CMR
INSTITUTE OF
TECHNOLOGY

USN | | | | | | | | | | |

Internal Assessment Test 2– March. 2025

Sub:	Cloud Computing							Sub Code:	22MCA332
Date:	12/3/2025	Duration:	90 min's	Max Marks:	50	Sem:	III	Branch:	MCA

Note : Answer FIVE FULL Questions, choosing ONE full question from each Module

		MARKS	OBE	
			CO	RBT
	PART I			
1	Why virtualization techniques have gained a renewed interest recently?	[10]	CO3	L2
	OR			
2	Explain in detail about the characteristics of virtualized solution.	[10]	CO3	L2
	PART II			
3	Explain in detail the Popek and Goldberg theorem in virtualization	[10]	CO3	L2
	OR			
4	List and discuss different types of hardware virtualization techniques	[10]	CO3	L2
	PART III			
5	What are the problems of virtualization (disadvantages) and how to solve them?	[10]	CO3	L3
	OR			
6	Explain Microsoft Hyper-V architecture	[10]	CO3	L3
	PART IV			
7	With a neat diagram explain VMware Virtualization Solutions	[10]	CO3	L3
	OR			
8	OR Discuss VMware full virtualization reference model	[10]	CO3	L2
	PARTV			
9	Explain Application Level virtualization	[10]	CO3	L2
	OR			
10	What is Hypervisor? With a neat diagram explain the types of hypervisor.	[10]	CO3	L2

Q1) Why virtualization techniques have gained a renewed interest recently?

Virtualization technologies have gained renewed interested recently due to the confluence of severalphenomena:

(a) Increased performance and computing capacity.

The high-end side of the PC market, where supercomputers can provide immense compute power that can accommodate the execution of hundreds or thousands of virtual machines.

(b) Underutilized hardware and software resources.

Hardware and software underutilization is occurring due to (1) increased performance and computing capacity, and (2) the effect of limited or sporadic use of resources.

Computers today are so powerful that in most cases only a fraction of their capacity is used by an application or the system. Using these resources for other purposes after hours could improve the efficiency of the IT infrastructure.

(c) Lack of space.

Companies such as Google and Microsoft expand their infrastructures by building data centers as large as football fields that are able to host thousands of nodes. Although this is viable for IT giants, in most cases enterprises cannot afford to build another data center to accommodate additional resource capacity. This condition, along with hardware underutilization, has led to the diffusion of a technique called server consolidation

(d) Greening initiatives.

Maintaining a data center operation not only involves keeping servers on, but a great deal of energy is also consumed in keeping them cool. Infrastructures for cooling have a significant impact on the carbon footprint of a data center. Hence, reducing the number of servers through server consolidation will definitely reduce the impact of cooling and power consumption of a data center. Virtualization technologies can provide an efficient way of consolidating servers.

(e) Rise of administrative costs.

The increased demand for additional capacity, which translates into more servers in a data center, is also responsible for a significant increment in administrative costs. Computers—in particular, servers—do not operate all on their own, but they require care and feeding from system administrators.

These are labor-intensive operations, and the higher the number of servers that have to be managed, the higher the administrative costs. Virtualization can help reduce the number of required servers for a given workload, thus reducing the cost of the administrative personnel.

Q2) Explain in detail about the characteristics of virtualized solution.

The characteristics of virtualized solutions are:

1 Increased security

2 Managed executions

3 Portability

1. Increased security

The virtual machine represents an emulated environment in which the guest is executed. All the operations of the guest are generally performed against the virtual machine, which then translates and applies them to the host. This level of indirection allows the virtual machine manager to control and filter the activity of the guest, thus preventing some harmful operations from being performed. For example, applets downloaded from the Internet run in a sandboxed 3 version of the Java Virtual Machine (JVM), which provides them with limited access to the hosting operating system resources. Both the JVM and the .NET runtime provide extensive security policies for customizing the execution environment of applications.

2 Managed executions

Virtualization of the execution environment not only allows increased security, but a wider range of features also can be implemented. In particular, sharing, aggregation, emulation, and isolation

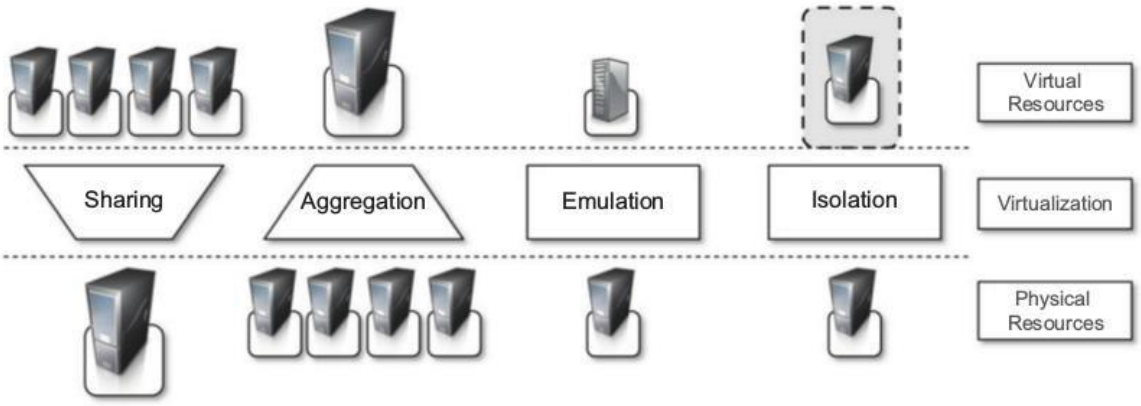

FIGURE 3.2

Functions enabled by managed execution.

are the most relevant features (see Figure 3.2).

(a) Sharing - Virtualization allows the creation of a separate computing environments within the same host. In this way it is possible to fully exploit the capabilities of a powerful guest, which would otherwise be underutilized.

(b)Aggregation - Not only is it possible to share physical resource among several guests, but virtualization also allows aggregation, which is the opposite process. A group of separate hosts can be tied together and represented to guests as a single virtual host.

(c)Emulation - Guest programs are executed within an environment that is controlled by the virtualization layer, which ultimately is a program. This allows for controlling and tuning the environment that is exposed to guests. For instance, a completely different environment with respect to the host can be emulated, thus allowing the execution of guest programs requiring specific characteristics that are not present in the physical host.

(d) Isolation - Virtualization allows providing guests—whether they are operating systems, applications, or other entities—with a completely separate environment, in which they are executed. The guest program performs its activity by interacting with an abstraction layer, which provides access to the underlying resources.

3 Portability

The concept of portability applies in different ways according to the specific type of virtualization considered. In the case of a hardware virtualization solution, the guest is packaged into a virtual image that, in most cases, can be safely moved and executed on top of different virtual machines.

In the case of programming-level virtualization, as implemented by the JVM or the .NET runtime, the binary code representing application components (jars or assemblies) can be run without any recompilation on any implementation of the corresponding virtual machine. This makes the application development cycle more flexible and application deployment very straightforward: One version of the application, in most cases, is able to run on different platforms with no changes.

Q3) Explain in detail the Popek and Goldberg theorem in virtualization

Popek and Goldberg provided a classification of the instruction set and proposed three theorems that define the properties that hardware instructions need to satisfy in order to efficiently supportvirtualization.

THEOREM 1

For any conventional third-generation computer, a VMM may be constructed if the set of sensitive instructions for that computer is a subset of the set of privileged

This theorem establishes that all the instructions that change the configuration of the system resources should generate a trap in user mode and be executed under the control of the virtual machine manager.

THEOREM 2

A conventional third-generation computer is recursively virtualizable if:

- It is virtualizable and
- A VMM without any timing dependencies can be constructed for it.

Recursive virtualization is the ability to run a virtual machine manager on top of another virtual machine manager. This allows nesting hypervisors as long as the capacity of the underlying resources can accommodate that. Virtualizable hardware is a prerequisite to recursive virtualization.

THEOREM 3

A hybrid VMM may be constructed for any conventional third-generation machine in which the set of user-sensitive instructions is a subset of the set of privileged instructions.

There is another term, hybrid virtual machine (HVM), which is less efficient than the virtual machine system. In the case of an HVM, more instructions are interpreted rather than being executed directly. All instructions in virtual supervisor mode are interpreted. Whenever there is an attempt to execute a behavior-sensitive or control-sensitive instruction, HVM controls the execution directly or gains the control via a trap. Here all sensitive instructions are caught by HVM that are simulated.

Q4) List and discuss different types of hardware virtualization techniques

Hardware virtualization techniques (System Level techniques)

- Hardware-assisted virtualization
- Full virtualization
- Paravirtualization
- Partial virtualization

- **Hardware-assisted virtualization (Hardware – Processors)**

  This term refers to a scenario in which the hardware provides architectural support for building a virtual machine manager able to run a guest operating system in complete isolation.

- **Full virtualization (Running Operating System – no modified OS)**

  Full virtualization refers to the ability to run a program, most likely an operating system, directly on top of a virtual machine and without any modification, as though it were run on the raw hardware.

- **Paravirtualization (Thin virtual machine - modified OS)**

  Paravirtualization techniques expose a software interface to the virtual machine that is slightly modified from the host and, therefore, guests need to be modified. The aim of paravirtualization is to provide the capability to demand the execution of performance-critical operations directly on the host

- **Partial virtualization**

  Partial virtualization provides a partial emulation of the underlying hard- ware, thus not allowing the complete execution of the guest operating system in complete isolation. Partial virtualization allows many applications to run transparently, but not all the features of the operating system can be supported.

Q5) What are the problems of virtualization (disadvantages) and how to solve them?

(a) Performance degradation

Performance is one of the major concerns in using virtualization technology. Since virtualization interposes an abstraction layer between the guest and the host, the guest can experience increased latencies (delays).

For instance, in the case of hardware virtualization, where the intermediate emulates a bare machine on top of which an entire system can be installed, the causes of performance degradation can be traced back to the overhead introduced by the following activities:

- Maintaining the status of virtual processors
- Support of privileged instructions (trap and simulate privileged instructions)
- Support of paging within VM
- Console functions

(b) Inefficiency and degraded user experience

Virtualization can sometime lead to an inefficient use of the host. Some of the specific features of the host cannot be exposed by the abstraction layer and then become inaccessible. In the case of hardware virtualization, this could happen for device drivers: The virtual machine can sometime simply provide a default graphic card that maps only a subset of the features available in the host. In the case of programming-level virtual machines, some of the features of the underlying operating systems may become inaccessible unless specific librariesare used.

(c) Security holes and new threats

Virtualization opens the door to a new and unexpected form of phishing. The capability of emulating a host in a completely transparent manner led the way to malicious programs that are designed to extract sensitive information from the guest. The same considerations can be made for programming-level virtual machines: Modified versions of the runtime environment can access sensitive information or monitor the memory locations utilized by guest applications while these are executed.

Q6) Explain Microsoft Hyper-V architecture

Hyper – V is an infrastructure virtualization solution developed by Microsoft for server virtualization. As the name recalls, it uses a hypervisor-based approach for hardware virtualization, which leverages several techniques to support a variety 2008 R2 that installs the hypervisor as a role within the server.

1. Architecture

Hyper – V supports multiple and concurrent execution of the guest operating system by means of partitions. A partition is a completely isolated environment in which an operating system is installed and run. Fig. 3.17 provides an overview of the architecture of Hyper – V. Hyper – V takes control of the hardware, and the host operating system becomes a virtual machine instance with special privileges, called parent partition. The parent partition (also called root partition) is the only one that has direct access to the hardware, it runs the virtualized stack, host all the drivers required to configure guest operating systems and creates child partitions through the hypervisor. Child partitions

are used to host guest operating systems and do not have access to the underlying hardware, but their interaction with it is controlled by either the parent partition or the hypervisor itself.

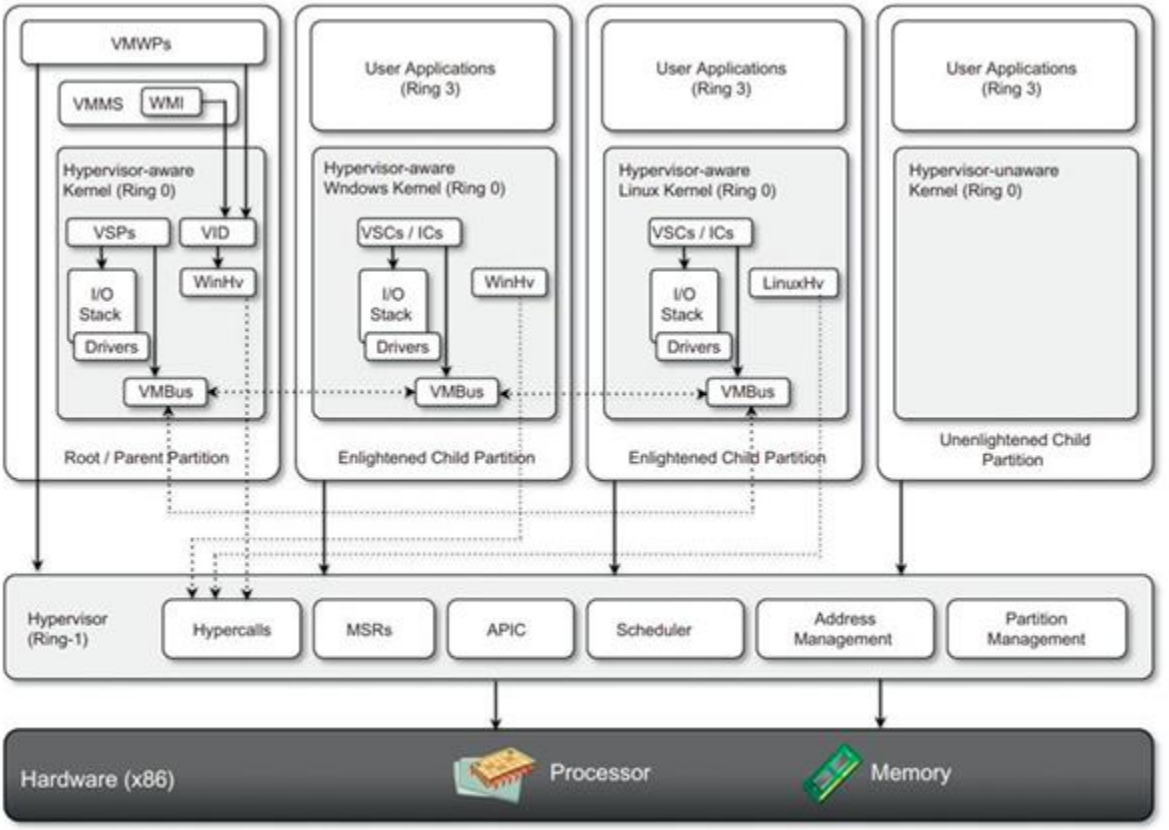

FIGURE 3.17
Microsoft Hyper-V architecture.

(a) Hypervisor: the hypervisor Is the component that directly manages the underlying

(b) hardware (processors and memory). It is logically defined by the following components:

- Hypercalls interface
- Memory service routines (MSRs)
- Advanced programming interrupt controller (APIC)
- Schedular
- Address manager
- Partition manager

Hypercalls interface

This is the entry point for all the partitions for the execution of sensible instructions. This interface is used by drivers in the partitioned operating system to contact the hypervisor by using the standard windows calling convention. The parent partition also uses interface to create children partitions.

Memory service routines (MSRs)

These are the set of functionalities that control the memory, and its access from partitions. By leveraging hardware-assisted virtualization, the hypervisor uses the input output memory management unit (I/O MMU or IOMMU) to fast track the access to devices from partitions, by translating virtual memory addresses.

Advanced programming interrupt controller (APIC)

It represents the synthetic interrupt controller, which manages the virtual processor signals coming from the underlying hardware when some event occurs. The hypervisor is responsible of dispatching, when appropriate, the physical interrupts to the synthetic interrupt controller.

Schedular

It schedules the virtual processors to run on available physical processors. The scheduling is controlled by policies that are set by the parent partition.

Address manager

It is used to manage the virtual network addresses that are allocated to each guest operating system.

Partition manager

It creates the partition, finalization, destruction, enumeration, and configurations. Its services are available through the hypercalls interface API.

c) Enlightened I/O and Synthetic Devices: It provides inter-partition communication channel rather than traversing the hardware emulation stack provided by the hypervisor. VMBus - an inter-partition communication channel that is used to exchange data between partitions for guest operating system. VSPs (Virtual Service Providers)– these are kernel level drivers that are deployed in the parent partition and provides access to the corresponding hardware devices. VSCs – these VSPs

interact with VSCs which represent the virtual device drivers (synthetic drivers) seen by the guest operating systems in the children partitions.

(d) Parent Partition: The parent partition is also the one that manages the creation, execution and destruction of children partitions. VID (virtualization Infrastructure Driver) – which controls the access to the hypervisor and also allows the management of virtual processor and memory. VMWPs (Virtual Machine Worker Process) – which manages the children partition by interacting with the hypervisor through the VID.

(e) Children Partition: Children partitions are used to execute guest operating systems. These are isolated environments, which allow a secure and controlled execution of guests.

2. Cloud Computing and Infrastructure Management

Hyper-V constitutes the basic building block of Microsoft virtualization infrastructure. It contributes to create a ful-featured platform for server virtualization. The basic features offered by Hyper-V with management capabilities includes:

- Management portal for the creation and management of virtual instances
- Virtual to virtual (V2V) and physical to virtual (P2V) conversions
- Delegated administration
- Library functionality and deep PowerShell integration
- Intelligent placement of virtual machines in the managed environment and
- Host capacity management

Q7) With a neat diagram explain VMware Virtualization Solutions

i. End-user (Desktop) Virtualization

VMware supports virtualization of operating system environments and single applications on end-user's computers. Specific VMware software – VMware Workstation, for windows operating systems and VMware Fusion, for Mac OS X environments – is installed in the host operating system to create virtual machines and manage their execution. Besides the creation of an isolated computing

environment, the two products allow a guest operating system to leverage the resources of the host machine (USB devices, folder sharing and integration with the GUI of the host operating system. Fig. 3.13 provides an overview of the architecture of these systems.

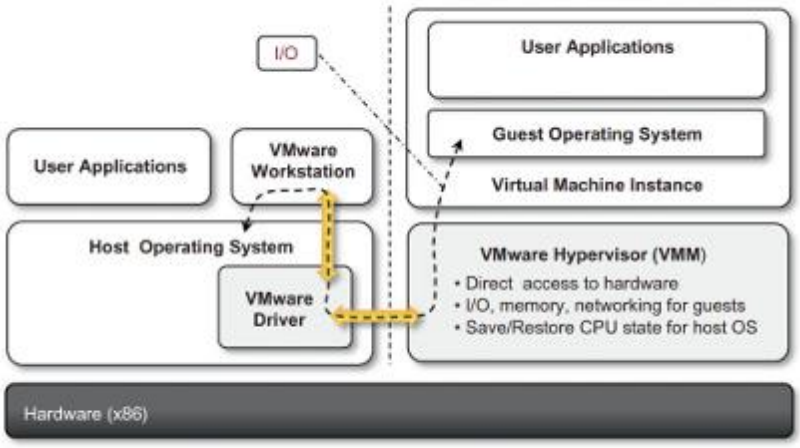

FIGURE 3.13

VMware workstation architecture.

The virtualization environments is created by an application installed in the guest operating system, which provides the guest operating system with full hardware virtualization of the underlying hardware. This is done by installing a specific driver in the operating system that provides two main services.

- It deploys a virtual machine manager that can run in privileged mode.
- It provides hooks for the VMware application to process specific I/O requests eventuallyby relaying such requests to the host operating system via system calls.

ii. Server Virtualization

Initial support for server virtualization was provided by VMware GSX server, which replicates the approach used for end-user computers and introduces remote management and scripting capabilities. The architecture of VMware GSX server is depicted in Fig. 3.14.

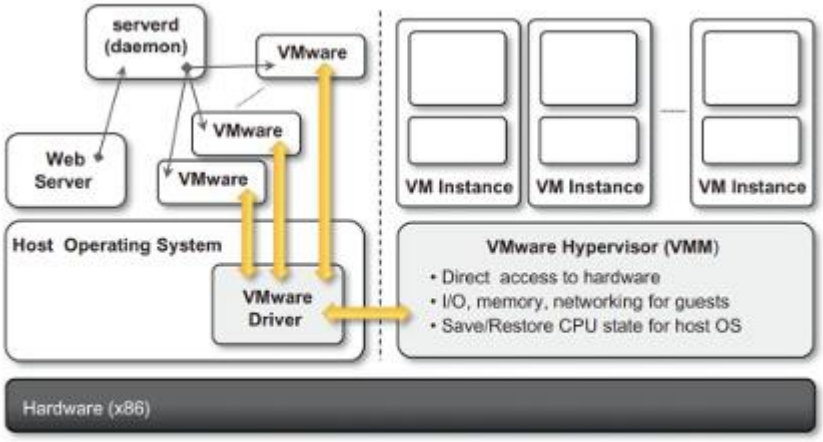

FIGURE 3.14

VMware GSX server architecture.

The architecture is designed to serve the virtualization of web servers. A daemon proce called served, controls, and manages VMware application processes. These applications are then connected to the virtual machines instances by means of the VMware driver installed on the host operating system. Virtual machine instances are managed by the VMM as described previously. User request for virtual machine management and provisioning are routed from the web server through the VMM by means of served.

VMware ESX Server and its enhanced version, VMWare ESXi Server, are examples of the

hypervisor-based approach. The base of the infrastructure is the VMkernel, which is a thin Portable Operating System Interface (POSIX) compliant operating system that provides the minimal functionality for processes and thread management, file system, I/O stacks, and resource scheduling. The kernel is accessible through specific APIs called User world API. These APIs are utilized by all the agents that provide supporting activities for the management of virtual machines. Remote management of an ESXi server is provided by the CIM Broker, a system agent that acts as a gateway to the VMkernel for clients by using the Common Information Model (CIM)8 protocol. The ESXi installation can also be managed locally by a Direct Client User Interface (DCUI), which provides a BIOS-like interface for the management of local users.

iii. Infrastructure virtualization and cloud computing solutions

VMware provides a set of products covering the entire stack of cloud computing from infrastructure management to software as a service solution hosted in the cloud. Fig. 3.16 gives an overview of the different solutions offered and how they relate to each other.

A collection of virtualized data centers are turned into a infrastructure-as-a-service cloud by VMware **vCloud**, which allows service providers to make available to end users a virtual computing environments, on demand, on a pay-per-use basis.

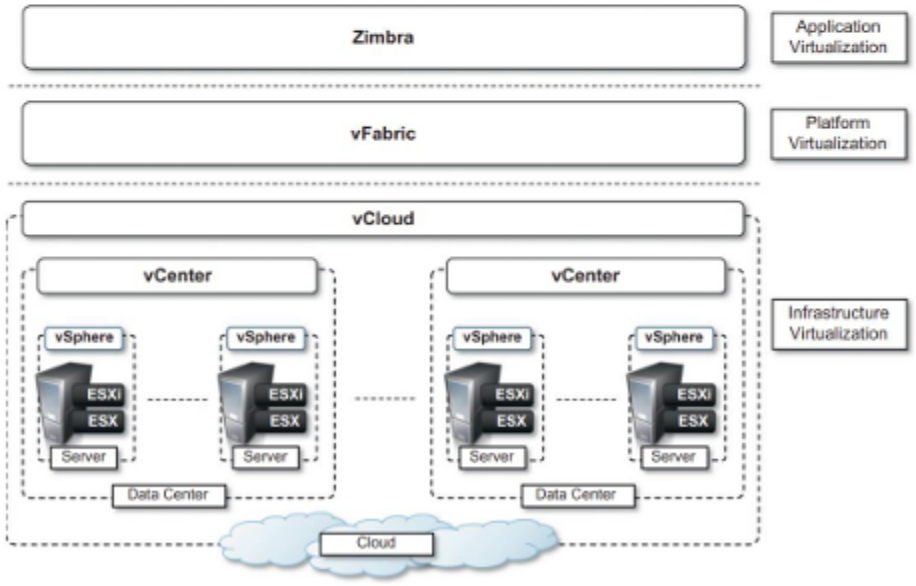

FIGURE 3.16
VMware Cloud Solution stack.

Q8) Discuss VMware full virtualization reference model

VMware: Full Virtualization

In full virtualization primary hardware is replicated and made available to the guest operating system, which executes unaware of such abstraction and no requirements to modify. Technology of VMware is based on the key concept of Full Virtualization. Either in desktop

environment, with the help of type-II hypervisor, or in server environment, through type-I hypervisor, VMware implements full virtualization. In both the cases, full virtualization is possible through the direct execution for non-sensitive instructions and binary translation for sensitive instructions or hardware traps, thus enabling the virtualization of architecture like x86.

(a) Full Virtualization and Binary Translation

VMware is widely used as it tends to virtualize x86 architectures, which executes unmodified on-top of their hypervisors. With the introduction of hardware-assisted virtualization, full virtualization is possible to achieve by support of hardware. But earlier, x86 guest operating systems unmodified in a virtualized environment could be executed only with the use of dynamic binary translation.

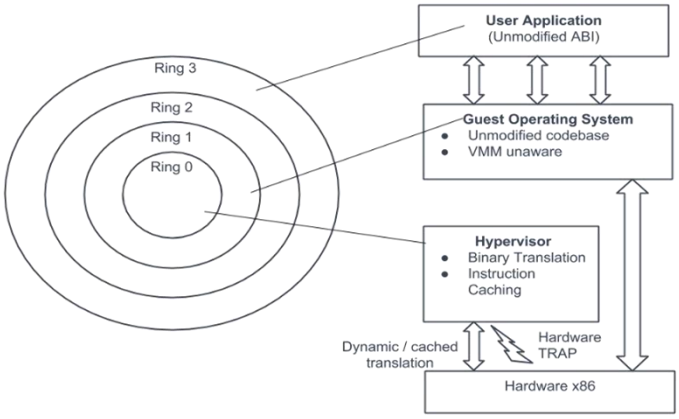

Figure – Full Virtualization Reference Model

The major benefit of this approach is that guests can run unmodified in a virtualized environment, which is an important feature for operating system whose source code does not exist. Binary translation is portable for full virtualization. As well as translation of instructions at runtime presents an additional overhead that is not existed in other methods like

paravirtualization or hardware-assisted virtualization. Contradict, binary translation is only implemented to a subset of the instruction set, while the others are managed through direct execution on the primary hardware. This depletes somehow the impact on performance of binary translation.

Advantages of Binary Translation –

1. This kind of virtualization delivers the best isolation and security for Virtual Machine.
2. Truly isolated numerous guest OS can execute concurrently on the same hardware.
3. It is only implementation that needs no hardware assist or operating system assist tovirtualize sensitive instruction as well as privileged instruction.

Disadvantages of Binary Translation –

1. It is time consuming at run-time.
2. It acquires a large performance overhead.
3. It employs a code cache to stock the translated most used instructions to enhance theperformance, but it increases memory utilization along with the hardware cost.
4. The performance of full virtualization on the x86 architecture is 80 to 95 percent that ofthe Host machines.

Q9) Explain Application Level virtualization

- Application-level virtualization is a technique allowing applications to be **run in runtime environments that do not natively support all the features required by such applications**. In this scenario, **applications are not installed in the expected runtime environment but are run as though they were.**
- Emulation can also be used to execute program binaries compiled for different hardware architectures. In this case, one of the following strategies can be implemented:

a. Interpretation. In this technique every source instruction is interpreted by an emulator for executing native ISA instructions, leading to poor performance. Interpretation has a minimal start-up cost but a huge overhead, since each instruction is emulated.

b. Binary translation. In this technique every source instruction is converted to native instructions with equivalent functions. After a block of instructions is translated, it is cached and reused. Binary translation has a large initial overhead cost, but over time it is subject to better performance, since previously translated instruction blocks are directly executed.

- Application virtualization is a good solution in the case of missing libraries in the host operating system; in this case a replacement library can be linked with the application, or library calls can be remapped to existing functions available in the host system. Another advantage is that in this case the virtual machine manager is much lighter since it provides a partial emulation of the runtime environment compared to hardware virtualization.

- One of the most popular solutions implementing application virtualization is Wine, which is a software application allowing Unix-like operating systems to execute programs written for the Microsoft Windows platform. Wine features a software application acting as a container for the guest application and a set of libraries, called Winelib, that developers can use to compile applications to be ported on Unix systems.

- Wine takes its inspiration from a similar product from Sun, Windows Application Binary Interface (WABI), which implements the Win 16 API specifications on Solaris. A similar solution for the Mac OS X environment is CrossOver, which allows running Windows applications directly on the Mac OS X operating system. VMware ThinApp, another product in this area, allows capturing the setup of an installed application and packaging it into an executable image isolated from the hosting operating system

Q10) What is Hypervisor? With a neat diagram explain the types of hypervisor

Hypervisors

A fundamental element of hardware virtualization is the **hypervisor**, or virtual machine manager (**VMM**). It recreates a hardware environment in which guest operating systems are installed. There are two major types of hypervisors: Type I and Type II (see Figure 3.7).

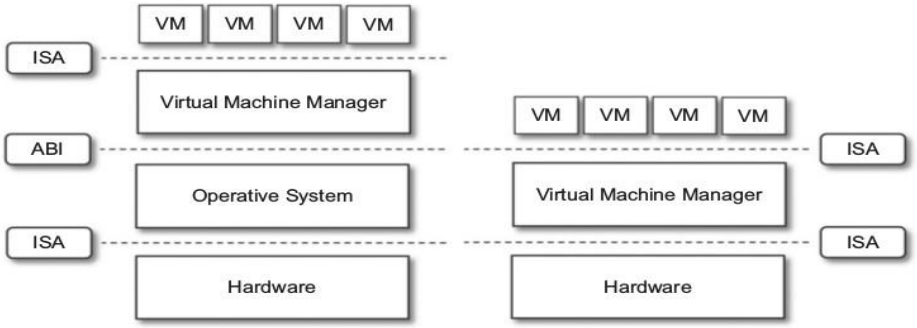

FIGURE 3.7

Hosted (left) and native (right) virtual machines. This figure provides a graphical representation of the two types of hypervisors.

- **Type I hypervisors (native virtual machine)** It run directly on top of the hardware. Therefore, they take the place of the operating systems and interact directly with the ISA interface exposed by the underlying hardware, and they emulate this interface to allow the management of guest operating systems. This type of hypervisor is also called a **native virtual machine** since it runs natively on hardware.
- **Type II hypervisors (hosted virtual machine)** It requires the support of an operating system to provide virtualization services. This means that they are programs managed by the operating system, which interact with it through the ABI and emulate the ISA of virtual hardware for guest operating systems. This type of hypervisor is also called a **hosted virtual machine** since it is hosted within an operating system.
- **Virtual machine manager (VMM)** is internally organized as described in Figure 3.8.

  Three main modules, **dispatcher**, **allocator**, and **interpreter**, coordinate their activity to emulate the underlying hardware.

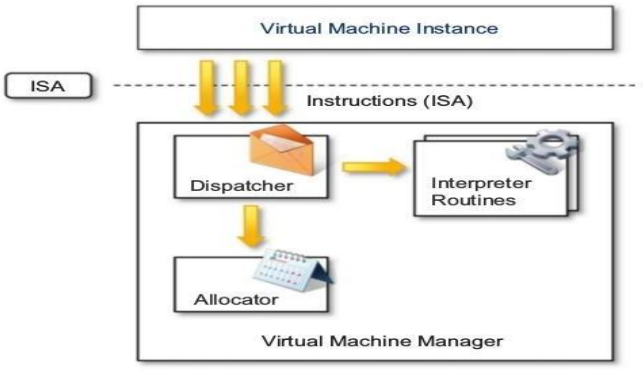

FIGURE 3.8

A hypervisor reference architecture.

- The **dispatcher** constitutes the entry point of the monitor and reroutes the

instructionsissued by the virtual machine instance to one of the two other modules.

- The **allocator** is responsible for deciding the **system resources to be provided to the VM**.
- The **interpreter module** consists of interpreter routines. These are executed whenever avirtual machine executes a privileged instruction.

The design and architecture of a virtual machine manager, together with the underlying hardware design of the host machine, determine the full realization of hardware virtualization, where a guest operating system can be transparently executed on top of a VMM as though it were run on the underlying hardware.