible bottleneck. Section 3.5 above describes software you use to monitor servers and uncover performance issues and bottlenecks.

Here are resources that could be bottlenecked and ways to resolve their performance issues.

Processors -- **processor utilization** is the percent of the time the processor is active executing user processes or threads (work). How high is too high for processor utilization depends on the workload and the intended use of the server. Some workloads require quicker response than others, and some workload produce spiky or uneven processor utilization. In many systems processor utilization above 65% indicates saturation. Possible solutions include:

- Add processors (for example, to an SMP computer)
- Add servers and split the workload between them (**load balancing**)
- Move CPU-intensive apps to other servers (another form of **load balancing**)
- Implement scalability solutions like clustering
- Turn off unneeded operating system services
- Turn off unneeded features that soak up CPU like encryption or compression
- Add memory (on certain systems, especially Unix and Linux, extra memory can help compensate for a weak processor)
- Tune applications to use less CPU resource (example: tune database programs to require less sorting and CPU-intensive operations)
- Remove software RAID (replace it with hardware RAID if desired)
- Set process priorities properly to ensure best use of available CPU
- Ensure applications and drivers are working properly

Memory – operating systems like Windows server, Unix, Linux, and NetWare logically extend memory through a technique called **virtual memory**. This operating system-based strategy creates a **memory map** consisting of both actual physical or **real memory** and a **paging file** (Windows) or **swap file** (Linux and Unix) on disk that is used as if it were memory. The disk area plus real memory together comprise **virtual memory**. With a large virtual memory the operating system can do more work than if it were constrained by the size of real memory alone.

If real memory is a bottleneck many systems will show **excessive swapping** or **thrashing** – a state in which most of the processor is devoted to the overhead of managing paging rather than doing real work for user processes or threads. Another problem to look for is a **memory leak**, a situation in which memory is not properly reclaimed by the operating system after being used by a user process. A memory leak will sometimes show up in monitoring tools when the system shows that it has less real memory that it actually does (reflecting the loss of memory to the leak).

Possible solutions to memory bottlenecks include:

- Identify and fix any memory leak
- If the system can productively use more virtual memory, increase the size of the **swap space** (Unix and Linux) or **paging file** (Windows server)
- If the paging file or swap space is heavily used you need to add more real memory if possible
- Add memory or memory boards
- Upgrade the motherboard to accept more memory
- Use faster memory
- Ensure applications are using memory effectively
- Distribute memory-hog applications to another server that has more memory (load balance)
- Double-check POST and server logs to ensure all memory is being recognized and used properly. It is possible to insert a memory stick into a motherboard that either does not recognize it or utilizes only half the memory on the stick – without recognizing this situation unless you investigate with memory tools or by looking at the logs.

Disks – How you tune disk bottlenecks often depends on what the disks are used for. Tuning disks on a file server, for example, is very different than tuning disks used by a database server. And tuning disks in a cluster or a RAID subsystem may be entirely different than tuning JBOD disks connected to a single server. Nevertheless, here are some ideas for tuning disk bottlenecks. In general disks shouldn't be busy more than 50% of the time servicing read or write requests. You'll have to determine which solution might apply to your problem based on the use of the server and its disks:

- Reduce the number of disks you place on a single host adapter
- Place more-frequently accessed data on the fastest disks
- Add more disk drives and load-balance the data across them
- Verify that disk access is being properly adjudicated on shared-disk or in server clusters
- Defragment disks
- Consider how RAID might be impacting disk performance
- Archive little-used data to other media, freeing up more disk space
- Check to ensure no drive errors are occurring. Remember that many drives use some form of **Self-Monitoring Analysis and Reporting Technology** or **SMART** technology and will report or log errors.
- If using SCSI remember that set-up can be error-prone. Double-check to ensure the chains are set up properly according to vendor documentation. What worked fine on one server may be totally inappropriate for another.

Network -- Network utilization should not normally exceed 30% or so. Higher values can be problematic for collision-detection/avoidance systems like Ethernet. Here are possible solutions to network bottlenecks:

- **Adapter Teaming** – install two or more network adapters in a server make them appear logically as one with **adapter teaming**. There are two approaches to adapter teaming:
 - **Adaptive Fault Tolerance** or **AFT** – this implements automatic failover to a secondary adapter if the primary fails. AFT usually supports up to four adapter groups or “teams” with two to four adapters per team.
 - **Adaptive Load Balancing** or **ALB** – up to four server adapters works as a team handling the load for a single network address. Balancing is automatic. Also called **asymmetric port aggregation**.
- **Multi-homing** means installing two or more **Network Interface Cards** or **NIC's** in one server and treating each interface as a separate subnet. In contrast to Adaptive Fault Tolerance the adapters are assigned to different IP addresses. In contrast to Adaptive Load Balancing there is no automatic load balancing with multi-homing. The load depends on the traffic coming in.
- Upgrade NIC's to faster equipment. Remember that the overall speed of the network is directly tied to the performance of the NIC's.
- Server placement sometimes causes bottlenecks based on server networks, so simply moving servers sometimes solves the problem.

Practice Questions

Chapter 1

1. Which two of the following statements about PCI-Express are true? Choose two answers.
 - A. It uses a parallel link like PCI but unlike PCI-X.
 - B. It has larger connectors than PCI cards.
 - C. It is intended to replace PCI and PCI-X.
 - D. It enables point-to-point serial communications.
 - E. It is slot-compatible with PCI.

2. You made a mistake on some settings in the CMOS configuration panels at machine start-up, but you don't remember which parameters you set. How can you reset the BIOS/CMOS configuration parameters back to their original values? Select the best answer.
 - A. Assuming you're working with a Windows server, when the system is booting, press F8 and enter Safe Mode. This allows you to reset these values directly by editing the Registry.
 - B. Go into Windows or Linux where you will find a panel for resetting these values under a HARDWARE DEVICES panel.
 - C. Go back into the configuration panels at boot time and guess #try resetting values until you get the correct one(s) reset. You may have to reboot multiple times during this process.
 - D. Go back into the configuration panels at boot time and select the RESET CONFIGURATION TO DEFAULTS option.

3. Match the terms with their proper definitions:

A. SMP	_____	1. Computer with a single CPU
B. Multi-processor	_____	2. Processor with more than one CPU
C. Multi-core processor	_____	3. Appears as a single operating system
D. Uni-processor	_____	4. Multiple equal processors share disk access
E. Single-system image	_____	5. Multiple equal processors with dedicated disks
F. MPP	_____	6. Computer with more than one processor.

Select the best answer.

- A. 1 = D, 2 = C, 3 = E, 4 = A, 5 = F, 6 = B
- B. 1 = D, 2 = C, 3 = E, 4 = F, 5 = A, 6 = B
- C. 1 = E, 2 = C, 3 = D, 4 = A, 5 = F, 6 = B
- D. 1 = D, 2 = B, 3 = E, 4 = A, 5 = F, 6 = C

4. Put the following statements into the correct order to reflect UPS operation:
1. An inverter converts DC power back to clean AC power.
 2. The UPS uses a rectifier to convert AC power to DC power.
 3. Power passes through the transformer to the load equipment.
 4. The UPS receives AC power.
 5. Some DC power is siphoned off to charge the battery.
- Select the best answer.
- A. 1, 2, 3, 4, 5
- B. 5, 4, 3, 2, 1
- C. 4, 2, 5, 1, 3
- D. 4, 2, 5, 3, 1
5. What is bus mastering and why is it beneficial? Select the best answer.
- A. Bus mastering allows the motherboard bus (the bus master) to direct the North and South Bridge buses. Without it, you would not have synchronous operation among these buses.
- B. Bus mastering allows devices to initiate and communicate across the bus with little or no processor involvement, to speed performance and conserve the processor resource.
- C. Bus mastering is the fundamental technology for coordinating operations across processors in SMP and NUMA systems. Without it, you would not have synchronous operation among the processors.
- D. Bus mastering is the fundamental technology for coordinating operations across processors in MPP systems. Without it, you would not have synchronous operation among the processors.

Chapter 2

1. You go into a Linux partition manager like GParted or QtParted because you want to increase the size of a disk partition that is only 2 gigabytes in size. The software informs you that even though you have sufficient space on the disk, it cannot increase the size of the partition you selected. Which of the following could possibly be the problem? Select the best answer.
- A. You cannot increase the size of the partition because GParted and QtParted tools cannot perform this task.
- B. You must increase the size of the file system first, before you increase the partition size.
- C. You cannot increase the size of the partition because it is an old FAT-type partition which is limited to a maximum size of 2 gb.
- D. You cannot increase the size of the partition because you are required to run a CHKDSK or SCANDISK command first.

2. Which of the following performance monitoring tools run under Linux, and which run under Windows?

- ▶ Performance Monitor
- ▶ GNOME System Monitor
- ▶ KDESystem Guard (aka KsysGuard)
- ▶ The Task Manager

Select the best answer.

- A. Windows - Performance Monitor Linux - GNOME System Monitor, KDE System Guard, The Task Manager.
- B. Windows - The Task Manager Linux - GNOME System Monitor, KDE System Guard, Performance Monitor.
- C. Windows - Performance Monitor, The Task Manager Linux - GNOME System Monitor, KDE System Guard.
- D. Windows - Performance Monitor, The Task Manager, GNOME System Monitor Linux - KDE System Guard.

3. Which statement best describes how you use Windows' Performance Monitor (once called the System Monitor)? Select the best answer.

- A. You add counters, the entities you want to monitor. Then you add objects, the values to measure for those entities. The Performance Monitor then creates a log of measurements over time that you view through the Event Viewer.
- B. You add objects, the entities you want to monitor. Then you add counters, the values to measure for those entities. The Performance Monitor then creates a log of measurements over time that you view through the Event Viewer.
- C. You add counters, the entities you want to monitor. Then you add objects, the values to measure for those entities. The Chart View then graphically displays these measurements over time.
- D. You add objects, the entities you want to monitor. Then you add counters, the values to measure for those entities. The Chart View then graphically displays these measurements over time.

4. Which of the following statements accurately describes the system requirements for Windows Server 2008 R2 on x64 servers? Select the best answer.

- A. 1 GHz processor
256 MB memory
20 GB disk
- B. 1.4 GHz processor
512 MB memory
32 GB disk
- C. 2 GHz processor
1 GB memory
60 GB disk
- D. 2 GHz processor
2 GB memory
120 GB disk

5. Match each server with its proper description:

- | | |
|-----------------|--|
| A. DHCP server | _____ 1. Transfers files across the Internet |
| B. Proxy server | _____ 2. A kind of storage server |
| C. NTP server | _____ 3. Provides network time and synchronizes clocks |
| D. RAS server | _____ 4. Dynamically assigns IP addresses to clients |
| E. NAS server | _____ 5. Provides remote access |
| F. FTP server | _____ 6. Intermediary that caches web pages |

Select the best answer.

- A. 1 = F, 2 = E, 3 = C, 4 = B, 5 = D, 6 = A
- B. 1 = F, 2 = E, 3 = D, 4 = B, 5 = C, 6 = A
- C. 1 = B, 2 = E, 3 = C, 4 = A, 5 = D, 6 = F
- D. 1 = F, 2 = E, 3 = C, 4 = A, 5 = D, 6 = B

Chapter 3

1. Which two statements about LUNs are not true? Choose the best two answers.

- A. If two computers access the same device, both must refer to that device by the same LUN.
- B. In current SCSI a LUN is a 64-bit identifier.
- C. LUNs are used with SCSI and iSCSI but not with Fibre Channel.
- D. A LUN is the identifier of a SCSI, iSCSI or Fibre Channel logical unit.
- E. LUN masking is a technique that makes a LUN available to some hosts and unavailable to other hosts.
- F. A LUN is not the only way to identify a logical unit.

2. You've just installed and configured a new server with several SCSI drive chains. You go home for the night, and then get a call at 3 a.m. that the server suddenly stopped communicating with two of the drive strings. What might be the problem? Select the best answer.
- A. You forgot to configure the SCSI drive chains properly in Windows. You need to remotely access Windows and reconfigure the drives.
 - B. The SCSI drives might be improperly terminated.
 - C. You made the mistake of improperly mixing different kinds of SCSI drives on the same chain.
 - D. There might be a power problem with the server. Check the power supply and whether power is at a consistent level.

Chapter 4

1. A brand new application project is in the works. You'll be setting up and configuring six servers for this project. The project is in the early stages and you want to ensure that the servers you install and configure will meet the project team's performance and availability expectations. How can you accomplish this? Select the best answer.
- A. Meet with the project team and develop a baseline.
 - B. Meet with the project team and develop an SLA.
 - C. Meet with the project team and develop an HBA.
 - D. Meet with the project team and give them all the printed material the vendors shipped with the hardware.
 - E. Meet with the project team and ensure your manager knows their manager.
2. Two weekends from now, you'll need to shut down all servers for a particular application team for maintenance. What should you do? Select the best answer.
- A. Send out a broadcast message from the servers to all logged-in users immediately before shutting them down, then send out another broadcast message immediately after you bring the servers back up.
 - B. Call the project team manager to inform him or her immediately before shutting down the servers, then call him or her back to tell them the servers are available again immediately after you bring the servers back up.
 - C. You need take no special action because the users need to realize that server maintenance is sometimes required and they have to learn to work around it.
 - D. You need to meet with the project team and its management well in advance to establish an agreed upon time for the servers to be down. Then use the project team's existing procedures to notify everyone about the window of unavailability for their servers.
 - E. Tell your manager about your plans and let him or her handle it. That's the manager's job.

3. The vendor has certified and made available a major operating system upgrade for your servers. How should you go about applying it to the servers? Select the best answer.
- A. Since the vendor has certified the upgrade, apply it to your production, development, and test servers according to whatever schedule is most convenient to both the users and yourself.
 - B. Apply the upgrade first to test server(s), then to development server(s), then to production server(s).
 - C. Develop an upgrade plan agreed to by your manager and those who use the servers. Typically the plan requires upgrading test server(s) first, then development server(s), then production server(s).
 - D. Since the vendor has certified the upgrade, apply it to all servers at once (production, development, and test) in order to ensure minimal disruption and downtime.
 - E. Since the vendor has certified the upgrade, apply it to all production and development servers at once to ensure minimal disruption and downtime. You can always apply it to less critical test servers later.

Chapter 5

1. What is the difference between a hot site and a cold site? Select the best answer.
- A. A hot site has backup data on hand but no ready-to-use hardware or software, while a cold site has backup data and also the ready-to-use hardware and software.
 - B. A hot site has ready-to-use hardware and software on hand but no backup data, while a cold site has backup data but no ready-to-go hardware or software.
 - C. A hot site has ready-to-use hardware, software, and backup data in place, while a cold site does not have backed up copies of the data or ready-to-use hardware and software.
 - D. A hot site has ready-to-use hardware, software, and backup data in place, while a cold site has the ready-to-use hardware and software but not the backup data.
2. Match the kinds of tape with their formats:
- | | |
|-----------|---|
| A. Travan | ___ 1. 8MM helical scan technology |
| B. DLT | ___ 2. Backward compatible with QIC formats |
| C. DAT | ___ 3. 1/2" wide |
| D. AIT | ___ 4. Standards for DAT tape encoding |
| E. DDS | ___ 5. Based on audiocassette technology |
- Select the best answer.
- A. 1 = D, 2 = A, 3 = B, 4 = E, 5 = C
 - B. 1 = D, 2 = A, 3 = B, 4 = C, 5 = E
 - C. 1 = C, 2 = A, 3 = B, 4 = E, 5 = D
 - D. 1 = D, 2 = B, 3 = A, 4 = C, 5 = E

3. Why should you test your DRP? Select the two best answers.

- A. Testing enables you to identify and remove any problems with the DRP prior to use.
- B. Testing is allocated for in the budget.
- C. Even though you know the plan is good, managers require typically testing.
- D. Testing is the only way to know it works.
- E. Testing shows the project team that disaster planning is important.

Chapter 6

1. One of your servers is experiencing transient memory errors. Of the possible causes listed, which two are most likely? Select the two best answers.

- A. Overheating.
- B. Power loss.
- C. You've filled all the memory slots in the server.
- D. Bad memory.
- E. Incorrect operating system configuration.
- F. High humidity.

2. As more users are added to a server, you find that its performance suddenly slows. You run a new benchmark and compare it to your original baseline, with these results:
BEFORE: AFTER: Processor utilization 32% 38% Disk utilization 7% 10% Paging rate 1100 pps 5575 pps Network utilization 17% 19% Which action might correct the performance problem? Select the best answer.

- A. Add another processor or update the existing processor.
- B. Add more disk.
- C. Add more memory.
- D. Add another NIC and configure the NICs for load balancing.

3. You have a small LAN server that's been sitting in a closet, working reliably for years. Nobody has ever backed it up. One day you learn that it has crashed. You find that this Windows computer appears to start up fine and make it through the POST tests, but it will not boot into Windows. You try Safe Mode, booting into the Recovery Console, the Last Known Good Configuration, and other similar tricks to boot Windows, but to no avail. Windows simply will not boot up at all. Meanwhile, your boss informs you that it is critical to save the data from this computer. Given that the disk drive spins when you start the computer, which two techniques might you try to retrieve the critical data off this IDE/ATA disk? Select the best two answers.
- A. Unless backups exist, there is no way to retrieve the data from the disk drive.
 - B. Remove the disk drive with the data from the LAN server, and place it into another computer as a secondary disk drive (with the jumpers and cable connections properly set). Then boot this second computer and retrieve and save the data from the disk.
 - C. Place a Live Linux CD or DVD into the CD/DVD drive and boot the computer into Linux. While in Linux, perform a full install onto the drive that has the data you want to retrieve. After you've done this, remove the Live Linux CD/DVD, and reboot the computer into the new Linux install. Now you can access the data on the drive you want to save.
 - D. Place a Live Linux CD or DVD into the CD/DVD drive and boot the computer into Linux. While in Linux, mount the drive that has the data on it. Immediately copy and save the data.
 - E. A Live Linux CD/DVD cannot be used to save the data from the Windows computer because Linux cannot access Windows disks. Therefore you must figure out how to boot Windows on the down LAN server. Search the web for advanced techniques to do this since the standard approaches you've already tried have failed.
4. You want to shrink a partition using a Linux partition manager like QtParted or GParted but are unable to. Which of the following might be the reason? Select all answers that apply.
- A. Partition managers like QtParted or GParted cannot shrink partitions.
 - B. The partition is mounted and in use. You must unmount the partition before you can resize it.
 - C. Not every kind of partition can be made smaller by QtParted or GParted (depending on its partition type and resident file system).
 - D. The file system on the partition could have an error. Until the error is corrected, the partition cannot be made smaller.
 - E. The partition manager cannot shrink the partition because it is a Windows NTFS partition.
 - F. The partition manager cannot shrink the partition because it is a Windows FAT partition.

5. Your company used an outside firm to disassemble, package, ship, and re-assemble some servers to a new location. As you check up on one of the servers you find that its SCSI disk chain is not working properly. You notice that the boot disk is set to ID 7 and that the other disks are numbered with consecutive IDs from 6 to 1. The SCSI host adapter is set to ID 0. What should you do? Select the best answer.
- A. Leave all SCSI IDs as is. The problem has to be somewhere else.
 - B. Set the SCSI boot disk to ID 0.
 - C. Set the host adapter to SCSI ID 7.
 - D. Set the SCSI boot disk to ID 0 and the host adapter to SCSI ID 7.

Answers & Explanations

Chapter 1

1. Answers: C, D

Explanation A. Incorrect. A defining characteristic of PCI-Express is that it uses point-to-point serial communications rather than parallel links.

Explanation B. Incorrect. PCI-Express connectors are smaller than traditional PCI connectors.

Explanation C. Correct. PCI-Express was intended by its designers to replace all previous bus extension card standards, including PCI and PCI-X.

Explanation D. Correct. PCI-Express gets great speed by defining point-to-point serial links.

Explanation E. Incorrect. PCI-Express and PCI use different sized slots specifically so that you cannot accidentally place a PCI-Express card into a PCI motherboard slot.

2. Answer: D

Explanation A. Incorrect. The Windows Registry does not control the BIOS/CMOS settings and cannot change them. You need to go back into the configuration panels at boot time and select the RESET CONFIGURATION TO DEFAULTS option. This resets all configuration values to their defaults.

Explanation B. Incorrect. You can not reset BIOS/CMOS values through any operating system. These values are set at a level beneath the operating system (which uses them and depends on them to be correct). You need to go back into the configuration panels at boot time and select the RESET CONFIGURATION TO DEFAULTS option. This resets all configuration values to their defaults.

Explanation C. Incorrect. While this approach could solve the problem, there is a better way. Go back into the configuration panels at boot time and select the RESET CONFIGURATION TO DEFAULTS option. This resets all configuration values to their defaults.

Explanation D. Correct. You can reset all parameters back to their defaults by selecting the BIOS/CMOS panel option that resets all configuration values to their defaults. Note the downside #this will reset all values# including some you might not want to revert to defaults.

3. Answer: A

Explanation A. Correct. You must know the precise definitions to correctly answer this question; otherwise you might mistakenly match a true characteristic to one of the terms instead of its definition. For example: "6. Computer with more than one processor" is true for several of the terms but it is only definitive for "B. Multiprocessor."

Explanation B. Incorrect. This answer confuses SMP with MPP. SMP is Symmetric Multiprocessing, many equal processors that all share disk access. MPP is Massively Parallel Processing, many equal processors that each has their own dedicated disk. Thus item 4 = A and 5 = F (not 4 = F and 5 = A). The correct answer is: 1 = D, 2 = C, 3 = E, 4 = A, 5 = F, 6 = B.

Explanation C. Incorrect. This statement confuses the definitions of "D. Uni-processor" and "E. Single System Image." A single system image is not defined as a "computer with a single CPU" but rather as a system that "appears as a single system image." Thus item 1 = D and 3 = E (not 1 = E and 3 = D). The correct answer is: 1 = D, 2 = C, 3 = E, 4 = A, 5 = F, 6 = B.

Explanation D. Incorrect. This statement confuses the definitions of "B. Multi-processor" and "C. Multi-core processor." The former is any computer that has more than one CPU while the latter is a single die or chip that contains more than one processor core. Thus item 2 = C and 6 = B (not 2 = B and 6 = C). The correct answer is: 1 = D, 2 = C, 3 = E, 4 = A, 5 = F, 6 = B.

4. Answer: C

Explanation A. Incorrect. First the UPS receives AC power, and then it uses a rectifier to convert it to DC power. At this point some power is siphoned off to keep the battery charged. Then the inverter converts DC power back into clean AC power, and this power passes through the transformer to the load equipment. The correct sequence is: 4, 2, 5, 1, 3.

Explanation B. Incorrect. First the UPS receives AC power, and then it uses a rectifier to convert it to DC power. At this point some power is siphoned off to keep the battery charged. Then the inverter converts DC power back into clean AC power, and this power passes through the transformer to the load equipment. The correct sequence is: 4, 2, 5, 1, 3.

Explanation C. Correct. First the UPS receives AC power, and then it uses a rectifier to convert it to DC power. At this point some power is siphoned off to keep the battery charged. Then the inverter converts DC power back into clean AC power, and this power passes through the transformer to the load equipment.

Explanation D. Incorrect. First the UPS receives AC power, and then it uses a rectifier to convert it to DC power. At this point some power is siphoned off to keep the battery charged. Then the inverter converts DC power back into clean AC power, and this power passes through the transformer to the load equipment. The correct sequence is: 4, 2, 5, 1, 3.

5. Answer: B

Explanation A. Incorrect. Bus mastering refers to how components can exclude the processor from their communications and thereby speed up performance.

Explanation B. Correct. Bus mastering eliminates the need to include the processor in initiating all communications and thereby speeds the system and conserves the processor resource.

Explanation C. Incorrect. Bus mastering is not a SMP or NUMA concept, but rather a concept that applies to how devices can communicate across the bus within a single system without processor oversight, thereby speeding communications. Bus mastering eliminates the need to include the processor in initiating all communications and speeds the system and conserves the processor resource.

Explanation D. Incorrect. Bus mastering is not a MPP concept, but rather a concept that applies to how devices can communicate across the bus within a single system without processor oversight, thereby speeding communications. Bus mastering eliminates the need to include the processor in initiating all communications and speeds the system and conserves the processor resource.

Chapter 2**1. Answer: C**

Explanation A. Incorrect. GParted and QtParted are quite capable of increasing partition sizes#this is one of their major uses. Any competitive partition management tool can increase partition sizes.

Explanation B. Incorrect. You normally increase the size of the disk partition, then increase the size of the file system. (Tools like GParted and QtParted only require a single action on your part to accomplish both tasks, but internally this is the sequence of operations they apply.)

Explanation C. Correct. Even if you have sufficient disk space, you still will not be allowed to increase the size of old FAT partitions beyond their maximum allowable sizes. For FAT16 the allowable maximum volume size was 2 gb. For FAT32 the maximum volume size varies but it could possibly be 2 gb. If you can, you may want to change the partition to a more modern definition like NTFS after saving the partition data. Then you can create a much larger partition without the 2 gb partition limitation.

Explanation D. Incorrect. Tools like GParted and QtParted do not require running aCHKDSK or SCANDISK command prior to extending partition size. However, it is always a good practice to ensure the file system of any partition you wish to increase in size is clean before attempting to increase partition size.

2. Answer: C

Explanation A. Incorrect. The Task Manager is a monitoring system that is built into Windows and not available under Linux. The correct answer is therefore:
Windows - Performance Monitor, The Task Manager
Linux - GNOME System Monitor, KDE System Guard

Explanation B. Incorrect. The Performance Monitor is a monitoring system that is built into Windows and not available under Linux. The correct answer is therefore:
Windows - Performance Monitor, The Task Manager
Linux - GNOME System Monitor, KDE System Guard

Explanation C. Correct. This answer correctly matches the system monitoring tools with the operating systems on which they run.

Explanation D. Incorrect. The GNOME System Monitor is a Linux system monitoring tool that runs under the Linux GNOME user interface. It does not run under Windows. The correct answer is therefore:
Windows - Performance Monitor, The Task Manager
Linux - GNOME System Monitor, KDE System Guard

3. Answer: D

Explanation A. Incorrect. This answer confuses counters and objects. Objects are the entities you want to monitor, and counters are the values to measure for those entities. The Performance Monitor then displays the values graphically over time through its Chart View (not through a textual log file you view through the Event Viewer).

Explanation B. Incorrect. The Performance Monitor does not create a log file you view through the Event Viewer; instead, you see the results plotted over time via the ChartView graphical user interface. This answer is correct in its statement concerning objects and counters.

Explanation C. Incorrect. This answer confuses counters and objects. Objects are the entities you want to monitor, and counters are the values to measure for those entities. This answer is correct in stating that you view the results through the Chart View.

Explanation D. Correct. This answer properly describes how you use Windows' Performance Monitor. You add objects, the entities you want to monitor. Then you add counters, the values to measure for those entities. The Chart View then graphically displays these measurements over time.

4. Answer: B

Explanation A. Incorrect. Windows Server 2008 R2 minimally requires a 1.4 GHz processor, 512 MB of memory, and 32 GB disk. The requirements listed in this answer are too low for all three resources.

Explanation B. Correct. Windows Server 2008 R2 minimally requires a 1.4 GHz processor, 512 MB of memory, and 32 GB disk.

Explanation C. Incorrect. Windows Server 2008 R2 minimally requires a 1.4 GHz processor, 512 MB of memory, and 32 GB disk. The requirements listed in this answer Software 29 are too high for all three resources.

Explanation D. Incorrect. Windows Server 2008 R2 minimally requires a 1.4 GHz processor, 512 MB of memory, and 32 GB disk. The requirements listed in this answer are too high for all three resources.

5. Answer: D

Explanation A. Incorrect. This answer confuses the definitions of DHCP and Proxy servers. That is, "4. Dynamically assigns IP addresses to clients" should be "A. DHCPserver," and "6. Intermediary that caches web pages" should be "B. Proxy server." The correct assignments are: 1 = F, 2 = E, 3 = C, 4 = A, 5 = D, 6 = B.

Explanation B. Incorrect. This answer confuses DHCP and Proxy servers, and also NTP and RAS servers. That is, item 4 should be A, 6 should be B, 3 should be C, and 5 should be D. The correct assignments are: 1 = F, 2 = E, 3 = C, 4 = A, 5 = D, 6 = B.

Explanation C. Incorrect. This answer confuses FTP and Proxy servers. That is, "6. Intermediary that caches web pages" should be "B. Proxy server" and "1. Transfers files across the Internet" should be "F. FTP server." The correct assignments are: 1 = F, 2 = E, 3 = C, 4 = A, 5 = D, 6 = B.

Explanation D. Correct. This answer correctly matches the various servers to their brief definitions.

Chapter 3

1. Answers: A, C

Explanation A. Correct. This statement is not true. If two computers access the same device they may well refer to it using different logical unit numbers, or LUNs.

Explanation B. Incorrect. The statement is true. LUNs are divided into four 16-bit components that represent a multilevel addressing scheme (typically only the first of these is used).

Explanation C. Correct. The statement is not true. Fibre Channel uses LUNs too.

Explanation D. Incorrect. The statement is true. It is the standard definition for a LUN or logical unit number, as used in disk storage.

Explanation E. Incorrect. The statement is true. LUN masking makes a LUN available to some hosts but not others.

Explanation F. Incorrect. The statement is true. Other ways to refer to a LUN or logical unit include its SCSI Device ID and the disk serial numbers.

2. Answer: B

Explanation A. Incorrect. The communications between server and SCSI chains would not work for a while and then fail if this were a problem. A more likely cause is that the SCSI drives might be improperly terminated.

Explanation B. Correct. If you have intermittent or mysterious communication failure on SCSI drive chains, one of the first things to verify is how those chains are terminated.

Explanation C. Incorrect. The communications between server and SCSI chains would not work for a while and then fail if this were the problem. A more likely cause is that the SCSI drives might be improperly terminated.

Explanation D. Incorrect. Although inconsistent power levels can cause odd problems if undetected, a more likely cause of intermittent or mysterious communication failure on SCSI drive chains is improper termination. The best answer is that the SCSI drives might be improperly terminated.

Chapter 4

1. Answer: B

Explanation A. Incorrect. You'll definitely want to create a baseline after you install and configure the servers. But while the baseline helps you track and ensure good performance and availability over time, it does not help you define and agree upon what those expectations are with the project team you support. The correct answer is to meet with the project team and develop an SLA, or service level agreement. The SLA will define agreed-upon performance and availability targets.

Explanation B. Correct. An SLA is a Service Level Agreement. Its purpose is to set mutually agreed upon expectations between yourself and the project team you support, for parameters such as performance and availability.

Explanation C. Incorrect. Recall that an HBA is a host bus adapter. This has no relevance to the question. The correct answer is to meet with the project team and develop an SLA, or service level agreement. The SLA will define agreed-upon performance and availability targets.

Explanation D. Incorrect. This action does nothing to ensure you meet the project team's performance and availability expectations for its servers. The correct answer is to meet with the project team and develop an SLA, or service level agreement. The SLA will define agreed-upon performance and availability targets.

Explanation E. Incorrect. While it is a good idea to have the managers meet one another, the way to ensure that you are meeting the project team's expectations in regards to server performance and availability is to jointly develop and agree upon an SLA, or service level agreement. The correct answer is therefore to meet with the project team and develop an SLA. The SLA will define agreed upon performance and availability targets.

2. Answer: D

Explanation A. Incorrect. If you take the servers down without prior advance notification to the user community #and without their approval and buy-in# you are very likely headed for friction with this project team. Depending on the project and their use of the servers, they might (rightly) be angry at your actions. The correct answer is to meet with the project team and its management well in advance to establish an agreed upon time for the servers to be down. Then use the project team's existing procedures to notify everyone about the window of unavailability for their servers.

Explanation B. Incorrect. If you take the servers down without prior advance notification to the user community #and without their approval and buy-in# you are very likely headed for friction with this project team. Depending on the project and their use of the servers, they might (rightly) be angry at your actions. The correct answer is to meet with the project team and its management well in advance to establish an agreed upon time for the servers to be down. Then use the project team's existing procedures to notify everyone about the window of unavailability for their servers.

Explanation C. Incorrect. If you take the servers down without prior advance notification to the user community #and without their approval and buy-in# you are very likely headed for friction with this project team. Depending on the project and their use of the servers, they might (rightly) be angry at your actions. The correct answer is to meet with the project team and its management well in advance to establish an agreed upon time for the servers to be down. Then use the project team's existing procedures to notify everyone about the window of unavailability for their servers.

Explanation D. Correct. You must secure the project team management's agreement to your plans well in advance. Then you must be sure to use the project team's preferred communications methods to keep everyone abreast of your progress, including notification as to when the servers become available again.

Explanation E. Incorrect. You need to talk with your manager about your plans and ensure he or she also agrees with them, but in very few organizations can you entirely and completely off-load responsibility for coordinating outages with your customers to your manager. You are the one supporting the project team and their servers, so you must ensure the project team and its manager is notified well in advance of your plans#and that they agree to them. The correct answer is to meet with the project team and its management well in advance to establish an agreed upon time for the servers to be down. Then use the project team's existing procedures to notify everyone about the window of unavailability for their servers.

3. Answer: C

Explanation A. Incorrect. Typically you apply major upgrades to test server(s) first, then development server(s), then production server(s). The fact that the vendor has certified and released the upgrade does not guarantee it will run problem-free on your servers. You should roll out the upgrade in a planned fashion to minimize and address any problems it might cause.

Explanation B. Incorrect. This is only a part of the correct answer. The correct answer adds “develop an upgrade plan,” which is a necessary part of the upgrade roll-out. Having a written upgrade plan in advance helps you reduce the chances of problems occurring during the upgrades. Typically the plan requires upgrading test server(s) first, then development server(s), then production server(s), according to some agreed upon schedule.

Explanation C. Correct. Developing an upgrade plan helps reduce the chances that problems occur due to the upgrade and minimizes any problems that do occur. Typically this means that major upgrades are applied first to test server(s), then development server(s), and finally production server(s).

Explanation D. Incorrect. You need to develop an upgrade plan. Then, you’ll typically apply changes to test server(s), then development server(s), and finally production server(s) to minimize the impact of any problems that are found in the upgrade software. A basic IT principle is that you apply any upgrades to non-production server(s) first to identify and shake out any problems or bugs.

Explanation E. Incorrect. You need to develop an upgrade plan. Then, you’ll typically apply changes to test server(s), then development server(s), and finally production server(s) to minimize the impact of any problems that are found in the upgrade software. A basic IT principle is that you apply any upgrades to non-production server(s) first to identify and shake out any problems or bugs.

Chapter 5

1. Answer: C

Explanation A. Incorrect. A hot site has ready-to-use hardware, software, and backup data in place, while a cold site does not have backed up copies of the data or ready-to-use hardware and software.

Explanation B. Incorrect. A hot site has ready-to-use hardware, software, and backup data in place, while a cold site does not have backed up copies of the data or ready-to-use hardware and software.

Explanation C. Correct. A hot site has ready-to-use hardware, software, and backup data in place, while a cold site does not have backed up copies of the data or ready-to-use hardware and software.

Explanation D. Incorrect. A hot site has ready-to-use hardware, software, and backup data in place, while a cold site does not have backed up copies of the data or ready-to-use hardware and software.

2. Answer: A

Explanation A. Correct. This answer correctly matches all the technologies with their descriptions.

Explanation B. Incorrect. This answer confuses DDS, which are standards for encoding on DAT tapes, with DAT itself. Item 4 should be E and 5 should be C. The correct answer is: 1 = D, 2 = A, 3 = B, 4 = E, 5 = C.

Explanation C. Incorrect. This answer confuses AIT with DAT. That is, item 1 should be D and 5 should be C. The correct answer is: 1 = D, 2 = A, 3 = B, 4 = E, 5 = C.

Explanation D. Incorrect. This answer confuses the terms for items 2 and 3, and those for items 4 and 5. The correct answer is: 1 = D, 2 = A, 3 = B, 4 = E, 5 = C.

3. Answers: A, D

Explanation A. Correct. Testing enables you to get any kinks out of the DRP or disaster recovery plan and update it accurately prior to actually having to rely on it in a disaster recovery situation.

Explanation B. Incorrect. While testing the DRP or disaster recovery plan may be allocated for in the budget, the reasons for testing it are to get out any bugs prior to relying on it and to ensure it works when you need it.

Explanation C. Incorrect. While testing the DRP or disaster recovery plan may be required by management, the real reasons for testing it are to get out any bugs prior to relying on it and to ensure it works when you need it.

Explanation D. Correct. You must test the DRP or disaster recovery plan because it is the only way to know it will work when you need it. Imagine if your first “test” of the plan was in when a disaster occurs, and then you discover it has bugs in it!

Explanation E. Incorrect. While testing the DRP or disaster recovery plan may help impress upon the project team that disaster planning is important, the real reasons for testing it are to get out any bugs prior to relying on it and to ensure it works when you need it.

Chapter 6

1. Answers: A, D

Explanation A. Correct. One of the main causes of transient memory problems is overheating. Ensure the system is properly cooled and that all fans are working.

Explanation B. Incorrect. Power loss would not cause transient memory errors. It could cause server shutdown or a switch to backup power. The correct answers are overheating and bad memory.

Explanation C. Incorrect. There is no reason that simply using all the memory slots in a server would cause transient memory errors. The server has been designed so that you can use all the memory slots. The correct answers are overheating and bad memory.

Explanation D. Correct. Bad or failing memory could be a cause of transient memory errors. Take out the memory and test it, and then replace any bad memory.

Explanation E. Incorrect. System configuration should not cause transient memory errors even if it is incorrect in some manner. Transient memory errors are a hardware issue, not one of software configuration. The two most likely causes are overheating and bad memory.

Explanation F. Incorrect. High humidity alone will not cause transient memory errors, although overheating could. The correct answers are overheating and bad memory.

2. Answer: C

Explanation A. Incorrect. The chart shows that processor utilization has not increased enough to identify this as the cause of the performance problem. The correct answer is that the chart shows the paging rate has increased dramatically since the baseline. Therefore, increasing memory is the best way to resolve this problem.

Explanation B. Incorrect. The chart shows that disk utilization has not increased enough to identify this as the cause of the performance problem. The correct answer is that the chart shows the paging rate has increased dramatically since the baseline. Therefore, increasing memory is the best way to resolve this problem.

Explanation C. Correct. The chart shows that the paging rate has increased dramatically since the baseline. Increasing real memory is the best way to resolve this problem.

Explanation D. Incorrect. The chart shows that network utilization has not increased enough to identify this as the cause of the performance problem. The correct answer is that increasing memory is the best way to resolve this problem.

3. Answers: B, D

Explanation A. Incorrect. As long as the disk drive spins and works okay, you should be able to retrieve its data. One possibility is to remove the disk drive and place it into another server as a secondary drive (resetting its jumper appropriately), then read its data. Another is to boot with a Live Linux CD and read the data from the drive.

Explanation B. Correct. Assuming you know which drive has the data you need to save, and assuming that drive works okay, you can remove the drive, insert it into a secondary machine as a secondary drive (setting the jumpers and connecting the cables properly), boot the secondary machine, and retrieve and save the data from the drive.

Explanation C. Incorrect. This approach might work, but there is no need to install Linux to the hard drive to save its data. In fact, this could over-write the data you want to save #if you let Linux do a "destructive install." A better approach is to simply boot Linux as a Live CD/DVD, mount the disk with the data you need to save, and access and save that data.

Explanation D. Correct. A very effective way to save data from Windows computers that will not boot into Windows is to boot into a Live Linux CD/DVD. Then mount the disk with the data you need to save and access and save the data.

Explanation E. Incorrect. Linux can read Windows NTFS (and FAT) disk drives. A very effective way to save data from Windows computers that will not boot into Windows is to boot into a Live Linux CD/DVD. Then mount the disk with the data you need to save and read and save the data.

4. Answers: B, C, D

Explanation A. Incorrect. Partition managers like QtParted or GParted can shrink partitions (although not in every case, as this question notes).

Explanation B. Correct. You can only resize unmounted partitions using most partition managers.

Explanation C. Correct. QtParted or GParted can resize most kinds of partitions, but not all, depending on the partition type and the kind of file system it uses.

Explanation D. Correct. Most partition managers will refuse to shrink any partition that fails to scan without errors. You need to eliminate these errors before you can successfully shrink the partition.

Explanation E. Incorrect. Partition managers like QtParted or GParted can resize all kinds of Windows partitions including NTFS.

Explanation F. Incorrect. Partition managers like QtParted or GParted can resize all kinds of Windows partitions including FAT types (like FAT32 and FAT16).

5. Answer: D

Explanation A. Incorrect. Especially on older SCSI you'll want to be sure the boot disk is ID 0 and the SCSI adapter is ID 7. While you can't be certain this will fix the problem it is certainly what you'd want to try first.

Explanation B. Incorrect. If you set the boot hard disk to ID 0 you'll also have to set the SCSI adapter to some other ID (since it is currently set at 0). The correct answer is to set the SCSI boot disk to ID 0 and the host adapter to SCSI ID 7.

Explanation C. Incorrect. If you set the host adapter to ID 7 you'll also have to set the boot disk to some other ID (since it is currently set at 7). The correct answer is to set the SCSI boot disk to ID 0 and the host adapter to SCSI ID 7.

Explanation D. Correct. Especially on older SCSI you'll want to be sure the boot disk is ID 0 and the SCSI adapter is ID 7. While it is not guaranteed that this will fix the problem, it is certainly the first thing you'd do.