Virtualization
Unit-II

Q.16	When should server Virtualization be used? When should it not be used? Explain
A	• The server virtualization technology you choose must meet all of your requirements. For you to make the right choice, you must be aware of and understand what the major choices are and which technologies offer the most features based on your own requirements. • While there are several different manufacturers offering virtualization engines, there are really only three that offer extensive feature sets: VMware, which was the founder of server virtualization on x86 technologies; Citrix, which, with its purchase of Xen Source in 2007, now offers the XenServer line of products; and Microsoft, who now offers virtualization through the Hyper-V code in Windows Server 2008. • Most other vendors offer technologies that are based on the Xen open-source code. This means that vendors such as Oracle, Sun, Virtual Iron, and more, really offer nothing much different than Citrix, since all of their offers are based on the same code. But in order for you to choose the right product for your needs, you must first be aware of what makes up what each vendor offers. Then, you will look to pricing and implementation costs. • Once you begin to understand which product best fits your bill, you must look to what you need to do and how best it can be implemented. • Once again, looking to the offerings of each vendor will assist you in the process. Finally, your decision will be crystallized through the analysis of the metrics for each major hypervisor—how many virtual machines can run on one host, how hosts are managed, which operating systems are supported in the virtual layer, and so on—which will help position each vendor in the marketplace and view its long-term value.
Q.17	**Explain Emulation, Simulation, and Virtualization.**
A	• Emulation is a concept that allows one environment to act or behave as if it were another environment. This could also be described as sophisticated impersonation. An environment is an execution platform, operating system, or hardware architecture. • Instructions are interpreted from the executing environment into instructions that the real, underlying environment understands. Emulation is used for running legacy environments, operating system development, and software testing. • Emulated environments incur a high performance penalty when compared to virtualized systems due to the overhead of the interpreter. • Common implementations of emulation include: Bochs, MAME, Virtual PC for Mac, WINE. • Simulation is a concept in which an environment imitates another environment. This imitation simply accepts pre-defined inputs and provides pre-defined responses. This is arguably the easiest or least complex concept to implement. • An environment is an execution platform, operating system, or hardware architecture. Simulators are used differently than both emulation and virtualization. They are primarily used in hardware and microchip design and prototyping. • By doing this, testing can be done on hardware and microchips yet to be built! This reduces the costs and risks associated with mistakes being made on hardware and chips before they are fabricated. • Common implementations of simulation include: Cadence and Synopsis, Simics, SimOS.

	• The use of emulators and simulators have their places, however virtualization is the only technology that enables revolutionary capabilities in the datacentre. • Virtualization provides the benefits necessary to give IT organizations the ability to save costs on hardware and increase the efficiency of server deployments, provisioning, and management. Virtualization also enables physical hardware independence, which gives IT the flexibility and freedom of not being locked in to a single vendor's hardware solution.
Q.18	**Explain Hardware Partitioning.**
A	• Hardware partitioning is a highly-specialized hardware technology that allows the computing resources of a single, physical computer to be divided into multiple partitions (often called hard partitions), each of which can host its own, isolated instance of an operating system. • Hardware partitioning has existed for quite some time, originating in high-end mainframe systems from IBM. Today, there are several hardware partitioning technologies available, although each implementation is proprietary and requires very specific server hardware and software to be used. • In some implementations, only one or two very specific operating systems are supported. In general, all of the required components of a system featuring hardware partitioning are only available from a single vendor, due to their proprietary nature. One of the key advantages of hardware partitioning is its very efficient resource sharing and management capabilities. These systems are much more efficient than equivalent software partitioning systems because the resource management between hard partitions is handled using separate hardware components (chips, circuits, memory, storage, etc.). • The specialized software (sometimes referred to as microcode) that performs the actual resource management resides in the specialized resource management hardware components as well. As a result, the available performance in each hard partition is maximized and remains unaffected by the resource management system's overhead. • This is very different from software partitioning technologies where the partitioning occurs in software that is executed using the same hardware that is being managed and shared. Another advantage, available in some implementations of hardware partitioning, is electrical isolation of each hard partition. • Electrical isolation in hardware partitioning systems allows a hardware fault to occur in one hard partition while not affecting any other hard partition in the same system. Systems offering hardware partitioning technologies are usually mid-range to high-end computing systems that are generally very scalable (usually scaling up) and robust. • Hardware partitioning systems have several disadvantages: expensive, proprietary hardware and software, additional costs incurred by the support and maintenance of the proprietary hardware, limited support for various operating systems, limited hardware portability for an existing installed base of hard partitions, and vendor lock-in. • Proprietary hardware and software systems almost always have additional costs for installation, training, support, and maintenance due to the lack of expertise of most IT organizations with these types of systems. Often vendors will only allow their services organization to perform the installation and support of these systems. • Hardware partitioning systems generally only allow one type of operating system to be installed; of course, each hard partition supports a separate instance of that operating system. There are some systems that are more flexible and support more than one operating system, but it is almost always limited to operating systems provided by the vendor.

- Aside from limited operating system support, hardware partitioning systems have very limited portability of existing partitions. Generally, these partitions may only be moved to systems comprised of the same vendor's hardware because of the lack of complete hardware abstraction.

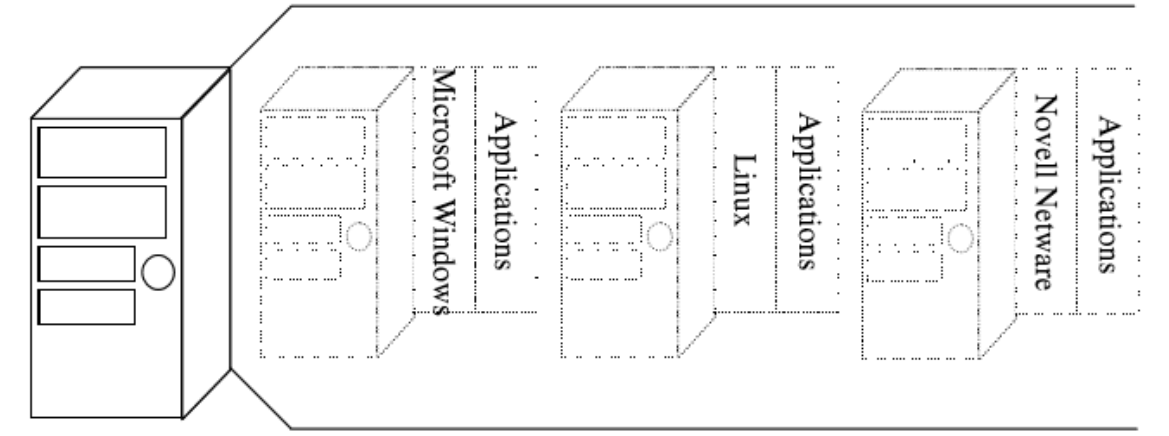

Figure 2.2 Hardware Partitioning.

Q.19 Explain Software Partitioning

A

- Software partitioning is a software-based technology that allows the resources of a single, physical computer to be divided into multiple partitions (also called soft partitions or virtual machines), each of which can host its own, isolated instance of an operating system.
- Software partitioning is generally similar to hardware partitioning in that multiple instances of operating systems may coexist on a single physical server. The major difference between hardware and software partitioning is that in software partitioning, the isolation of each soft partition and the management of the shared resources of the computer are completely handled by a special software layer called a Virtual Machine Monitor (VMM) or Hypervisor.
- The VMM, as well as each operating system within each soft partition, all consume computing resources from the same set of hardware, thus software partitioning incurs overhead that does not exist in hardware partitioning.
- The overhead produced by the VMM varies from each implementation of software partitioning systems, but always has an impact on the performance of each soft partition. Depending on how resources are managed by the VMM, it is conceivable that the computing resources consumed by each soft partition could also impact the VMM's performance as well.
- Software partitioning implementations exist on mid-range to high-end computing systems as well as commodity server and workstation computers. Server virtualization (and the term virtualization) as described in this book refers directly to server-based software partitioning systems, generically referred to as virtualization platforms. Software partitioning systems are generally implemented in one of two ways.
- They are either hosted as an application in an existing operating system installed on a physical computer or they are installed natively on a physical computer without an operating system. When software partitioning systems are hosted in an existing operating system, they gain the advantages of leveraging that operating system's resource management capabilities, its hardware

compatibility, and application programming interfaces (APIs).

- This allows the software partitioning system to be smaller and potentially easier to write and support. This configuration also imposes the deficiencies and inefficiencies of the host operating system upon the software partitioning system as well as the additional resource consumption of the host operating system. Hosted software partitioning systems generally have more overhead and less performance than their native counterparts.

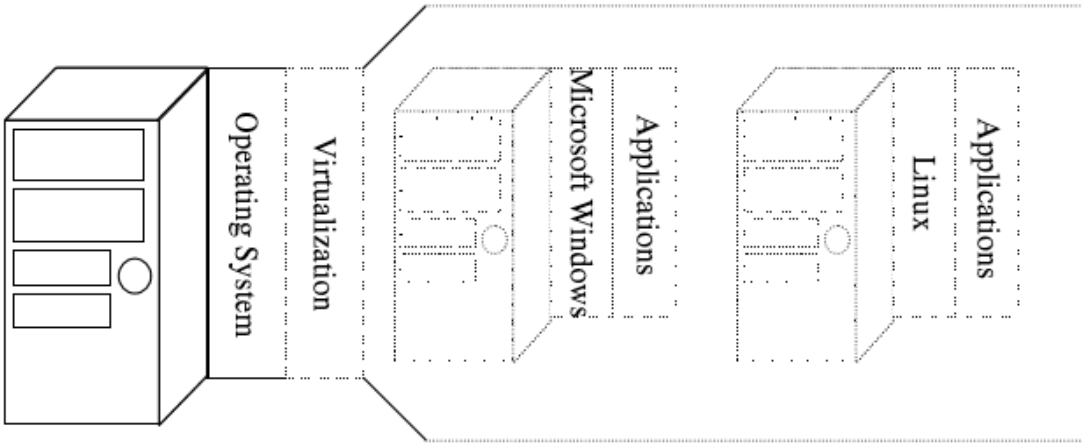

Figure 2.3 Software Partitioning Hosted in an Existing Operating System.

Q.20 What is Application Partitioning? Explain.

A.

- Application partitioning is a software-based technology that allows the operating system resources on which an application depends to be placed in an alternate container within the operating system without the application's knowledge.
- These resources are said to be virtualized by the application partitioning system. The isolated application can then be executed in multiple instances simultaneously in the same operating system, by one or more users, without the application instances interfering with one another. Each instance of the application has no knowledge that the other instances exist and the application does not require any modifications to be hosted by the application partitioning system.
- The primary advantage of an application partitioning system is that any application, regardless if it was designed to be used by a single user or multiple users, can be centrally managed and made available in a distributed fashion.
- A single server can execute many instances of the application and each application instance state is written into a separate container. Each container is automatically handled by the application partitioning system.
- Application partitioning can consolidate a single application from multiple desktop computers and servers onto a single server and the application can be managed much like a single instance of the application.
- The operating system itself is not completely abstracted from the application, only certain subcomponents such as data storage facilities (fi le systems), therefore only applications normally run on the operating system being used are allowed to be hosted under the application partitioning system.

	• Application partitioning systems are available from the following companies (listed with their product name): Citrix—Metaframe Presentation Server, Ensim—VPS, Softricity—SoftGrid, SWSoft—Virtuozzo, Tarantella—Secure Global Desktop. 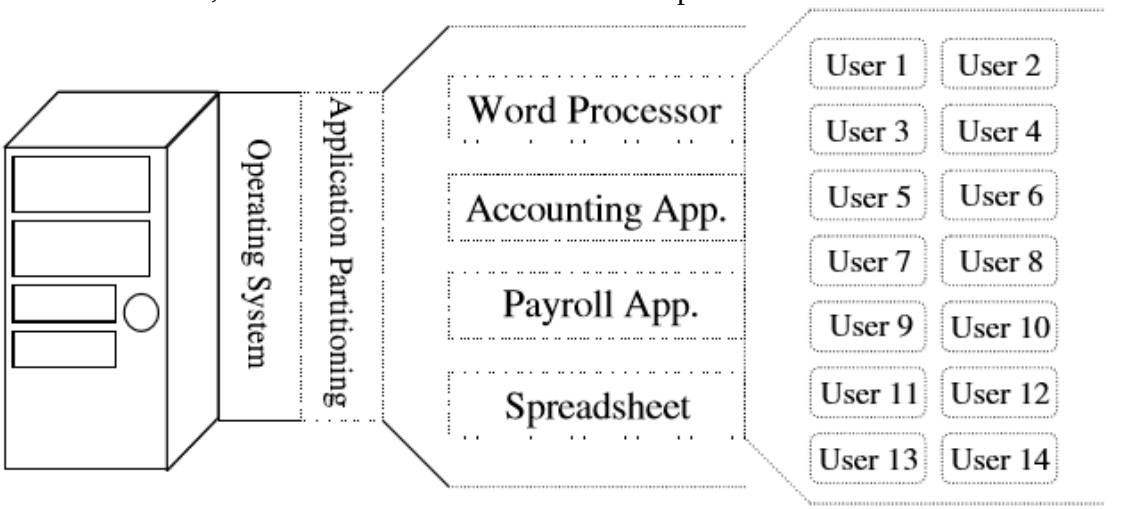 **Figure 2.5 Application Partitioning.**
Q.21	**Explain Resource Partitioning and Service Partitioning.**
A.	• Resource partitioning is a software-based technology that abstracts how certain operating system resources are allocated to application instances or individual processes executing within the operating system. • This technology is used to control resource consumption of applications and processes, allowing more granular control than what is provided by the operating system. Resource partitioning systems also allow the resource consumption to be controlled not only at the application or process level, but also by the combination of application or process and user account. • Resource partitioning systems enable the operating system to become Quality-of-Service enabled. Application instances or processes can be given parameters, which allow certain minimum and maximum levels of resource utilization such as CPU, memory, and disk I/O to be effectively controlled and managed. • Just as in application partitioning, resource partitioning does not abstract the entire operating system from an application, therefore only applications that would normally run on the operating system being used are allowed to be controlled by the resource partitioning system.

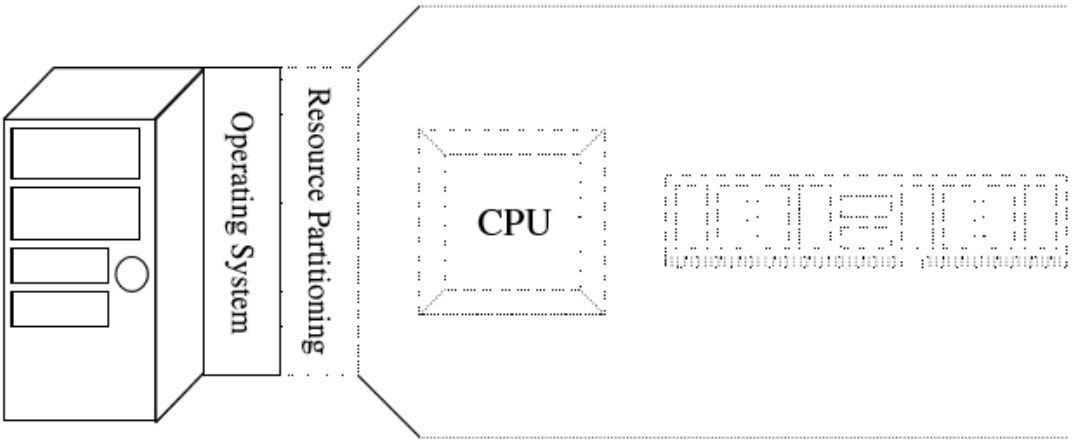

Figure 2.6 Resource Partitioning.

- Service partitioning is a software-based technology in which a single application instance provides multiple, isolated instances of a service.
- Each instance of the service usually appears to consumers of that service to be a dedicated application instance (and often a dedicated server instance). Abstraction between the service application and the operating system is not required. The abstraction occurs on top of the application instance, which allows multiple instances of the application's service to coexist.
- The level of isolation between service instances can vary greatly between implementations of service partitioning, and in some systems can even be controlled, providing complete isolation at one extreme and no isolation at the other extreme of the isolation configuration settings.

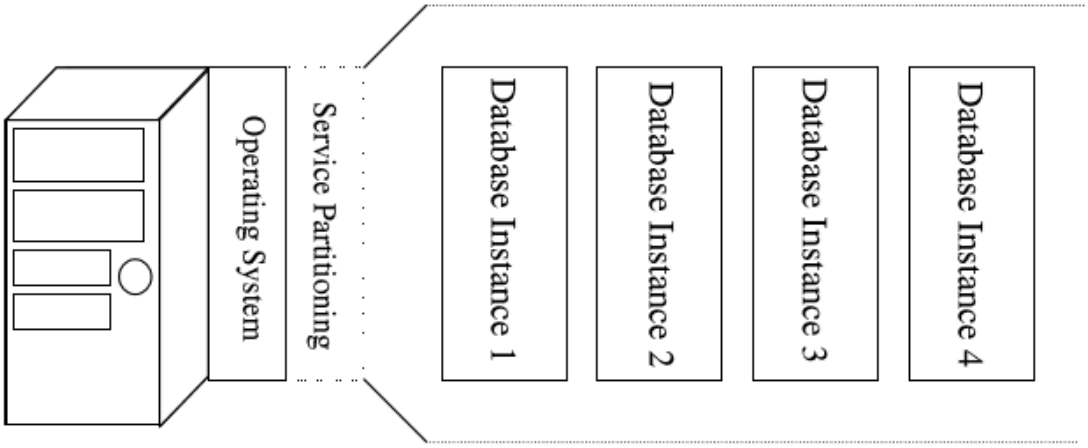

Figure 2.7 Service Partitioning.

Q.22	**What are Host Servers, Virtual machines and guest operating system? Explain**
A.	• The host server is probably the most familiar component. The host server is the physical, bare metal, x86 server computer on which the virtualization layer resides. • If the host server has an operating system installed, the operating system is referred to as the host operating system or host OS. The host server is referred to as a "host" because it hosts the virtualization layer and one or more guest servers or virtual machines.

	• A virtual machine is a single, isolated execution environment that is created through software by the virtualization layer, which appears to be almost exactly like a physical x86 computer. A virtual machine, also referred to as a VM or virtual server, represents a single virtual computer. A compatible operating system and applications may be installed into a virtual machine. • Although the VMM creates and manages the virtual machines, virtual machines also require resources of a physical host server in order to run. These resources are allocated through the VMM and the virtual machine and by default do not have any knowledge of the existence of the host server or the virtualization layer. • It simply "sees" the virtual hardware exposed by the VMM as if it were real hardware. Some of the virtual hardware is emulated or simulated by the virtualization layer and some of it actually maps to real hardware, exposed through the VMM in a very controlled manner. • A guest server is another term for a virtual machine. It is a term that is used to easily differentiate a physical server, also called a host server, from a virtual server, also called a guest server. Guest servers have an operating system installed that is compatible with the virtual hardware and often with the host server. • This operating system does not have to be the same operating system installed on the host server (i.e., if the host server has an operating system). An instance of an operating system that resides in a guest server is referred to as the guest operating system or guest OS. • Guest servers are isolated from each other and, by default, have no knowledge of the virtualization layer or other guest servers. • This isolation is referred to as Host/Guest isolation. The term guest server is often used interchangeably with the term virtual machine. The term guest operating system or guest OS always refers to an instance of an operating system installed in a guest server (or virtual machine).
Q.23	**Give an overview of the following virtual hardware:** **Processor, Memory, Hard Disk Drive**
A.	**Processor** • The processor is the most significant piece of hardware in a computer. Also referred to as the central processing unit (CPU), the processor is often thought of as the brains of a computer. Its primary function is to perform complex calculations on data. • In conventional x86 servers, there are often one, two, or four processors, with high-end servers having eight or more. Traditionally, operating systems are responsible for sending instructions to the physical processor(s) to perform requested commands. • Most server virtualization platforms support uniprocessor virtual machines, that is, virtual machines with a single virtual processor. However, the virtual processor is not a fully emulated processor. • Completely emulating the processor would cause a severe performance bottleneck. Instead, the virtualization platform generally traps instructions issued by virtual machines and either passes the instruction through to the physical processor or emulates the instruction by issuing one or more different instructions to the physical processor and returning the expected result to the virtual processor. **Memory** • Computers typically use random access memory (RAM) as the temporary, high speed storage location of data that is currently being used by the processor(s). Memory is many thousand times faster than disk access, so programs and data are loaded one time from disk into memory before

	being used by the processors. • Just as their physical counterparts, virtual machines require memory as well. As of this writing, 3.6GB of memory is the maximum amount of memory that can be assigned to a single virtual machine in most virtualization platforms. • This limitation exists because 32-bit x86 architecture systems support a maximum memory size per process of 4GB and the virtualization platform consumes some of that memory as overhead. Currently, there is no virtual physical address extensions feature (or virtual PAE) designed to overcome this limitation on 32-bit systems. • The virtual memory allocated to a virtual machine usually maps directly to physical memory of the host server. The term virtual memory in this context should not be confused with virtual memory as in a paging or swap fi le as used by many operating systems to extend the amount of physical memory available by using hard disk storage as additional memory. **Hard Disk Drives** • Hard disk drives are the most commonly used hardware components to provide long-term, non-volatile data storage. Unlike a computer's memory, its hard disk drives, or disks, are able to persist data for long periods of time, even when the computer is powered off. • Modern hard disk drives provide massive amounts of storage that is very inexpensive relative to the cost of memory chips. The most significant difference between hard disk drives and memory for data storage is that hard disk drives are typically the slowest device in the data access chain, usually thousands of times slower than memory. • Data is often first read from a hard disk drive and placed into the computer's memory before it can be used by the processor. When data is changed in memory, it often must be written back to the hard disk drives for permanent storage. • There are two common types of hard disk interfaces, IDE and SCSI. These interfaces support the disk drives that are directly connected to a physical computer. They differ from each other in physical aspects such as the types of physical connectors, cables, and electrical signals, as well as in software including the protocol used to transfer data between the hard disk drive and the computer to which it is connected.
Q.24	**Give an overview of the following virtual hardware:** **Floppy Disk Drive, DVD Drive, USB Port**
A.	**Floppy** • Disk Drives Computers are slowly moving away from the use of floppy disk drives; however, they still refuse to completely disappear in either the physical or virtual world. The floppy disk drive has remained relatively standard and unchanged for many years. • A virtual floppy disk drive can read and write to a physical floppy diskette that is inserted into the host server's floppy disk drive. • By mapping the host server's floppy disk drive to the virtual machine, the guest operating system can interact with the host floppy disk drive as if it were physically attached to the virtual machine. • One important distinction between the physical and the virtual floppy disk drive is that a virtual floppy disk can be mapped to a floppy disk image fi le located on the host server. A floppy disk image file (also known as a floppy image) can be created by most virtualization management utilities or by using third-party tools. • By mounting or attaching a virtual floppy image to a virtual machine rather than to the physical

drive, huge performance increases in both read and write speed will be realized. This performance increase is due to the information being read from a fi le residing on a hard disk drive instead of from a physical floppy disk drive with very slow access, read and write times.

- Virtual floppy disk drives are needed in some cases to assist the installation of some guest operating systems.

DVD-ROM Drives

- The CD/DVD drive has taken the place of the floppy drive as the preferred storage of choice. CD and DVD discs have become the standard medium for distribution of large quantities of data.
- The average CD can hold between 650 and 700MB while a standard DVD can hold 4.7GB. Compared to the 1.44MB that the floppy disk provides, it is easy to see why the CD/DVD has replaced the floppy disk. CD/DVD media, like most things in the computer industry, follow certain standards and guidelines.
- CD-ROM discs are the media of choice when it comes to moving data between virtual machines or when operating systems or applications need to be installed on a virtual machine. Like the virtual floppy drive, a virtual CD/ DVD drive can read physical media that is inserted into the host server.
- By mapping the host server's CD/DVD drive to the virtual machine, the guest operating system will interact with the host drive, as if the media were physically inserted into the VM. Also like the virtual floppy drive, there is an important distinction to be made between the virtual and the physical CD/DVD drive.
- A virtual CD/ DVD drive can be mapped to a physical fi le located on the host's fi le system. This file, as explained above, has been defined as an ISO image. By mounting to a virtual ISO image rather than a physical drive, a huge performance boost in reading the disc is achieved.
- At the physical level, fi le reads from a hard disk are much faster than fi le reads from a CD-ROM.

USB Ports

- A relative newcomer to the computer port scene is the universal serial bus (USB). The USB design consists of a single host controller and multiple devices connected in a tree like formation using special hub devices. At least one hub, the root hub, always exists.
- The root hub is attached directly to the host controller. The hardware that contains the host controller and the root hub is called the Host Controller Device (HCD).
- The HCD is defined by the hardware manufacturer. Found in many of today's computer systems is the Intel specifications, the Universal Host Controller Interface (UHCI) and the Extended Host Controller Interface (EHCI). Devices that attach to the bus can either be custom devices or class devices.
- A custom device requires a custom driver to accompany it in order for the device to be used. On the other hand, a class device follows the same behavior and interface descriptors, so that the same device driver may be used for devices that are members of the same class. One advantage to using USB over serial or parallel ports is that a single USB port can be used to connect up to 127 peripheral devices.
- Common uses of the USB port are connecting devices such as mice, keyboards, printers, scanners, removable hard drives, and networking components. Unlike a SCSI bus, devices on a USB chain do not require a terminator. Another advantage of using USB, certain devices that require a low power draw can receive their power from the bus, thereby reducing the need for extra power sources.
- USB was originally designed for ease of use. By allowing devices to chain from a single host, the

	need to add multiple expansion cards to the computer's bus was removed. Additionally, USB added improved plug and play ability as well as the capability of hot swapping devices without the need for a reboot of the system. Because of these types of features, USB is quickly becoming the standard of choice for peripheral manufacturers of devices such as printers and scanners.
Q.25	**Give an overview of the following virtual hardware:** **Network adapter, Serial port, Parallel Port, Keyboard and Mouse.**
A.	**Network Adapters** • There are some distinctions between the host network adapter and the guest network adapter. Depending on configuration, these virtual network adapters can be isolated into a separate private network that only communicates with other virtual network adapters in that private network. • They can be isolated by themselves so that no communication occurs with other machines. Or they can be joined to a network in which they only communicate with the host machine. Another option is for several machines to be part of a private network that also allows them to talk to the host. • Virtual network adapters can also talk directly through the host network adapter as if they were a separate network interface on the host's network. And finally, it is possible for the virtual network adapter to communicate through the host network adapter as if it were the host adapter (through a process called NAT). **Serial Ports** • A serial port is a port, or an interface, on a computer system where information is transmitted in or out one bit at a time. Throughout the computer system's history, most serial ports conformed to the RS-232 or RS-422 standards. • Using a simple cable and this general purpose interface, the serial port can be used to connect almost any type of device. Through the serial port, the computer can be connected to devices such as terminals, modems, keyboards, mice and printers. • Originally specified as a 25 pin D-type connector, most of these pins were unused since data was only being transmitted one bit at a time. Therefore, these devices were replaced with parallel port devices, and later, USB and Fire wire devices. **Parallel Ports** • A parallel port is a port, or an interface, on a computer system where data is transmitted in or out eight bits at a time. The data is transmitted in parallel, across more than one wire at a time. Contrary to the serial port, the parallel port is able to transmit one bit of data across multiple wires, rather than a across a single cable. • Within virtualization, a virtual machine may be configured to have a virtualized parallel port. Like the serial port, the virtual parallel port can be mapped to the host's parallel port. The virtual machine can then make use of the host's attached devices. While not every physical device will work within the virtual machine, many do. • The obvious choice of using a parallel port in a virtual machine is to allow the virtual machine to send output to a physical printer. Other uses may include mapping the parallel port to make use of a data drive or an application security dongle. **Keyboard And Mouse** • The keyboard is the device that allows the entering of data into a computer. Over the years, the keyboard has been modified and updated, however it has remained fairly standard in its layout. The English language keyboard has followed the QWERTY key scheme. • This standard layout of letters, numbers, and punctuation gets its name from the first six letter

keys of the top row: QWERTY.

- A mouse is a pointing device for computers that is typically fitted with one or more buttons that have different functions depending on the operating system or application running. The mouse's 2D motion is then translated into the movement of the cursor or pointer on a display screen.
- The device derived its name through its appearance. Early mice had a mouse like shape and a cord that resembled a tail. Additionally, the movement of the device and the pointer on the display screen seemed to scurry around much like a mouse.

Q 26 **Explain fixed disk file, dynamic disk file and differencing disk file and undo disk file.**

A.

Fixed Disk File

- A fixed disk is a virtual hard disk drive that most closely resembles a physical hard disk drive. The fixed disk is represented by a fixed-extent file residing on the host server. When created, the file is sized immediately to its maximum size, consuming an equal amount of storage from the host server.
- Like all virtual hard disk drives, the fixed disk is empty upon creation. For example, when a 10GB fixed virtual hard disk drive is created, a static 10GB fi le is created on the host server.
- This fi le represents the fixed disk and will contain all of its data. The file on the host server will always remain 10GB in size.
- It will not change, even when data is written to the fixed disk or when data is deleted from the fixed disk. Fixed disks will typically provide better performance than dynamic disks because there is no overhead due to file growth.

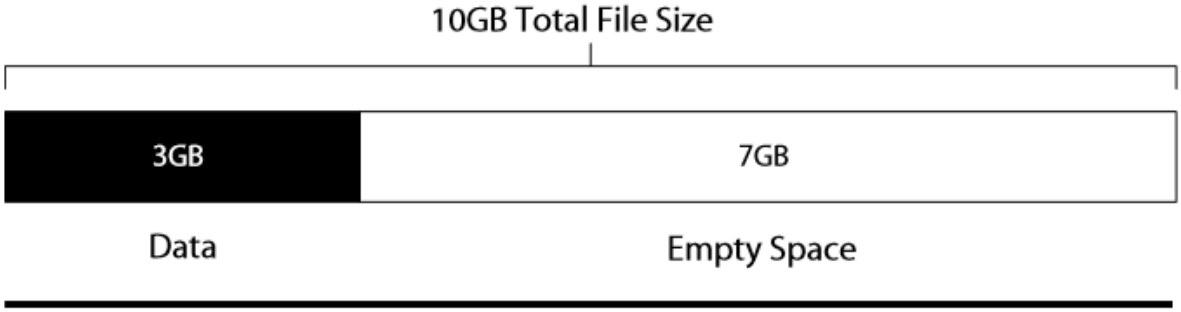

Figure 3.1 Fixed Disk File.

Dynamic Disk File

- A dynamic disk is a virtual hard disk drive that begins as a sparse file, consuming only the amount of storage from the host server that is needed, and grows as new data is written to the virtual hard disk drive. When created, the maximum size of the disk is specified, but the file representing the disk is sized only to the size needed to store its data at that point in time.
- The dynamic disk cannot grow in excess of its maximum size. For example, when a 10GB dynamic virtual hard disk drive is created, the maximum size of 10GB is stored within the dynamic disks' internal data structures. The file on the host server representing the dynamic disk at first may only be a few kilobytes in size.
- Once the process of installing a guest operating system begins, the dynamic disk will grow as new data is written to fill the disk. When files are deleted from a dynamic disk, the file on the host server representing the dynamic disk will not change (it will not shrink). Instead, the deleted sectors are marked and are reused as necessary.

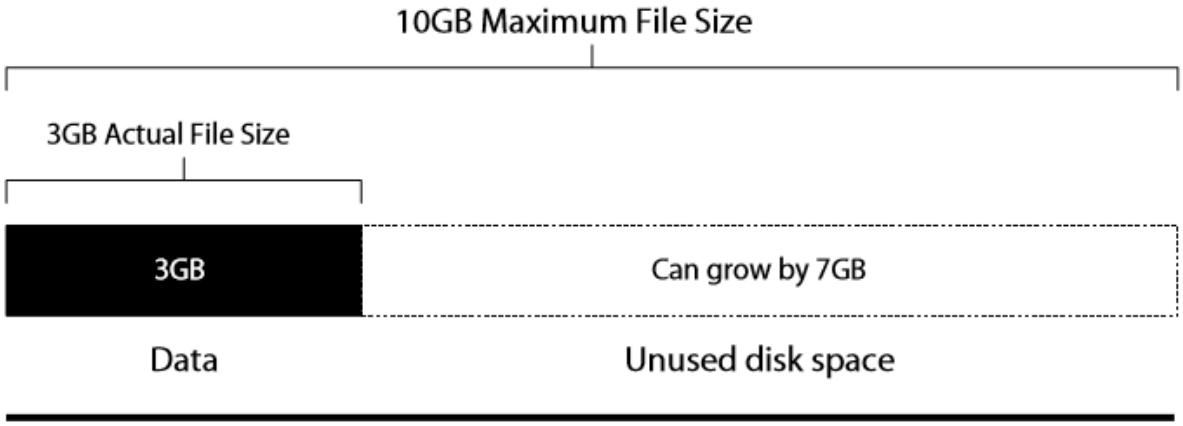

Figure 3.2 Dynamic Disk File.

Differencing Disk

- A differencing disk is a special type of dynamic virtual hard disk drive that is dependent on a parent disk. Differencing disks share a parent/child relationship with another virtual hard disk drive in which the differencing disk is the child.
- The combination of the differencing disk and its parent virtual hard disk drive create a single virtual hard disk drive as it appears to the virtual machine to which the diff erencing disk is attached.
- Once a differencing disk is linked to a parent disk, the parent virtual hard disk drive becomes read-only. Any changes made thereafter are saved into the differencing disk, which grows as needed. The parent of the differencing disk, however, controls the maximum disk size overall.
- When a read on the disk occurs, the most recent changes are always returned. First the read occurs in the parent disk and then in the differencing disk, returning the most current data from the differencing disk, if it exists in the differencing disk.
- If the parent disk of a differencing disk is ever altered in any way, the differencing disk will become invalid, as its relationship to the parent disk depends upon the state of the parent disk at the time they are linked.

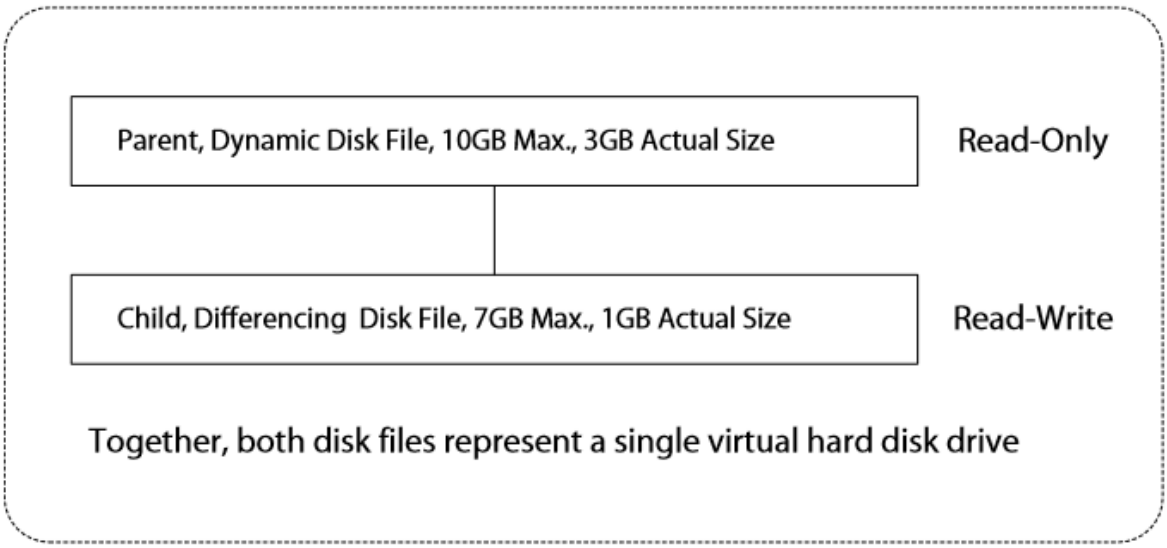

Figure 3.3 Differencing Disk Files.

Q.27 How does virtualization solve business problems?

A

- Aside from the pure novelty of running more than one operating system on a single computer at the same time, server virtualization technology was designed to be used to help solve business problems.
- Virtualization can be used in a very straight-forward manner much like physical servers and when used in this way can help reduce physical server hardware costs by consolidating multiple servers onto a single server.
- This usage is commonly referred to as server consolidation. Virtualization provides new features that are either unavailable or unobtainable on physical server hardware. One of the most profound benefits of virtualization is the abstraction between virtual machines and the physical hardware, thereby making virtual machines portable. In this way, virtualization simplifies and enhances legacy server and application support, while providing new opportunities in disaster recovery and high-availability scenarios.
- Because virtual machines are much easier to deploy than physical servers and because the software nature of virtualization technology lends itself well to automation of provisioning, on-demand and adaptive computing are more easily realized.
- Server virtualization can be used in many ways to help solve business problems. Server consolidation can reduce data center costs while enabling the virtualized servers to be backed up in an easier manner helping with disaster recovery.
- The lifetime of legacy servers and applications can be significantly extended and the performance increased by moving them to virtual machines. Disaster recovery strategies can be enhanced through the use of server virtualization, which can dramatically reduce the time and complexity required to restore vital systems in the event of a catastrophic event, improving the chances of business continuance.
- Server virtualization can also help reduce the costs involved in creating highly-available systems, also aiding in business continuance through increased fault tolerance. In addition to cost reduction, server virtualization opens up new opportunities by providing the capabilities to easily

	create adaptive and on-demand computing systems increasing an organization's productivity. • Server virtualization is not the answer to all business problems, however, and some systems do not lend themselves well to be run within virtual machines, including performance-sensitive applications, graphics-intensive applications, and applications requiring specialized hardware. • It is important to take these limits into consideration when planning for server consolidation projects, legacy server and application support, disaster recovery, and high availability solutions. All in all, virtualization can save the organization thousands to millions of dollars through more efficient and effective use of hardware.
Q.28	**What is server consolidation? How can organizations save money through server consolidation?**
A.	• Server consolidation allows many physical servers to be virtualized and hosted on as little as a single physical server leveraging virtualization. Most servers today are far underutilized, running between 8 percent and 12 percent utilization. • Server consolidation means reducing the number of servers in your datacentre, creating larger servers that host multiple workloads. Many organizations today have several servers, each performing a dedicated task. • These servers include file servers, print servers, e-mail servers, Web servers, database servers, and other application servers. It is not uncommon to have applications that require or strongly recommend a dedicated server, usually because the software does not particularly coexist well with other applications installed within the same operating system. • In these situations, organizations will purchase small dedicated servers to host these applications, many of which have a low rate of utilization because the application only has a small amount of users or because it may only be used infrequently, such as once a month. • The problem with this scenario is that the organization must incur a capital expenditure of $6,000 or more and that the processor, memory, and disk storage of the new dedicated server are highly underutilized and essentially wasted. Over time, the data center may support many small, highly underutilized servers hosting dedicated applications, also incurring the data center hosting costs (ping, power, and pipe) for each server. • Server virtualization can help organizations save money in the datacenter through server consolidation. Many, if not all, of the organization's highly underutilized application servers may be moved from physical hardware into virtual machines. • Applications requiring a dedicated server still meet that requirement because a virtual machine is a dedicated server with its own, fully isolated virtual hardware and operating system. • Many of these virtual machines may be hosted on a single physical server by means of the virtualization platform. As new dedicated application servers are needed, instead of purchasing new physical servers, virtual machines may be created on existing virtualization host servers. • The density of virtual machines, which may be hosted on a single virtualization host server, depends on several factors including the number of processors, the amount of memory, and disk storage space on the physical host server, the virtualization platform (installed on the host server), and the amount of resources consumed by the virtual machines. It is not uncommon to have 10:1 or higher server consolidation ratios for highly underutilized servers.
Q.29	**Enumerate the reasons for existence of legacy system in an organization.**
A.	• A common pain point of many information technology organizations is that, over time, legacy

	hardware and software must be maintained. • The term legacy commonly refers to either out-dated and possibly unsupported computing systems or systems comprised of components of a previous version. Replacing legacy systems with new systems is always a challenge, both technically and in business terms. • Legacy systems exist because of many reasons. These may include a very high cost to completely replace the system due to a large investment in capital expenditure, training, and customization of the system. Legacy systems may also be sensitive, mission critical resources that cannot easily be replaced. • Sometimes legacy systems must be maintained because they cannot be replaced because it was purchased from a third-party company no longer in business or from one that no longer supports the system, or possibly because of a lack of expertise in the system's internals. • In some cases, it costs less to maintain legacy systems than it does to replace them, but over time legacy systems typically become more unusable because they often are not able to scale up because of costs or technology limitations. • For example, an organization may have invested in a custom software application designed to run under the Microsoft Windows NT 3.51 Server operating system many years ago. The application in this example was written in such a way that it will not work with any other version of the Windows operating system. The application is becoming a productivity bottleneck because its utilization has grown over the years while running on the same hardware. • The organization wishes to upgrade the server hardware on which the application and operating system are installed upon, but cannot because the newer hardware available today does not have the appropriate support for the older operating system. There are no available chipset and storage drivers that will work with the older operating system while allowing it to use top-of-the-line processors, memory, and disks to solve the legacy application's ability to scale up. • Additionally, because the application was custom-written, there are no newer versions of the application available and the costs to switch to a different application that provides like functionality is cost prohibitive. • Server virtualization can solve legacy server and application support issues. Migrating the legacy servers to virtual machines inherently abstracts the physical hardware from the legacy software, including the operating system and the applications, allowing the legacy hardware to be discarded or reused elsewhere as needed. • Because the legacy software now uses virtualized hardware, it can be moved to any host server as necessary, making the legacy server portable. Aside from its newly gained portability, the legacy server migrated to a virtual machine can be hosted on and use any hardware that is supported by the virtualization platform and the host server. • Continuing the previous example, the organization decides to migrate the legacy server to a virtual machine. They choose a virtualization platform that allows them to host the migrated server on modern, high-end server equipment, which addresses the bottleneck issues while adding the portability necessary to move the server to better hardware in the future, if the need arises. • Additional savings in terms of hosting can be realized by migrating legacy servers to virtual servers in order to decommission older, inefficient hardware and by increasing overall server utilization.
Q.30	**What is disaster recovery? What is high availability? How can virtualization help in disaster recovery and high availability?**

A

Disaster recovery

- Disaster recovery is a term used to describe the strategy and processes used by an information technology organization to ensure that in the event of a disaster that damages or destroys the organization's computing infrastructure, a minimal recovery of the computing infrastructure and data can be accomplished in a reasonable amount of time to provide business continuance.
- Disaster recovery strategies are almost always an afterthought in most organizations, probably because it is like paying for insurance. When an event occurs, however, it is the disaster recovery strategy that can save a company from going out of business, thus providing business continuance.
- Although disaster recovery strategies encompass a wide range of considerations, one of the most common activities is data backup. Backing up the data of an organization is critically important to help not only with disaster recovery, but also to help to recover from day to day mistakes, such as when a user inadvertently deletes an important document, or to help protect data lost from a malicious intruder.
- There are many theories and methods of performing data backup. Most commonly, only an application's data is backed up instead of the application itself or the operating system. This is typically done to save space within the backup system (and to reduce the cost of the backup system) because the application's code and the operating system itself changes rarely.
- It is also traditionally more difficult to restore an operating system and applications from backups versus reinstalling them and then applying data afterwards. Because of these methods, the restoration of systems becomes a much harder, longer task.
- Although complete system restoration is rare (hopefully), the amount of time it takes to restore systems is critical to business continuance. Although disaster recovery strategies address business continuance in the event of a catastrophic disaster, they do not address the needs of fault tolerance and system robustness.

High-availability

- High-availability strategies provide business continuance through the use of fault-tolerant systems, usually by implementing redundant system components such as RAID storage systems or clustered servers.
- Most highly available systems must be available more than 99 percent of the time, meaning that the systems must be able to recover from a number of different component failures. One of the most common methods of providing a highly available system is through the use of clustered servers.
- In a server cluster, there are at least two servers that have the same applications and configuration and share access to the same data. When the active server in the cluster experiences a fault (such as a failed network adapter or hard drive failure), one of the other servers in the cluster take over and become the active server in order to provide nonstop services.
- One of the issues associated with server clusters is the additional cost of server hardware that is mostly unused while it is waiting for a failure to occur. Server virtualization can provide an inexpensive method of implementing server clusters by using virtual machines in a cluster.
- Virtual machines can be clustered with other virtual machines or other physical servers. Using virtual machines in server clusters works well with server consolidation. While the primary server in the cluster is active, the secondary server, a virtual machine, will usually consume very few resources, allowing it to be placed on a host server along with other servers.
- If the primary server in a server cluster is also a virtual machine, it should be placed on a separate physical host server in order to maximize the cluster's uptime.

Q.31	**What is adaptive computing? What is on-demand computing? Explain**
A.	**Adaptive computing** • Adaptive computing consists of server systems that have the ability to autonomously reconfigure themselves to address changing requirements. Adaptive computing is also referred to as autonomous computing, grid computing, on-demand computing, or utility computing. • Server virtualization can work well with adaptive computing initiatives because of the ease of virtual machine provisioning. • For example, consider a bank of Web application servers in a load balanced cluster. The Web application utilization rises and the overall performance of the application decreases. The system then allocates resources on a virtualization host server on which it creates two additional virtual machines using the same Web application server image. • Once the two virtual machines have been created they are booted up and added dynamically to the existing cluster. The two additional servers help spread the application's workload over more computing resources, thus increasing the overall application performance. • When the application's utilization falls off, the two additional servers are no longer needed and they are powered off and deleted. This type of adaptive computing can be applied to many applications that share a common set of virtualization host server resources on which to dynamically create virtual machines. • In addition to dynamically responding to needs, adaptive computing systems can have capacity scheduled in order to help optimize computing resource utilization. For instance, during the week an application may have five virtual machines on which to perform its work, but over the weekend, three of the virtual machines may be scheduled to be reconfigured to work with a different application to help with back-end processing. • Adaptive computing scenarios can be achieved with physical hardware, including traditional server and blade servers, but typically at an increased cost and increased level of complexity as compared to using server virtualization. **On-Demand Computing** • Virtualization technology easily facilitates on-demand computing systems that can quickly allocate one or more servers to a group of users for a short time and reclaim and reuse those resources once the users' time has expired or if they have released the resources. • This type of system is commonly used to provide virtual labs used by system administrators and developers for testing software configurations. It is also used to provide virtual training labs within a learning management system, whether it is a physical class or an e-learning system accessed over the Internet. • Once again, the ease of provisioning virtual machines supports on-demand systems very well, but it is the virtual machine's ability to discard changes made to their virtual hard disk drives, resetting them to a pristine state, where the most value is gained. Building these types of systems without using virtualization is very difficult and expensive.
Q.32	**What are the limitations of server virtualization? Explain.**
A.	• Virtualization technology opens the data center to new possibilities that may provide cost savings and new types of functionality, but virtualization does have its limits. Not every server or application is well suited to be run in a virtual machine. • Some applications are highly performance-sensitive, such as databases, data warehousing

	applications, business intelligence, reporting, and many others. These applications usually require multiple processors and massive amounts of memory. • At the present time, server virtualization platforms support virtual machines with either a single processor or at most dual processors and up to 3.6GB of memory. In the near future, these limitations will likely be overcome, although for performance-sensitive applications, it may still not be enough to justify running them as virtual machines. • This is because virtual machines are not only sharing resources with other virtual machines on the same host server, but because the virtualization platform itself incurs a small amount of overhead that can negatively impact performance-sensitive applications. • Virtualization technology does not support graphics-intensive applications very well at this time. Graphics-intensive applications and games usually require the use of high-performance video cards. In the virtual machine, the video card is a virtual device implemented in software and at this time does not support advanced graphics acceleration features. • Even if the virtual video card device did have support for advanced graphics acceleration, it would invariably place more overhead on the system's physical processors. Much like the problems with graphics-intensive applications, applications requiring specialized hardware cannot be used within a virtual machine. • Specialized hardware refers to any nonstandard hardware device, commonly PCI cards. At the present time, virtualization platforms are generally closed systems and new types of virtual hardware cannot be created by third-parties. • In the future, this limitation may be overcome, at which time third-parties may have the ability to write their own virtual hardware devices in order to connect specialized physical hardware devices to virtual machines.
Q.33	**List and explain the considerations to be taken into account while evaluation server virtualization platforms.**
A.	• Below is a list of considerations to keep in mind when evaluating server virtualization platforms. 1. **Hardware Compatibility:** When considering server virtualization, it is important to understand that not all x86 server virtualization platforms will work with all x86 server hardware. Generally, MSVS and GSX Server will work with nearly any hardware that is supported by the host server's operating system, which can be Windows for MSVS and GSX Server or Linux for GSX Server. 2. 2. **Host Operating System**: Support When MSVS or GSX Server is being considered, it is important to know exactly what host operating system will be used to support the server virtualization platform. It is not just a simple matter of choosing Linux or Windows, but of knowing which version, edition, and revisions may be required to be successful. When considering GSX Server for Linux, it is important to know which Linux distributions and kernel versions are compatible for use as a host operating system. 3.**Guest Operating System**: Support Although most server virtualization platforms can run most 32-bit, IA-32 (or x86) and some 64-bit, x64 operating systems in a virtual machine, the server virtualization platform vendors do not support every possible guest operating system. In some cases, such as MSVS, only a subset of the platform vendor's Windows operating system is officially supported. It is very prudent to not use unsupported configurations in a production environment to eliminate risk, even when the given configuration may work most of the time. 4. **Application Support**: Depending on the use case of the server virtualization system being deployed, the application support topic may or may not be a large issue. When virtual machines are planned on being used in a production environment, it is wise to find out if the intended

	production applications will be supported for production use by the server virtualization platform vendor or the vendor of the application. 5. **Capacity**: Server virtualization platforms also differ in the capacities they support. Capacity metrics including the number of supported processors, the minimum and maximum amount of memory, the maximum size of virtual hard disks, and the maximum number of supported virtual machines will affect the overall VM density of the system. 6. **Scalability**: Server virtualization systems are able to scale out very well. By adding more host servers, the overall virtual machine capacity is increased. But in terms of cost savings, it may be beneficial to scale each host server up rather than out. 7. **Features**: The various features that are supported by the different server virtualization platforms can also affect the decision making process. 8. **Stability**: Thee stability and robustness of the server virtualization platform can also play a large role in the decision making process, especially for mission critical production systems. The level of maturity in each server virtualization platform should be considered as well as the release history and the product's overall longevity. 8. **Maintenance and Uptime**: Another factor to consider is how the platform is maintained in regard to product updates such as security patches and product fixes. It is not uncommon for server virtualization platforms hosted on top of an existing operating system such as MSVS and GSX Server to require critical updates on a regular basis. 9. **Performance**: The performance of the server virtualization platforms is a very difficult metric to use to make a decision upon. Published reviews and benchmarks may be biased or may not even apply to the specific use case of the solution planning to be implemented. 10. **Ease of Use**: Another aspect to be aware of when comparing server virtualization platforms is the ease of use. Some platforms are very easy to learn and use, typically MSVS and GSX Server because they are hosted on an already familiar host operating system such as Windows or Linux. 11. **Manageability**: The available methods of configuring and managing the server virtualization hosts should also be taken into consideration when considering virtualization platforms. GSX Server and MSVS require a host operating system, such as Linux or Windows. 12. **Support**: The levels of available support from the server virtualization platform vendor may also be considered when comparing virtualization platforms. Terms of support agreements vary and some vendors may actually offer a limited amount of support and offer extended support facilities at an extra cost. 13. **Training**: The availability of quality training is another consideration. If an organization does not possess the proper skill sets to deploy or manage a server virtualization system, it may need to have its employees trained. The availability, quality, and costs involved in obtaining proper training on the various server virtualization platforms could have a significant impact on the overall decision. 14. **Cost**: Finally, most decisions almost always come down to the bottom line, the cost. The available server virtualization platforms diff er greatly in cost. Some such as MSVS and GSX Server also incur the cost for an additional operating system license used by the host server.
Q.34	**Explain the 3 C's while selecting the hardware necessary for server virtualization deployment.**
A.	• Selecting the hardware necessary for a server virtualization deployment may seem easy at first. But, after digging into the details, it soon becomes evident that there are many factors at work. • The difficulty lies in balancing the **cost, capabilities, and compatibility**, referred to as the 3 C's.

	• Cost and required capabilities should be referenced in the Use Case and Requirements documents. Compatibility is often a derivative of the selected server virtualization platform. • Before attempting to select the server hardware, it is important to know exactly what server virtualization platform will be used and what host operating system will be used, if applicable. It becomes easy to take the server architecture into the wrong direction by losing focus on any one of the 3 C's. • It is equally important to select server hardware that is compatible with the selected server virtualization platform and host operating system. This is usually less of an issue when the host operating system is Microsoft Windows because of the large availability of hardware drivers for the Windows operating system. • When selecting server hardware components for server virtualization hosts, it is important to know that a minimum of two network adapters is usually required. One network adapter is bound to the host server operating system and the other is bound to one or more virtual machines. • Virtual machines can share a single network adapter, but it is generally a bad idea to share the network adapter used by the host server with the virtual machines. Depending on the speed of the network adapters selected and the number of virtual machines that will be placed on a single host, it may be necessary to have more than one network card reserved for virtual machines. • If anyone virtual machine has any network bandwidth-intensive applications installed, it may be necessary to reserve one network adapter just for that virtual machine. It is not uncommon for host servers to have three or more network adapter cards installed, even two or three dual or quad port Ethernet server adapters. • When using Ethernet server adapters with more than one port, each port usually acts as a separate network adapter and therefore the host can use one port while the other ports may be used for virtual machines. • Another network adapter consideration is using load balancing, fault-tolerant network adapters and binding more than one network adapter or port together in a team to increase bandwidth and fault-tolerance. In this case, a minimum of two ports would be required for each teamed network adapter. • It is possible to use teamed network adapters for the host or the virtual machines. Teamed network adapters require explicit support from both the physical network adapter cards being used and the drivers installed into the host operating system.
Q.35	**Discuss the role of host servers in virtualization solution.**
A.	• The host server is probably the most familiar component. The host server is the physical, bare metal, x86 server computer on which the virtualization layer resides. If the host server has an operating system installed, the operating system is referred to as the host operating system or host OS. • The host server is referred to as a "host" because it hosts the virtualization layer and one or more guest servers or virtual machines. It is highly recommended that the role of host servers in a server virtualization solution be completely dedicated to server virtualization. • This is implicitly chosen for ESX Server deployments, but it may not be obvious to new users of MSVS or GSX Server. It may be tempting to try to purpose a Windows or Linux-based virtualization host server as a dual or multi-role server, handling file sharing, print serving, a Web application, or a database server alongside server virtualization. • Doing this will ultimately cause more server resources to be consumed by the host server

	operating systems and its applications, which will severely limit the number of active virtual machines as well as decrease their performance. • Virtualization host servers should be dedicated to the role of hosting virtual machines. It may be advantageous to place the necessary services on a server running in a virtual machine on the host server, depending on the application or services resource or performance needs. • Host servers should have their configuration minimally optimized by removing unneeded software packages and operating system components, disabling unneeded services, and by turning off unneeded operating system features, including enhanced graphical display features.
Q.36	**What are the common licensing schemes? How can they be adapted to virtualization?**
A.	• The rise of mainstream virtualization technology, such as x86 server virtualization, is causing confusion today regarding operating system and software licensing. Virtualization technology is causing a major paradigm shift in software licensing schemes, for better or worse. • Most software vendors have not yet attempted to adjust their licensing schemes to account for their software running inside of virtual machines. Some vendors are not even aware of how server virtualization might impact their current licensing schemes, not accounting for the architectural and philosophical changes that are brought about by virtualization technology. • Because of this lack of visibility into virtualization, there are some interpretations that must be made in order to stay in compliance with software licenses. It is highly recommended to contact the software company and to request an exception or written interpretation as to how their licensing scheme is affected or not affected by the fact that their software will be running inside virtual machines. • However, this may not always be practical. When it is not practical, there are some simple things that can be done to attempt to stay in compliance in most situations. 1. **Instance-based licensing**: Instance-based licensing is the most common type of software licensing scheme. Each time the software is installed, it requires a license. The consumer will purchase a fixed number of licenses (or seats) and the software may legally be installed on the same number of computers, whether physical or virtual. This is one of the easiest to manage licensing schemes and it works the same way on virtual machines as it does with physical computers. 2. **Computer-based licensing**: Computer-based licensing allows software to be installed more than once on a single computer, possibly even limitless instances, but only requires one license per computer. In some cases, it may be interpreted that some computer-based licensing schemes may allow the software to be installed an unlimited amount of times within multiple virtual machines on a given host and only have to acquire one license per host server. 3. **Processor-based licensing**: Processor-based licensing schemes incur a cost for each processor or a different cost based upon the number of processors installed in the computer. Depending upon the exact verbiage of the license agreement, the license may only apply to the number of physical processors installed in the computer or only to the number of processors available to the virtual machine. The licensing costs could also apply once per physical server or virtual server depending on how the license is written. This scenario is one of the most confusing of all licensing schemes as it could work in favour of the vendor or the consumer depending on the interpretation of the license. 4. **Fixed user-based licensing**: Some software is licensed by an amount of fixed or named users that will access the software. In this licensing scheme, one license is purchased for each exact user of the software, regardless of whether they are all using the software at the same time or

	not. This licensing scheme works the same way in virtual machines as it does on physical computers. 5. **Concurrent user-based licensing**: In the concurrent user-based licensing scheme, software licenses are required for the total number of simultaneous users connecting to or using the software. This licensing scheme is much more flexible than the fixed user-based licensing scheme. Concurrent user-based licensing works the same way in virtualization as it does on physical computers. 6. **Device connection-based licensing**: In device connection-based licensing schemes, a license is required for every device that connects to the software. This is usually found in enterprise server applications. The term device usually refers to any user, computer, or other device or account that connects to the software.
Q.37	**Discuss the supported and unsupported operating systems in VMware GSX, VMware ESX and Microsoft virtual Servers.**
A.	• The issue of running guest operating systems not officially supported by the server virtualization platform vendor sparks a lot of controversy. It is important to understand the implications of using software such as server virtualization in unsupported configurations, especially when considering doing this in a production system. • Depending on an organization's level of comfort with the officially unsupported guest operating system and the virtualization platform and depending on the use case of the server virtualization implementation, it may be well justified by the organization to absorb the amount risk associated in using an unsupported configuration. • Currently, GSX Server has official support for the widest range of guest operating systems, followed by ESX Server, followed by MSVS. What does having official support for a guest operating system from the virtualization platform vendor really mean? It does not necessarily mean that an unsupported guest operating system will not install and run properly in a virtual machine. • There are many unsupported IA-32 (or x86) guest operating systems that will run as expected within the available server virtualization platforms. Primarily, it means that the virtualization platform vendor will not be able to provide assistance in the advent of a malfunction, bug, or other issue that could arise. • The vendor will not assume responsibility or risk in these cases and will generally not spend time attempting to support such configurations, even if it is truly caused by a fault in their software. It costs a large amount of money to provide product support, and the vendors cannot support every possible configuration, just as a hardware provider cannot support every conceivable operating system and configuration for their systems. • Second, the guest enhancement software provided by the virtualization platform vendor (MSVS has Virtual Server Additions and ESX Server and GSX Server have VMware Tools) is likely not available for unsupported operating systems. • The existing Linux versions of these add-in software tools may not work properly in other Linux distributions or versions. The guest enhancement software usually consists of a set of optimized device drivers designed to help the guest operating system perform better with the specific set of virtual hardware exposed by the virtual machine as well as a background service that manages such things as communications with the host server, VM heartbeats, and time synchronization, among other features. • In general, unsupported guest operating systems will run slower than a supported operating

	system with the guest enhancement software installed. There will usually be performance and usability issues when attempting to access virtual machines that have an unsupported guest operating system installed using the virtualization platform's native remote control technology, such as Virtual Machine Remote Control in MSVS and the VMware Remote Console in ESX Server and GSX Server.
Q.38	**What are the Various support issues from vendors in virtualization? Explain**
A.	• Virtualization is a technology that has been around for quite some time, yet it is really just now becoming mainstream. Because of this, many people in the information technology industry, including many software vendors, are not virtualization understanding. • Support issues can arise from mixing virtualization with software applications from vendors that have little or no virtualization experience. Some software vendors will not support their products if their products are installed and used in virtual machines. • The basic argument that these vendors will use is that there are too many unknown variables in the way that the virtualization technology affects the guest operating system and other software applications installed within a virtual machine. • Of course, it may not be stated in such terms, but that is by far the most common argument. To the vendors, the virtual machine is an unknown, black box into which they have no visibility; therefore they often do not wish to support it. • These issues will usually come to light when a legitimate support call is made to the vendor and the vendor's support analysis discovers the use of virtualization or when a vendor sends out a consultant to implement their software and the consultant discovers that they will be installing software into a virtual machine. • When a legitimate support issue arises, it most often has nothing to do with the virtualization technology. Even so, many vendors will not be so easily convinced. They may completely refuse to support the software or they may be reasonable and will want the issue to be reproduced using physical hardware. With some vendors, it will be easy to negotiate a reasonable resolution when working with virtualization and with others there may be no compromise. • When placing legitimate support calls to a software vendor, it may be beneficial not to volunteer the fact that their software is running within a virtual machine. Only if the vendor queries as to what server equipment their software is being used on is it recommended to let them know that it is installed within a virtual machine. • At that time, it is also recommended to give them information regarding the virtualization platform, version, and possibly even the specification of the physical server hardware. Sometimes, a more reasonable vendor may even get the virtualization platform vendor's support organization involved to help them solve the issue. • As server virtualization continues to grow as a mainstream technology that is present in most data centers, more and more software vendors will also catch on and begin supporting their products within virtual machines. This may actually become a necessity as the software vendors' competition may use their support of virtualization as marketing leverage. • The outcome is a positive and natural evolution of the industry, which will further catapult virtualization technology's growth.
Q.39	**What are the reasons for unexpected server growth in server virtualization? How can this growth be combated?**

A.	• A common side effect that often takes place in server virtualization systems is unexpected server growth. Server growth is the unexpected addition of unplanned virtual machines. In some ways, this can be thought of as virtual server sprawl. • It primarily occurs after a server virtualization system has been deployed in the data center and the need arises to position new servers. Server virtualization makes it very easy to deploy new servers in a very short amount of time. • This is not necessarily bad unless it negatively impacts the existing virtual machines. Depending on the use case for the initial server virtualization deployment, it could be a very bad situation, such as creating normal enterprise application server virtual machines on a server virtualization system that was built for on demand computing. • In this situation, the application server virtual machines are essentially diminishing the overall capacity of the on-demand server system. Unexpected server growth not only impacts the production system, but can also increase operating system and software licensing costs. • Unexpected server growth also occurs out of the natural need for a development and test server virtualization environment. • Architects, developers, system administrators, and system engineers will always have a need for a disconnected, nonproduction server virtualization system in which they can design, develop, and test new methods of using the virtualization platform features, new server and network layouts, test new operating systems and software, construct virtual machine base images, and perform other nonproduction tasks and experiments without disturbing or impacting the production environment. • The size and scope of the nonproduction server virtualization environment that is needed can vary greatly depending on its planed usage, but as a base line, the host servers should generally be the exact same configuration of the production host servers in order to provide an minimally adequate nonproduction test system deployed alongside the production system. • Additionally, the development and test server virtualization environment should be separated from the production network either physically or by a firewall and it should have a dedicated switch to provide the interconnects between host servers. • The following guidelines can help combat unexpected server growth in production systems: Include an adequate development and test server virtualization environment in the deployment plans if there is not an existing facility. • Ensure that everyone that has direct access to the server virtualization host servers fully understands and adheres to the use case. If there is a need for more than one use case for server virtualization technologies, plan for each use case accordingly. Plan for about 20 percent more capacity than what is initially required to handle immediate growth.
Q.40	**What is virtual machine density? Enumerate the factors to consider estimating virtual machine density.**
A.	• The number of virtual machines residing on a single host server is referred to as virtual machine density, or VM density. When a high VM density is achieved, overall costs are generally lowered by sharing and better utilizing the host server's resources. • Many factors must be considered to properly estimate VM density. 1. The operating system and virtualization platform installed on each host server 2. The number of processors in each host server 3. The amount of available memory in each host server 4. The amount of available disk storage space on each host server

	5. The guest operating system and applications installed on each virtual machine 6. The number of virtual processors configured for each virtual machine 7. The amount of memory allocated to each virtual machine 8. The size of all virtual hard disks of each virtual machine 9. The amount of idle processor consumption of each virtual machine. 10. The expected utilization of each virtual machine 11. The acceptable performance of each virtual machine 12. The expected usage of each virtual machine • The operating system and virtualization platform installed on each host server defines the overall environment in which the virtual machines will reside. Together, they set the initial boundaries of the environment by their supported capabilities and features. • For example, if each host server will have Microsoft Windows Server 2003 Standard Edition as its operating system and Microsoft Virtual Server 2005 Standard Edition as the virtualization platform, the number of physical processors is limited to four, the maximum amount of physical memory is limited to 8GB, the maximum number of virtual processors is limited to one per virtual machine, the maximum amount of memory per virtual machine is 3.6GB, and no more than sixty-four virtual machines may be created on each host. • Furthermore, if each virtual machine is expected to have 1GB of memory, a maximum of seven virtual machines per host is possible because the host operating system consumes some of the 8GB of maximum memory. • In this example, 8GB is the maximum amount of memory because that is the maximum amount of memory supported by the version and edition of the Windows operating system being used. It is very important to be aware of these types of limitations early in the planning process. • The number of processors installed in each host server defines the maximum amount of processing power and parallelism capabilities. Some server virtualization platforms also restrict the maximum number of virtual machines per physical processor installed in the host server. • A general rule of thumb is that the more physical processors installed, more virtual machines can be created. In reality, this is only true to a certain point, especially when considering very generic virtual machines with low utilization and low performance metrics. • It is also recommended to have a minimum of two processors in each host server in order to achieve good overall system performance.
Q.41	**What are the data and system backup considerations to be taken into account during server virtualization?**
A.	• Adding new systems into existing data centers will always have an impact on existing data center management systems and the data backup systems are no exception. Server virtualization tends to greatly amplify this impact by potentially adding another level of complexity and larger than usual demand on existing data backup systems and strategies. • In some cases, server virtualization systems may not easily fit into existing data and system backup processes. The traditional concerns are the costs imposed by possibly needing additional backup system agents to install on the host server or the virtual machines, the storage capacity of the backup system, the backup time capacity of the backup system, and costs incurred by the need for additional persistent media and off site storage. • Besides the traditional backup issues, server virtualization will impose two dimensions of back up: backing up the host server and backing up the virtual machines residing on that host. System backup for virtualization host servers is usually much like system backup for traditional physical

	servers, especially when dealing with server virtualization platforms requiring a host operating system. • In addition to the core operating system, the virtual machine configuration and log files should be backed up. The virtual machines should always be backed up as well, but there are several different strategies regarding virtual machine backup that can have a huge impact on storage capacity, costs, and time. • ESX Server can more of a challenge when considering host server backup. Many existing backup agents are not compatible with ESX Server and its special VMFS fi le system, although new solutions are beginning to appear. In many cases, system administrators create custom backup scripts using shell scripting or Perl scripts that are scheduled using cron within the ESX Server Service Console. • When considering virtual machine back up, there are two major methodologies that are commonly used. One method is to back up each virtual machine's virtual hard disk fi les that reside of the host server on a scheduled basis, such as daily or weekly. • These fi les tend to be large, multi-gigabyte fi les. Virtual machine disk fi le sizes commonly range from 2GB to over 100GB. It is very expensive to back up every virtual machine daily not only because of the massive amount of storage capacity that is needed, but because, in many cases, the virtual machines must be powered off in order to back them up. • This is not totally true for every virtualization platform and every configuration, but it is the most reliable way to back up a virtual machine's disk fi les. In some configurations and virtualization platforms, such as ESX Server, it is possible to back up a live virtual machine. • Because the fi le sizes are large, the backups can take a long time to complete, especially if the network bandwidth between the host server and the backup storage devices is slow. Using a high-performance SAN can solve this problem, but it is very expensive in terms of hardware, software, and services costs. • Backing up each virtual machine's disk fi les provides a best-case scenario when the virtual machines must be restored. Depending of the level of importance of the virtual machines, the costs and time it takes to back up entire virtual machines daily or weekly may be well worth it after experiencing a system failure causing the backed up virtual machines to be restored. • It is also important to note that highly available virtual machines that are clustered make it easier to back up each virtual machine in the cluster because the services provided by the cluster will continue to run while the other nodes are backing up. • If the virtual machine ever needs to be restored, the virtual machine's disk fi les are restored to the host server first, which quickly establishes the base virtual machine and then operating system and data backups can be applied within the virtual machine using normal data restoration techniques.
Q.42	**What are the monitoring and the performance backup considerations to be taken into account during server virtualization?**
A.	**Monitoring Considerations** • Most established data centers will often have a centralized server monitoring system in place to keep track of the health and status of servers and applications. Once server virtualization becomes a part of the data center, it will likely affect the existing server monitoring system. • The server monitoring system will probably be able to monitor the host servers without any problems. At this time, many server monitoring solutions do not have any capabilities for monitoring virtual machines within the context of the virtualization platform.

	• They may be able to monitor virtual machines in a more limited method by installing the monitoring agents into each virtual machine, but the monitoring system would not know the virtual machines apart from the physical servers. • Some monitoring systems are just starting to produce add-ins that understand and support monitoring on some virtualization platforms. As time goes on and virtualization technology continues to grow inside the data center, better monitoring solutions will become available for virtualization. • Another way server virtualization may affect existing server monitoring systems is by the increased cost of monitoring the host servers and all of the virtual machines. Most monitoring systems are licensed by the capacity of servers they monitor. Once server virtualization is introduced into the data center, the number of servers that need to be monitored will tend to grow quickly, which will likely increase server monitoring costs in terms of requiring more licenses. • It is a good idea to be aware of this issue and to plan ahead for the costs of additional server monitoring. **Performance Considerations** • Performance can be very difficult to estimate or measure in advance when planning a new server virtualization deployment. It is always better to make conservative performance estimates, especially if performance is a sensitive issue. • To properly size for performance, real-world testing and benchmarking must be performed. Real-world testing and benchmarking refers to creating a test environment that is identical to that of the planned production environment with the exception of the scale. • The hardware and software must be installed and configured exactly as in a production system in order to obtain useful performance metrics. It is highly recommended that evaluation equipment be acquired, even for short evaluation-only terms, such as 30 days, so that in-house testing and benchmarking can be completed. • For some organizations, this may not be possible. Obtaining third-party white papers with similar benchmarking using similar equipment and server virtualization platforms may be the only option for some organizations on which to base their planning. • In this case, many assumptions are being made that could invariably be wrong; therefore there is much higher risk in using this option of gauging the performance of the use case. • Ensuring adequate performance is often important because if the applications running on the server virtualization systems do not perform well, the consumers of those applications may deem the deployment or the technology a failure, even when the systems are executing as planned. • If applications are planned to be migrated from an existing physical server environment to a server virtualization environment, performance testing is imperative in order to ensure that the final production configuration of the virtualized system meets or exceeds the existing expected performance levels.
Q.43	**What are the processor and memory considerations to be taken into account during server virtualization?**
A.	**Processor** • The processor is one of the most common performance bottlenecks in server virtualization and requires a great deal of attention. Processor cycles are a precious commodity due to their finite compute power. • Server virtualization platforms usually offer tools which can help manage processor resource

	usage across virtual machines, such as limiting a virtual machine's processor usage or giving another virtual machine a minimum guaranteed amount of processor resources. • The virtual machine layout across host servers in conjunction with the use case plays a major role in how processor usage is affected, which in turn affects performance. For instance, consider a training scenario where there are five host servers each hosting ten virtual machines. Each virtual machine uses an identical image, so they are all exactly the same. • This setup will allow 50 users to be trained concurrently, each having a single virtual machined assigned. Through the course of the training, the students are all performing the same tasks at the same time. • This can severely impact processor performance, since the processor load will likely be very high on each host because each virtual machine is asking for the same amount of processor resources simultaneously. • The exact type of applications being used within the virtual machines will also play a role in the amount of processor resources being used, but the use case amplifies the effect. The solution to bad performance in this scenario may be to reduce the VM density, which will increase the number of hosts needed. • The optimum VM density must be determined through testing. It may be five or even as few as three virtual machines per host. **Memory** • Memory-based performance bottlenecks can occur for different reasons. One of the most common causes of memory-based bottlenecks is a virtual machine that does not have enough memory allocated. • This typically causes the virtual machine to have a very high amount of page faults, swapping data between memory and the swap file. This is usually solved by adding memory to the virtual machine. • Sometimes memory cannot be added to a virtual machine because it has already used the maximum amount of memory assignable by the server virtualization platform, currently 3.6GB. • If a virtual machine's performance is severely impacted by memory constraints and it is using the maximum amount of memory assignable to a virtual machine, that server is most likely not a good candidate for server virtualization. • To correct the issue, the server's software components could either be distributed among several virtual machines or the server may not be a good candidate for server virtualization. Memory performance issues can also be caused by faulty memory modules. • It is recommended to run server diagnostics routines on each new server for at least 24 hours prior to use in order to fully test the hardware for faults.
Q.44	**What are the disk I/O considerations to be taken into account during server virtualization?**
A.	**Disk I/O** • Disk I/O bottlenecks are caused by the fact there are only a limited number of I/O operations that can occur on a disk subsystem within a given period of time. • When several virtual machines' virtual hard disk fi les are hosted on the same disk subsystem of the host server and they are placing a heavy load on their virtual disks, overall I/O performance will degrade system wide. • Virtual machines with consistently high I/O volumes should either be placed on a host with low volume I/O virtual machines or fewer virtual machines. Another option is to place high-volume

	virtual machines' virtual hard disk fi les onto their own dedicated disk on the host server, which also has a dedicated disk controller channel to the physical disk controller. • If a virtual machine's I/O demands are too great, it may be placed on a host with very few or no other virtual machines. In some cases, extremely high I/O demands may deem that a particular application is not a good candidate for server virtualization, but this is very rare. • Additionally, disk I/O resource controls may be available in the server virtualization platform. In this case, it may necessary to allocate more disk I/O resources to a single virtual machine using the platform's resource controls.
Q.45	**What are the network and security backup considerations to be taken into account during server virtualization?**
A.	**Network** • Network performance issues usually arise in server virtualization because many virtual machines typically share a single network adapter. This is especially true when using a Fast Ethernet (10/100) network. • Most modern server network adapters today use Gigabit Ethernet (10/100/1000). Gigabit Ethernet is preferred when planning server virtualization systems, but the network adapter, cabling, and switches must all support Gigabit Ethernet in order to take advantage of the increased bandwidth. • Although Gigabit Ethernet provides a large amount of bandwidth, there are times when too many virtual machines are sharing a single network adapter which may cause a networking bottleneck. • It is recommended to dedicate an entire network adapter to virtual machines running network bandwidth-intensive applications and to share a separate network adapter for other virtual machines on the same host server that have medium to low network bandwidth utilization. • Even though most servers have at least two Gigabit Ethernet adapter ports, it is recommended to have at least three to four Gigabit Ethernet adapter ports available on each host server. This is fairly easy to accommodate as two or four port server network adapters are readily available which only consume one PCI slot in the server. • Depending on the use case, it is not uncommon to have six or eight Gigabit Ethernet adapter ports on a host server. For each network adapter port on each server that will be used, a corresponding switch port must also be available. **Security Considerations** • Security is a very important consideration in the deployment plan. Aside from the additional costs that could be incurred from the needs for additional security software licenses such as antivirus client licenses for each host server and virtual machine, there are other considerations that could affect the overall architecture of the system. • The exposure of any host servers or virtual machines to the Internet is one of the most primary security concerns. It is highly recommended that host servers and virtualization platform management software are not placed on an Internet facing connection. • Host servers should always be placed behind a firewall on an internal network and should never be placed in a DMZ or directly connected to the Internet. • The reason for the recommendation is that if a host server is compromised, all of the virtual machines on that are host can easily be compromised or at the very least powered off and deleted. • There are techniques that can be used to harden the host server's operating system, and this should be done, but even using a security-optimized, hardened system is no substitute for good

	perimeter security and access control architectures and practices. • If host servers should need to be managed remotely, a reliable VPN connection should be used.
Q.46	**Give an outline of USE CASE document for server virtualization solution deployment. Enlist the questions which can help gather important to build the Use Case Document.**
A.	• The use case is the critical first step in building a solid foundation for a successful server virtualization solution deployment. It defines how and why a particular technology and solution, in this case server virtualization, will be used. • Most projects do not follow a formal process for creating a use case and this can cause many problems throughout the deployment. Creating the Use Case Document ensures alignment of stake holders in the project, business and technical, which ensures that a common goal and vision is understood and agreed upon by all. • It can also be used to show key business stakeholders that a problem exists and to propose the intended solution. The Use Case Document also serves as the foundation of the Requirements Document, which is explained later. • The Use Case Document should identify the specific business and technical stakeholders, describe in detail the problem being addressed, the proposed solution to the problem, and other details of the proposed solution such as impacts, costs, and returns. Use Case Document Outline • Below is a sample list of interview questions which will help gather important information which will be used to build the Use Case Document. 1. What is driving the necessity of the changes? 2. What is the desired outcome of the changes? 3. What is the expected life expectancy of the proposed solution? 4. What existing systems will be impacted by the changes? • Administration and Management • Data and System Backup • Disaster Recovery • Intranet • Monitoring • Networking • Security • Storage 5. What existing processes or job functions will be impacted by the changes? • Business • Development • Education/Training • Helpdesk/Support Staff • System Administration • Testing/QA 6. What business or technical problems might be anticipated with the changes? 7. What business constraints may impact the proposed solution? • Budget • Time • Human Resources 8. What technical constraints may impact the proposed solution? • Hardware limitations

	• Limited data center floor space • Network infrastructure limitations • Operating system restrictions • Inadequate skill sets and knowledge
Q.47	**Compare Microsoft Virtual Server, VMware GSX and VMware GSX server virtualization technologies based on the following areas: Software maturity, Host and guest operating system support, Training and certification, Management and user Interface, Ease of creating virtual machines, Hardware support.**
A.	• **Software Maturity (VMware):** Where VMware stands out is in their product maturity. They are considered to have released the first commercially available x 86 virtualization platforms to market back in February of 1999. So, they have been at this for quite some time. VMware GSX Server is a third-generation product with 1.0 having been released near the end of 2000. VMware GSX Server is currently at version 3.2.1. VMware ESX Server has been shipping for well over four years now and is currently at version 2.5.2. VMware products have not only been time tested; the company also claims to have well over two million registered users and an installation base in over 5,000 corporations. On the other hand, Microsoft's Virtual Server 2005 product is a 1.0 release that was released around October of 2004. While that may not seem like a very long time, it is important to remember that Microsoft purchased the product from Connectix, which had a release candidate of the product in early 2003. Microsoft's current version, Microsoft Virtual Server 2005 R2, was released to manufacturing on November 15, 2005. Within a short time frame, Microsoft has quickly gained acceptance and adoption of its virtualization products. • **Host and Guest Operating System Support (VMware GSX Server)**: VMware GSX Server by far has the largest matrix of supported host and guest operating systems. ESX Server is its own host operating system as it installs on bare metal. Its guest operating system support is also impressive, supporting a number of Microsoft and Linux derivatives. Microsoft Virtual Server 2005 on the other hand, only supports a small set of Microsoft operating systems for the host, and a small set of guest operating systems made up of entirely Microsoft operating systems. Microsoft Virtual Server 2005 R2 adds support for several 64-bit Microsoft x64 guest operating systems as well as support for Microsoft Windows Server 2003 SP1 (all editions) and Microsoft Windows Server 2003 R2 (all editions), in 32-bit and 64-bit (x64) versions as host and guest operating systems. Although Microsoft in April 2005 announced future plans to support additional third-party guest operating systems on Microsoft Virtual Server 2005 R2, the exact details of the specific guest operating systems that may be supported and the terms of support have not yet been released. At the time of this writing, Microsoft Virtual Server 2005 R2 only supports Microsoft guest operating systems. • **Training and Certification (VMware)**: Currently, VMware offers courses with extensive hands-on labs, case studies, and course materials designed to be used as on the job reference tools. Courses such as GSX Server System Management and Virtual Infrastructure with ESX Server and Virtual Center provide a solid classroom environment, which is required to become a VMware Certified Professional. Currently, Microsoft does not offer a certified curriculum or certification track for Virtual Server. • **Management and User Interface (VMware GSX Server)**: All three platforms offer a Web-based management interface to manage a single virtualization host at a time. In addition to the Web application, VMware GSX Server also offers the VMware Virtual Machine Console, which is a client-based application. And for an additional sum of money, both companies off er other

	products that can help manage a larger virtualization environment. In 2003, VMware launched a brand new product to help manage and provision an entire server virtualization farm, VMware Virtual Center. It manages a server farm made up of both GSX Server and ESX Server hosts and virtual machines. Not to be outdone, Microsoft introduced a management pack add-on to its Microsoft Operations Manager (MOM) solution to help manage and monitor Virtual Server environments. • **Ease of Creating Virtual Machines (VMware GSX Server)**: When using virtualization, one of the key components and often-used features is going to be creating virtual machines. And while all platform management interfaces have an option to do it, some are much better and more user friendly than others. Microsoft Virtual Server's new virtual machine configuration process can be a little confusing at first because there is no easy work flow to follow. But after the first two or three virtual machines are created, its lack of hand holding through the process becomes a distant memory and quickly becomes second nature. VMware ESX Server does a somewhat better job of stepping through the process; however, it is VMware GSX Server's new virtual machine Wizard that leads the pack. It can even make creating the first virtual machine seems simple. To quickly create a virtual machine, it offers a typical path that takes care of most of the environmental settings with default responses. To gain more control over the creation process, a custom path is offered where more options and settings are made available. • **Hardware Support (Microsoft Virtual Server and VMware GSX Server)**: Virtual Server and GSX Server provide the best hardware support of the three platforms. Both products can install on top of almost any x86 server hardware. ESX Server on the other hand has a huge downfall when it comes to the physical hardware that it supports. To deliver such high performance, the ESX Server architecture accesses devices directly through device drivers on the ESX Server machine rather than through a host operating system such as with GSX Server or Virtual Server. The price paid for the improved performance is the limiting factor of supporting only a specific set of devices and certain qualified server families/models. Specific hardware must be present to install and use the product.
Q.48	**Compare Microsoft Virtual Server, VMware GSX and VMware GSX server virtualization technologies based on the following areas: Technical support, Performance, Price**
A.	• **Technical Support (Microsoft Virtual Server)**: With the rapid increase in use of virtualization products throughout the industry, the question of technical support has come up on more than one occasion. With the introduction of virtualization, there are now multiple levels of technical support required. Support is still needed for the hardware, host operating system, and now support is needed for the virtualization platform software, and the virtualized guest operating system. The host operating system support for VMware GSX Server or Microsoft Virtual Server depends on the platform. Microsoft Virtual Server only installs on top of a Microsoft operating system. Therefore, the support is no different than the support received from Microsoft on a normal server installation of the operating system. GSX Server offers a choice, it can either be installed on top of a Microsoft operating system (again, providing the same level of support as on a normal server installation) or a Linux operating system. The Linux operating system would be covered under the same level of support agreement as if there was no virtualization software installed and it was running some other type of application. ESX Server, on the other hand, provides its own highly efficient hypervisor/kernel as the host operating system. It then makes use of a service console originally based on a modified Red Hat Linux operating system. The ESX Server operating system is supported by VMware technical support and is covered under the

	service agreement purchased with the product. • **Performance (VMware ESX Server)**: Without going into benchmark tests, the overall performance of the three products can be summed up succinctly: VMware ESX Server overall has the best performance. This is due to a number of reasons. Perhaps the biggest boost in performance is due to the fact that ESX Server runs within its own kernel, which means there is very little overhead from the host operating system allowing its performance to be near that of a physical server. Because of ESX Server's CPU affinity and granular resource control over CPU, memory, network, and disk I/O, the other platforms are unable to surpass the scalability that ESX Server offers. By being able to over allocate its memory, ESX Server is able to run a greater number of virtual machines at the same time. And by allowing virtual machine I/O to pass directly through to the host devices, performance is greatly exceeded over hosted virtualization that inserts the overhead of an intervening Windows or Linux host operating system. Unlike the other two products, ESX Server also supports virtual Symmetric Multi-Processing (SMP) on the guest operating system. On the negative side, GSX Server and Virtual Server are both in contention for processor and memory resources with the host operating system. This further separates them from ESX Server as far as performance is concerned. Although Virtual Server does offer better virtual machine control through CPU resource allocation, the control is nowhere near as sophisticated as what ESX Server offers. • **Price (Microsoft Virtual Server)**: All three virtualization platforms have their pros and cons. But let's face it, when it comes right down to it, sometimes the biggest factor in a company's decision making process is "how much is the solution going to cost." Each product can be purchased from a number of different resellers, vendors, or in some cases even directly from the manufacturer. To try and compare the cost of these products as fairly as possible, the prices were found from a single reseller. The most inexpensive platform was Microsoft Virtual Server. The reseller offers a single server 4CPU license copy for $476.28. Coming in next was VMware GSX Server. The same reseller offers a single server 2CPU license copy for $1,299.28. And finally, the most expensive (although highest performing) product was VMware ESX Server. The reseller offers a single server 2CPU license copy for $4,109.46. This is by no means the only packages offered from these manufacturers nor is this reseller the only one in the market. These prices are offered as a quick and easy way to compare what a single reseller is charging. Not to be confused, Microsoft Virtual Server also offers a 32CPU version, while VMware GSX Server offers an unlimited CPU version, and ESX Server offers numerous CPU count packages. Although pricing can fluctuate and vary from one reseller to the next, the general idea is that Microsoft Virtual Server is the most inexpensive, followed by VMware GSX Server, and then VMware ESX Server.
Q.49	**Discuss the selection of virtualization product (Microsoft Virtual Server, VMware GSX and VMware GSX Server) for: Software development and test lab automation, Legacy application migration and rehousing, department and branch server consolidation.**
A.	• **Software Development and Test Lab Automation**: Server virtualization is ideal for situations that require frequent server reconfiguration as is required with software development and testing. By providing automated configuration, integration, and management of virtual machines, a software development and test group can become more productive and offer better test coverage throughout the software life cycle. Both, Microsoft and VMware, provide an adequate feature set to accomplish this use case. By making use of Virtual Server's Undo disk, GSX Server's Snapshot, or ESX Server's REDO disk modes, a virtual machine can quickly be rolled back to a

	pristine state prior to the problem event. Virtual Server's differencing disks can easily help to create a variety of configurations from a single base disk, which can be extremely useful for things such as a large testing matrix or software patch testing. • **Legacy Application Migration and Re-hosting**: Many companies are still running applications that require older operating systems such as Microsoft Windows NT Server 4.0. Upgrading the application can either be too costly, or is not even offered as an option. There are two other problems associated with these legacy applications. First, support for Windows NT Server 4.0 as an operating system is coming to an end. Second, the servers running these applications are also aging and probably nearing end of life with the hardware vendor. Migrating these applications to newer equipment may prove difficult, as drivers for the new equipment may not be readily available from the vendor for an out-dated operating system. Even if drivers were available, it probably would not be cost effective to run an older operating system and application on a powerful new server. The server would be way underutilized for the amount of money invested. Migrating these applications to virtual machines enables efficient use of modern hardware resources. All three virtualization platforms can successfully support a legacy application and operating system inside of a virtual machine. However, if the legacy operating system is Microsoft Windows NT Server 4.0, Microsoft Virtual Server may be the better platform choice simply because of potentially better technical support. If the legacy application is running on a Linux operating system, either VMware platform will probably be better suited to re-host the application. • **Department and Branch Server Consolidation**: Many departments and branch offices have servers that are being underutilized. Another problem is the number of servers being managed because of the applications that are installed. They either need different operating systems, different versions of the same operating system, or they cannot be installed on a server with any other application (needs application isolation). These are all excellent examples of servers and applications that can be consolidated into virtual machines that run on far fewer physical servers. While all three platforms can perform this function, it may come down to a cost decision. Depending on the size of the department or branch and the number of applications and servers needing to be consolidated, VMware ESX Server may get priced out of the running. If the size is relatively small, and cost is a major factor in the decision making process, the scalability of ESX Server may not be needed, and therefore Virtual Server or GSX Server would be the better fi t.
Q.50	**Discuss the selection of virtualization product (Microsoft Virtual Server, VMware GSX and VMware GSX Server) for:** **Product data center consolidation, cost effective disaster recovery and high availability, guaranteed service level agreements.**
A.	• **Production Data Center Consolidation**: As in the previous topic, consolidating servers in a production data center can be a very important use case. The cost of doing business in a data center can become highly overwhelming when paying for ping, power, and pipe in a hosted facility. If the data center is internally supported, there are even more costs associated with the number of physical servers and rack space such as cooling, real estate costs, and general management. In both cases, fewer physical servers generally translate into less overhead and expenses. Server consolidation in a production data center differs slightly from a department or branch consolidation. The need for scalability, stability, and performance will almost always outweigh the cost in any well run facility. In this scenario, nothing surpasses VMware ESX server.

	• **Cost-effective Disaster Recovery (DR)**: Disaster recovery is an integral element of a business continuity strategy. Implementing a virtualization solution can help realize capital cost savings, improvements in time to recovery, as well as reliability of recovery. Rather than using a costly 1-to-1 mapping of production servers to DR servers, virtualization makes it possible for many production servers to be recovered on as few as one DR server. Virtualization does more than just reduce the number of servers needed in a disaster recovery site; it also provides a hardware agnostic platform that offers the flexibility of being able to restore systems without identical or near-identical hardware. Any of the three platforms can easily provide a disaster recovery solution for a company. There are limiting factors in the decision, such as the size of the implementation and the types of applications that are being recovered, as well as the amount of time the DR facility will be online once activated. If the disaster recovery period has a relatively short-lived window, either Virtual Server or GSX Server should be able to handle the task quite well. Unfortunately, predicting the length of time that a DR facility will have to remain online is about as easy as predicting when a disaster will strike. Therefore, it is more important to look at the amount of servers that will need to be serviced and the amount of resources that are needed by the applications. VMware ESX Server provides a more robust platform, and by making use of its resource allocation abilities along with its virtual SMP, it is able to provide more peace of mind that the failover will be able to provide the necessary processing power needed to meet the expectations of most applications for a longer period of time. • **Cost-effective High Availability (HA)**: A cost-effective high-availability solution has similar results to that of a cost effective disaster recovery solution. Maximizing server uptime has become increasingly more important and critical to the productivity and satisfaction of a company's customers or employees. What used to be sold as a premium, HA is more and more becoming an entitlement, something that end users have just come to expect. Using virtualization, network load balancing, standby, replication, and clustering of machines has become cost-effective and a lot easier to implement. All three platforms have the ability to provide each of these functions; however, combining the virtualization platform with third-party tools or add-on packages creates a more effective HA solution. One of the biggest problems with Virtual Server and GSX Server for Windows is uptime. Since both platforms have a Windows Server operating system underneath, the host server and all virtual machines will need to be taken down on a fairly regular basis for update and patch management. • **Guaranteed Service Level Agreements**: When consolidating virtual machines onto a single hardware platform, the virtual machines will probably require different server workloads based on the different applications that they are running. One application may be more CPU intensive while another application may require more memory. In this scenario, ESX Server provides the best guaranteed service level agreements by making use of its resource controls for virtual machine CPU, memory, disk I/O, and network I/O. ESX Server uses a proportional-share based policy to guarantee service levels. Resources can be divided into a number of shares where each virtual machine is then given a proportionate amount of the shares for each of these resources. Neither of the other two platforms has this level of granularity of resource control, and therefore cannot match the guaranteed service level agreement of ESX Server.
Q.51	**Explain the different hardware that can be seen by the guest operating systems hosted on Microsoft Virtual Server, VMware GSX and VMware GSX server.**
A.	**Refer Reference Book's pdf file Page Number-115-116**