USN	1	P	E			M	C	A		

PESIT Bangalore South Campus
Hosur Road, 1km before Electronic City, Bengaluru -560100
Department of Master of Computer Applications

INTERNAL ASSESSMENT TEST II

Date : 24/09/19	**Max Marks:** 40
Subject & Code: Cloud Computing (17MCA542)	**Section:** V MCA
Name of Faculty: Kanthimathi S	**Time:** 11:30-1:00 PM

Note: ***Answer FIVE full questions. Select one question from each part.***

	Part I	
Q 1	What is virtualization? Explain implementation levels of virtualization. 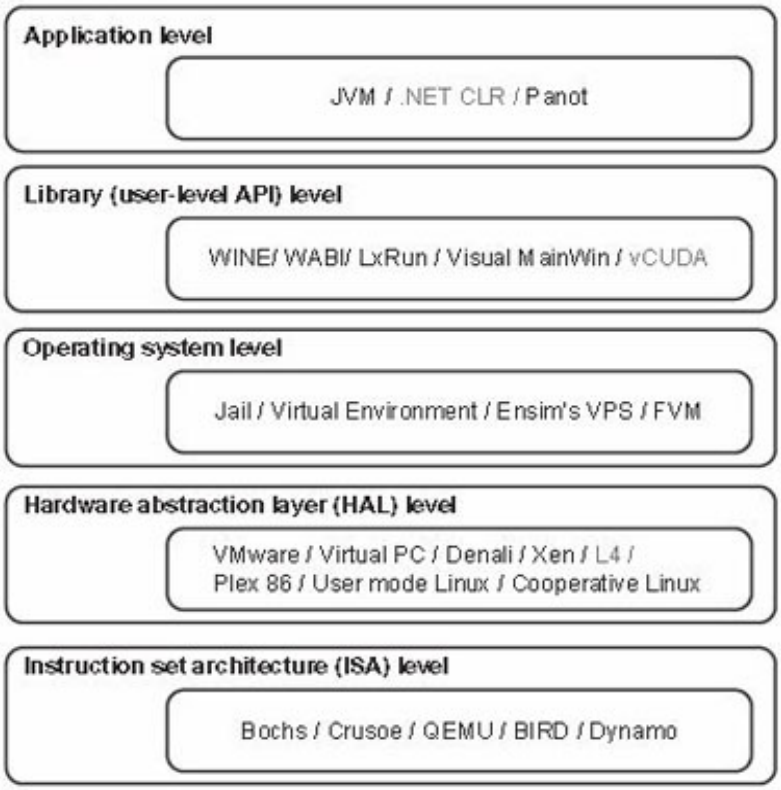 **FIGURE 3.2** Virtualization ranging from hardware to applications in five abstraction levels.	**8**
	OR	
Q 2	Discuss about Hypervisor and Xen architecture. • The hypervisor supports hardware-level virtualization (see Figure 3.1(b)) on bare metal devices like CPU, memory, disk and network interfaces. The hypervisor software sits directly between the physical hardware and its OS. • This virtualization layer is referred to as either the VMM or the hypervisor.	**8**

	• The hypervisor provides hypercalls for the guest OSes and applications. Depending on the functionality, a hypervisor can assume a micro-kernel architecture like the Microsoft Hyper-V. Or it can assume a monolithic hypervisor architecture like the VMware ESX for server virtualization. • A micro-kernel hypervisor includes only the basic and unchanging functions (such as physical memory management and processor scheduling). The device drivers and other changeable components are outside the hypervisor. • A monolithic hypervisor implements all the aforementioned functions, including those of the device drivers. Therefore, the size of the hypervisor code of a micro-kernel hypervisor is smaller than that of a monolithic hypervisor. • Essentially, a hypervisor must be able to convert physical devices into virtual resources dedicated for the deployed VM to use • Xen is an open source hypervisor program developed by Cambridge University. Xen is a micro-kernel hypervisor, which separates the policy from the mechanism. • Xen does not include any device drivers natively . It just provides a mechanism by which a guest OS can have direct access to the physical devices. • As a result, the size of the Xen hypervisor is kept rather small. • Core components • Domain 0 and Domain U • Compromising Domain 0 • Flexibility to modify.	
	Part II	
Q3	Explain about three types of clouds. • A **public cloud** is built over the Internet and can be accessed by any user who has paid for the service. – **Public clouds** are owned by service providers and are accessible through a subscription. • The callout box in top of **Figure 4.1** shows the architecture of a typical public cloud. – Many public clouds are available, including **Google App Engine (GAE)**, **Amazon Web Services** (**AWS**), **Microsoft Azure**, **IBM Blue Cloud**, and Salesforce.com's **Force.com**. – The providers of the aforementioned clouds are commercial providers that offer a *publicly accessible remote interface* for creating and managing **VM** instances . • A **private cloud** is built within the domain of an **intranet** owned by a single organization. • Therefore, **it is client owned and managed**, and its access is limited to the owning clients and their partners. – Its deployment was not meant to sell capacity over the Internet through publicly accessible interfaces. • A **private cloud** is supposed to deliver more efficient and convenient cloud services. • It may impact the cloud standardization, while retaining greater customization and organizational control. • A **hybrid cloud** is built with **both public** and **private clouds**, as shown at the lower-left corner of **Figure 4.1**. • **Private clouds** can also support a **hybrid cloud** model by supplementing *local infrastructure with computing capacity* from an external public cloud. – For example, the **Research Compute Cloud** (**RC2**) is a **private cloud**, built by IBM, that interconnects the computing and IT resources at eight IBM Research Centers scattered throughout the United States, Europe, and Asia. • A **hybrid cloud** provides access to clients, the partner network, and third parties.	**8**
	OR	
Q4	Illustrate the concept of Trust Management in virtualized data centers. A VMM changes the computer architecture. It provides a layer of software between the operating systems and system hardware to create one or more VMs on a single physical platform. A VM entirely encapsulates the state of the guest operating system running inside it. Encapsulated machine state can be copied and shared over the network and removed like a normal file, which proposes a challenge to VM security. In general, a VMM can provide secure isolation and a VM accesses hard-ware resources through the control	**8**

	of the VMM, so the VMM is the base of the security of a virtual system. Normally, one VM is taken as a management VM to have some privileges such as creating, suspending, resuming, or deleting a VM. Once a hacker successfully enters the VMM or management VM, the whole system is in danger. A subtler problem arises in protocols that rely on the "freshness" of their random number source for generating session keys. Considering a VM, rolling back to a point after a random number has been chosen, but before it has been used, resumes execution; the random number, which must be "fresh" for security purposes, is reused. With a stream cipher, two different plaintexts could be encrypted under the same key stream, which could, in turn, expose both plaintexts if the plaintexts have suffi-cient redundancy. Noncryptographic protocols that rely on freshness are also at risk. For example, the reuse of TCP initial sequence numbers can raise TCP hijacking attacks.	
	Part III	
Q 5	Explain in detail about the Live VM Migration steps and performance with diagram. Steps 0 and 1: Start migration. • This step makes preparations for the migration, including determining the migrating VM and the destination host. • Although users could manually make a VM migrate to an appointed host, in most circumstances, the migration is automatically started by strategies such as load balancing and server consolidation. Steps 2: Transfer memory. • Since the whole execution state of the VM is stored in memory, sending the VM's memory to the destination node ensures continuity of the service provided by the VM. • All of the memory data is transferred in the first round, and then the migration controller recopies the memory data which is changed in the last round. • These steps keep iterating until the dirty portion of the memory is small enough to handle the final copy. • Although precopying memory is performed iteratively, the execution of programs is not obviously interrupted. Step 3: Suspend the VM and copy the last portion of the data. • The migrating VM's execution is suspended when the last round's memory data is transferred. • Other nonmemory data such as CPU and network states should be sent as well. • During this step, the VM is stopped and its applications will no longer run. • This "service unavailable" time is called the "downtime" of migration, which should be as short as possible so that it can be negligible to users. Steps 4 and 5: Commit and activate the new host. • After all the needed data is copied, on the destination host, the VM reloads the states and recovers the execution of programs in it, and the service provided by this VM continues. • Then the network connection is redirected to the new VM and the dependency to the source host is cleared. • The whole migration process finishes by removing the original VM from the source host.	**8**

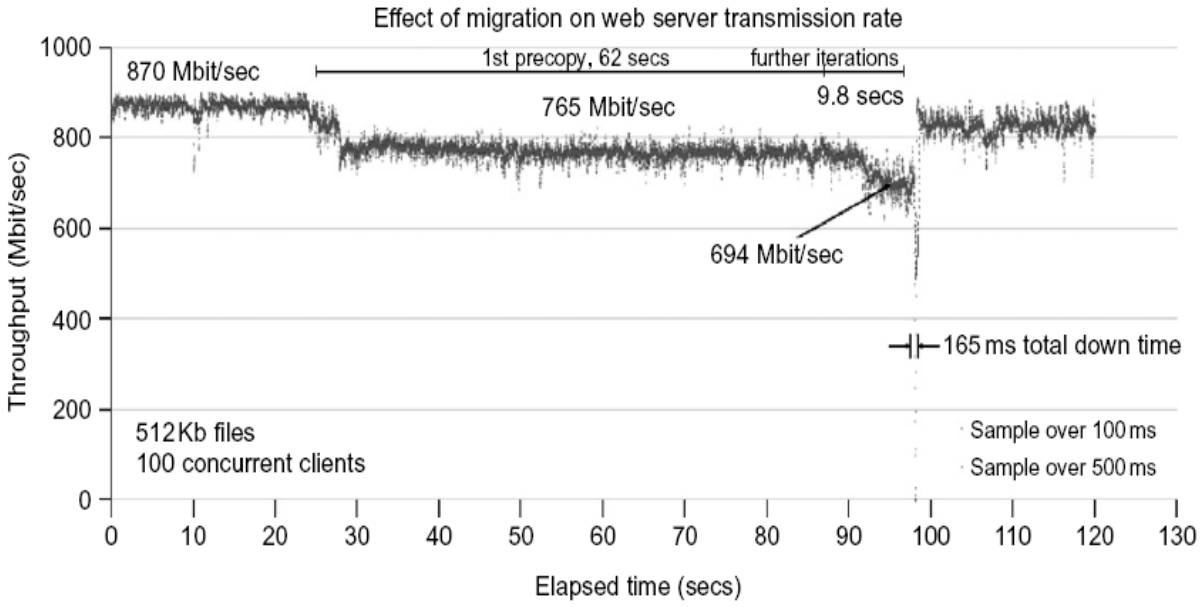

FIGURE 3.21

Effect on data transmission rate of a VM migrated from one failing web server to another.

(Courtesy of C. Clark, et al. [14])

OR

Q6 Write short notes on a) Cloud OS for virtualized Data center **4 4**

Data centers must be virtualized to serve as cloud providers. These VI managers and OSes are specially tailored for virtualizing data centers which often own a large number of servers in clusters. Nimbus, Eucalyptus and OpenNebula are all open source software available to the general public. Only vSphere 4 is a proprietary OS for cloud resource virtualization and management over data centers.These VI managers are used to create VMs and aggregate them into virtual clusters as elastic resources. Nimbus and Eucalyptus support essentially virtual networks. OpenNebula has additional features to provision dynamic resources and make advance reservations. All three public VI managers apply Xen and KVM for virtualization. vSphere 4 uses the hypervisors ESX and ESXi from VMware. Only vSphere 4 supports virtual storage in addition to virtual networking and data protection. We will study Eucalyptus and vSphere 4 in the next two examples.

b) Virtual Storage Management

In system virtualization, virtual storage includes the storage managed by VMMs and guest OSes. Generally, the data stored in this environment can be classified into two categories: VM images and application data. The VM images are special to the virtual environment, while application data includes all other data which is the same as the data in traditional OS environments.The most important aspects of system virtualization are encapsulation and isolation. Traditional operating systems and applications running on them can be encapsulated in VMs. Only one operating system runs in a virtualization while many applications run in the operating system

Part IV

Q 7 Compare and contrast **4 4**

a) Full vs Para Virtualization

- Does not need to modify guest OS
- Critical instructions are emulated by software through the use of binary translation.
- VMware Workstation applies full virtualization, which uses binary translation to automatically modify x86 software on-the-fly to replace critical instructions.

	• Binary translation slows down the performance. • VMWare ESXi and Microsoft Virtual Server are examples of full virtualization b) Physical vs Virtual Cluster • A physical cluster is a collection of servers (physical machines) interconnected by a physical network such as a LAN. • Virtual clusters are built with VMs installed at distributed servers from one or more physical clusters. • The VMs in a virtual cluster are interconnected logically by a virtual network across several physical networks • Each virtual cluster is formed with physical machines or a VM hosted by multiple physical clusters. The virtual cluster boundaries are shown as distinct boundaries	
	OR	
Q 8	Explain OS level Virtualization in detail. * It is slow to initialize a hardware-level VM because each VM creates its own image from scratch. * In cloud environment, thousands of VMs need to be initialized simultaneousl * Besides slow operation, storing the VM images also becomes an issue. * Full virtualization at the hardware level also has the disadvantages of slow performance. To reduce the performance overhead of hardware-level virtualization, even hardware modification is needed. * OS virtualization inserts a virtualization layer inside an operating system to partition a machine's physical resources. * Multiple isolated VMs within a single operating system kernel. (Virtual execution environment, Virtual private system or container) * VE has its own set of processes, file system, user accounts, network interfaces with IP addresses, routing tables, firewall rules and other personal settings.	**8**
	Part V	
Q 9	a) Define VMM with diagram of abstraction levels.	**4** **4**

Table 3.1 Relative Merits of Virtualization at Various Levels (More "X"'s Means Higher Merit, with a Maximum of 5 X's)

Level of Implementation	Higher Performance	Application Flexibility	Implementation Complexity	Application Isolation
ISA	X	XXXXX	XXX	XXX
Hardware-level virtualization	XXXXX	XXX	XXXXX	XXXX
OS-level virtualization	XXXXX	XX	XXX	XX
Runtime library support	XXX	XX	XX	XX
User application level	XX	XX	XXXXX	XXXXX

b) Write short notes on server consolidation in data centers.

Heterogeneous workloads can be roughly divided into two categories: chatty workloads and noninteractive workloads. Chatty workloads may burst at some point and return to a silent state at some other point.

A web video service is an example of this, whereby a lot of people use it at night and few people use it during the day. Noninteractive workloads do not require people's efforts to make progress after they are submitted. High-performance computing is a typical example of this.

At various stages, the requirements for resources of these workloads are dramatically different. However, to guarantee that a workload will always be able to cope with all demand levels, the workload is statically allocated enough resources so that peak demand is satisfied.

Server consolidation is an approach to improve the low utility ratio of hardware resources by reducing the number of physical servers. Among several server consolidation techniques such as centralized and physical consolidation, virtualization-based server consolidation is the most powerful.

Server virtualization enables smaller resource allocation than a physical machine

OR

Q 10 Discuss about Memory and Network migration. **8**

- It can be in a range of hundreds of megabytes to a few gigabytes in a typical system today, and it needs to be done in an efficient manner.
- The Internet Suspend-Resume (ISR) technique exploits temporal locality as memory states are likely to have considerable overlap in the suspended and the resumed instances of a VM.
- Temporal locality refers to the fact that the memory states differ only by the amount of work done since a VM was last suspended before being initiated for migration
- A migrating VM should maintain all open network connections without relying on forwarding mechanisms on the original host or on support from mobility or redirection mechanisms.
- To enable remote systems to locate and communicate with a VM, each VM must be assigned a virtual IP address known to other entities. This address can be distinct from the IP address of the host machine where the VM is currently located. Each VM can also have its own distinct virtual MAC address.
- The VMM maintains a mapping of the virtual IP and MAC addresses to their corresponding VMs.
- In general, a migrating VM includes all the protocol states and carries its IP address with it