USN | | | | | | | | | |

CELEBRATING 25 YEARS
CMRIT
CMR INSTITUTE OF TECHNOLOGY, BENGALURU.
ACCREDITED WITH A+ GRADE BY NAAC

Internal Assessment Test 1 – Sept 2019

Sub:	**Cloud Computing and its Applications**					Sub Code:	**15CS742**	Branch:	**CSE**
Date:	14/09/2019	Duration:	90 mins	Max Marks:	50	Sem / Sec:	VII/A, B, C		OBE

<u>Answer any FIVE FULL Questions</u>

		MARKS	CO	RBT
1 (a)	Compare Distributed Computing and Grid Computing.	[05]	CO2	L2
(b)	Demonstrate the usage of virtualization in Cloud Computing	[05]	CO2	L3
2 (a)	Discuss in detail about a real-life scenario under which you can rely on cloud services	[05]	CO3	L3
(b)	Describe the pros and cons of virtualization	[05]	CO1	L1
3 (a)	Discuss three open challenges in cloud.	[06]	CO1	L1
(b)	Imagine yourself as a cloud service provider. How do you design the virtual machine manager?	[04]	CO2	L3
4 (a)	Classify various types of cloud.	[05]	CO1	L2
(b)	Discuss the elements of virtualization in Xen	[05]	CO2	L2
5 (a)	Explain Goldberg and Popek criteria and theorems	[06]	CO2	L2
(b)	Describe about network, storage, desktop and OS virtualization	[04]	CO1	L2
6 (a)	Name three cloud computing services	[03]	CO1	L1
(b)	Explain cloud computing architecture with the help of a diagram	[07]	CO2	L2
7 (a)	Illustrate with example how community clouds are significant in the current computational world?	[05]	CO1	L3
(b)	Explain with examples how hybrid clouds are used in industry?	[05]	CO1	L2

CO PO Mapping

Course Outcomes		Modules covered	PO1	PO2	PO3	PO4	PO5	PO6	PO7	PO8	PO9	PO10	PO11	PO12	PSO1	PSO2	PSO3	PSO4
CO1	Explain cloud computing, virtualization and classify services of cloud computing	1,2	2	1	2	-	1	-	-	-	-	-	-	-	-	1	-	-
CO2	Illustrate architecture and programming in cloud	2,3,4	2	2	2	1	2	1	-	-	-	-	-	-	1	1	-	1
CO3	Describe the platforms for development of cloud applications	4,5	2	2	2	2	2	1	-	-	-	-	-	-	2	2	-	1
CO4	List the applications of cloud	4,5	2	2	2	2	2	1	-	-	-	-	-	-	2	2	-	1

COGNITIVE LEVEL	REVISED BLOOMS TAXONOMY KEYWORDS
L1	List, define, tell, describe, identify, show, label, collect, examine, tabulate, quote, name, who, when, where, etc.
L2	summarize, describe, interpret, contrast, predict, associate, distinguish, estimate, differentiate, discuss, extend
L3	Apply, demonstrate, calculate, complete, illustrate, show, solve, examine, modify, relate, change, classify, experiment, discover.
L4	Analyze, separate, order, explain, connect, classify, arrange, divide, compare, select, explain, infer.
L5	Assess, decide, rank, grade, test, measure, recommend, convince, select, judge, explain, discriminate, support, conclude, compare, summarize.

PROGRAM OUTCOMES (PO), PROGRAM SPECIFIC OUTCOMES (PSO)				CORRELATION LEVELS	
PO1	Engineering knowledge	PO7	Environment and sustainability	0	No Correlation
PO2	Problem analysis	PO8	Ethics	1	Slight/Low
PO3	Design/development of solutions	PO9	Individual and team work	2	Moderate/ Medium
PO4	Conduct investigations of complex problems	PO10	Communication	3	Substantial/ High
PO5	Modern tool usage	PO11	Project management and finance		
PO6	The Engineer and society	PO12	Life-long learning		
PSO1	Develop applications using different stacks of web and programming technologies				
PSO2	Design and develop secure, parallel, distributed, networked, and digital systems				
PSO3	Apply software engineering methods to design, develop, test and manage software systems.				
PSO4	Develop intelligent applications for business and industry				

CMR INSTITUTE OF TECHNOLOGY

Scheme Of Evaluation
Internal Assessment Test 1 – Sep 2019

Sub:	Cloud Computing and its Applications			**Code:**	15CS742
Date: 14/09/2019	Duration: 90mins	Max Marks: 50	**Sem:** VII A,B & C	**Branch:**	CSE

Note: Answer Any Five Question

Question #		Description	Marks Distribution		Max Marks
1	a)	Definition of Grid computing Definition of distributed computing Two unique differences related to work load management and support of multiple policies	1 1 1.5 + 1.5	5 M	10 M
	b)	Virtualization of a compute environment Virtualization of hardware interfaces Replication of run time environments Virtualization of software stack Components of virtualization in IaaS	1 1 1 1 1	5 M	
2	a)	Pay as you go Revenue comparison Maintenance cost improvement Increased agility Accessibility from any where	1 1 1 1 1	5 M	10 M
	b)	Pros: Managed execution; Portability; efficient use of resources Cons: Performance; phishing ; inefficient use of host	2.5 2.5	5 M	
3	a)	Cloud Interoperability and Standards Scalability and fault tolerance Security trust and privacy (or) Cloud definitio	2 2 2	6 M	10 M
	b)	VMM Design based on type of hypervisors Details about ISA Details about ABI Details about API	1 1 1 1	4 M	
4	a)	Public clouds Private clouds Hybrid clouds Community clouds	1 1 1.5 1.5	5 M	10 M
	b)	Ring diagram Para virtualization Concept of domains	2 2 1	5 M	

5	a)	Equivalence Resource Control Efficiency Theorem 1 Theorem 2 Theorem 3	1 1 1 1 1 1	6 M	10 M
	b)	Network Virtualization Storage Virtualization Desktop Virtualization Operating system level virtualization	1 1 1 1	4 M	
6	a)	IaaS PaaS SaaS	1 1 1	03 M	10 M
	b)	Classification based on services Middleware & Virtual Machine Manager Pure PaaS Pure SaaS Iaas(M)	1.5 2.5 1 1 1	07 M	
7	a)	Definition of Community clouds Picture depicting community cloud Any 3 examples: Healthcare Industry Energy and core industries Public Sector	1 1 1 1 1	05 M	10 M
	b)	Why Hybrid clouds Need of dynamic provisioning Examples like OpenNebula, InterGrid with EC2	1 1 3	05 M	

USN										

CMRIT — CMR INSTITUTE OF TECHNOLOGY, BENGALURU. — ACCREDITED WITH A+ GRADE BY NAAC

Internal Assessment Test 1 – September 2019

Sub:	**Cloud Computing and its Applications**		Sub Code:	**15CS742**	Branch:	**CSE**
Date:	14 / 09 / 2019	Duration: 90 min's	Max Marks: 50	Sem / Sec:	**VII/** A,B & C	OBE

Answer any FIVE FULL Questions

			MARKS	CO	RBT
1	(a)	Compare Distributed Computing and Grid Computing.	[05]	CO2	L2
	(b)	Demonstrate the usage of virtualization in Cloud Computing	[05]	CO2	L3
2	(a)	Discuss in detail about a real-life scenario under which you can rely on cloud services	[05]	CO3	L2
	(b)	Describe the pros and cons of virtualization.	[05]	CO1	L1
3	(a)	Discuss three open challenges in cloud.	[06]	CO1	L2
	(b)	Imagine yourself as a cloud service provider. Illustrate how do you design the virtual machine manager?	[04]	CO2	L3
4	(a)	Classify various types of clouds.	[05]	CO1	L2
	(b)	Discuss the elements of virtualization in Xen	[05]	CO2	L2

			MARKS	CO	RBT
5	(a)	Explain Goldberg and Popek criteria and theorem	[06]	CO2	L2
	(b)	Describe about network, storage, desktop and OS virtualization	[04]	CO1	L2
6	(a)	Name three cloud computing services	[03]	CO1	L1
	(b)	Explain cloud computing architecture with the help of a diagram.	[07]	CO2	L2
7	(a)	Illustrate with example how community clouds are significant in the current computational world?	[05]	CO1	L3
	(b)	Explain with examples how hybrid clouds are used in industry?	[05]	CO1	L2

1 (a)	Compare Distributed Computing and Grid Computing.	[05]	CO2	L2

Definition of Distributed Computing

Distributed Computing is an environment in which a group of independent and geographically dispersed computer systems take part to solve a complex problem, each by solving a part of solution and then combining the results from all computers. These systems are loosely coupled systems coordinately working for a common goal. It can be defined as

1. A computing system in which services are provided by a pool of computers collaborating over a network .

2. A computing environment that may involve computers of differing architectures and data representation formats that share data and system resources.

Definition of Grid Computing

The Basic idea between Grid Computing is to utilize the ideal CPU cycles and storage of millions of computer systems across a worldwide network function as a flexible, pervasive, and inexpensive accessible pool that could be harnessed by anyone who needs it, similar to the way power companies and their users share the electrical grid. There are many definitions of the term: Grid computing:

1. A service for sharing computer power and data storage capacity over the Internet
2. An ambitious and exciting global effort to develop an environment in which individual users can access computers, databases and experimental facilities simply and transparently, without having to consider where those facilities are located.
3. Grid computing is a model for allowing companies to use a large number of computing resources on demand, no matter where they are located.

Traditional distributed computing can be characterized as a subset of grid computing. some of the differences between these two are

1. Distributed Computing normally refers to managing or pooling the hundreds or thousands of computer systems which individually are more limited in their memory and processing power. On the other hand, grid computing has some extra characteristics. It is concerned to efficient utilization of a pool of heterogeneous systems with optimal workload management utilizing an enterprise's entire computational resources(servers, networks, storage, and information) acting together to create one or more large pools of computing resources. There is no limitation of users, departments or originations in grid computing.

2. Grid computing is focused on the ability to support computation across multiple administrative domains that sets it apart from traditional distributed computing. Grids offer a way of using the information technology resources optimally inside an organization involving virtualization of computing resources. Its concept of support for multiple administrative policies and security authentication and authorization mechanisms enables it to be distributed over a local, metropolitan, or wide-area network.

1 (b)	**Demonstrate the usage of virtualization in Cloud Computing**	[05]	CO2	L3

Virtualization is another core technology for cloud computing. It encompasses a collection of solutions allowing the abstraction of some of the fundamental elements for computing, such as hardware, runtime environments, storage, and networking. Virtualization has been around for more than 40 years, but its application has always been limited by technologies that did not allow an efficient use of virtualization solutions. Today these limitations have been substantially overcome, and virtualization has become a fundamental element of cloud computing. This is particularly true for solutions that provide IT infrastructure on demand. Virtualization confers that degree of customization and control that makes cloud computing appealing for users and, at the same time, sustainable for cloud services providers.

Virtualization is essentially a technology that allows creation of different computing environments. These environments are called virtual because they simulate the interface that is expected by a guest. The most common example of virtualization is hardware virtualization. This technology allows simulating the hardware interface expected by an operating system. Hardware virtualization allows the coexistence of different software stacks on top of the same hardware. These stacks are contained inside virtual machine instances, which operate in complete isolation from each other. High-performance servers can host several virtual machine instances, thus creating the opportunity to have a customized software stack on demand. This is the base technology that enables cloud computing solutions to deliver virtual servers on demand, such as Amazon EC2, RightScale,
VMware vCloud, and others. Together with hardware virtualization, storage and network virtualization complete the range of technologies for the emulation of IT infrastructure.

Virtualization technologies are also used to replicate runtime environments for programs. Applications in the case of process virtual machines (which include the foundation of technologies such as Java or .NET), instead of being executed by the operating system, are run by a specific program called a virtual machine. This technique allows isolating the execution of applications and providing finer control on the resources they access. Process virtual machines offer a higher level of abstraction with respect to hardware virtualization, since the guest is only constituted by an application rather than a complete software stack. This approach is used in cloud computing to provide a platform for scaling applications on demand, such as Google AppEngine and Windows Azure. Having isolated and customizable environments with minor impact on performance is what makes virtualization an attractive technology. Cloud computing is realized through platforms that
leverage the basic concepts described above and provides on demand virtualization services to a multitude of users across the globe.

2 (a)	**Discuss in detail about a real-life scenario under which you can rely on cloud services**	[05]	CO3	L2

In a real life scenario, cloud services bring forth huge benefit in infrastructure.

This is mainly because the IT assets, namely software and infrastructure, are turned into utility costs, which are paid for as long as they are used, not paid for up front. Capital costs are costs associated with assets that need to be paid in advance to start a business activity. Before cloud computing, IT infrastructure and software generated capital costs, since they were paid up front so that business start-ups could afford a computing infrastructure, enabling the business activities of the organization. The revenue of the business is then utilized to compensate over time for these costs. Organizations always minimize capital costs, since they are often associated with depreciable values. This is the case of hardware: a server bought today for $1,000 will have a market value less than its original price when it is eventually replaced by new hardware. To make profit, organizations have to compensate for this depreciation created by time, thus reducing the net gain obtained from revenue. Minimizing capital costs, then, is fundamental. Cloud computing transforms IT infrastructure and software into utilities, thus significantly contributing to increasing a company's net gain. Moreover, cloud computing also provides an opportunity for small organizations and start-ups: these do not need large investments to start their business, but they can comfortably grow with it. Finally, maintenance costs are significantly reduced: by renting the infrastructure and the application services, organizations are no longer responsible for their maintenance.

Increased agility in defining and structuring software systems is another significant benefit of cloud computing. Since organizations rent IT services, they can more dynamically and flexibly compose their software systems, without being constrained by capital costs for IT assets.

End users can benefit from cloud computing by having their data and the capability of operating on it always available, from anywhere, at any time, and through multiple devices. Information and services stored in the cloud are exposed to users by Web-based interfaces that make them accessible from portable devices as well as desktops at home.

---

2 (b)	**Describe the pros and cons of virtualization.**	[05]	CO1	L1

<u>Pros</u>

Managed execution and isolation are perhaps the most important advantages of virtualization. In the case of techniques supporting the creation of virtualized execution environments, these two characteristics allow building secure and controllable computing environments. A virtual execution environment can be configured as a sandbox, thus preventing any harmful operation to cross the borders of the virtual host. Moreover, allocation of resources and their partitioning among different guests is simplified, being the virtual host controlled by a program. This enables fine-tuning of resources, which is very important in a server consolidation scenario and is also a requirement for effective quality of service.

Portability is another advantage of virtualization, especially for execution virtualization techniques. Virtual machine instances are normally represented by one or more files that can be easily transported with respect to physical systems. Moreover, they also tend to be self-contained since they do not have other dependencies besides the virtual machine manager for their use. Portability and self-containment simplify their administration. Java programs are "compiled once and run every where"; they only require that the Java virtual machine be installed on the host. The same applies to hardware-level virtualization. It is in fact possible to build our own operating environment within a virtual machine instance and bring it with us wherever we go, as though we had our own laptop. This concept is also an enabler for migration techniques in a server consolidation scenario. Portability and self-containment also contribute to reducing the costs of maintenance, since the number of hosts is expected to be lower than the number of virtual machine instances. Since the
guest program is executed in a virtual environment, there is very limited opportunity for the guest program to damage the underlying hardware. Moreover, it is expected that there will be fewer virtual machine managers with respect to the number of virtual machine instances managed.

Finally, by means of virtualization it is possible to achieve a more **efficient use of resources**. Multiple systems can securely coexist and share the resources of the underlying host, without interfering with each other. This is a prerequisite for server consolidation, which allows adjusting the number of active physical resources dynamically according to the current load of the system, thus creating the opportunity to save in terms of energy consumption and to be less impacting on the environment.

<u>Cons</u>

Performance is definitely 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. 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. These concerns are becoming less and less important thanks to technology advancements and the ever-increasing computational power available today.

Virtualization can sometimes lead to an **inefficient use of the host**. In particular, 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 sometimes 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 libraries are used.

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. In the case of hardware virtualization, malicious programs can preload themselves before the operating system and act as a thin virtual machine manager toward it. The operating system is then
controlled and can be manipulated to extract sensitive information of interest to third parties. Examples of these kinds of malware are BluePill and SubVirt.

---

3 (a)	**Discuss three open challenges in cloud.**	[06]	CO1	L2

Cloud Definition

there have been several attempts made to define cloud computing and to provide a classification of all the services and technologies identified as such. One of the most comprehensive formalizations is noted in the NIST working definition of cloud computing. It characterizes cloud computing as on-demand self-service, broad network access, resource-pooling, rapid elasticity, and measured service; classifies services as SaaS, PaaS, and IaaS; and categorizes deployment models as public, private, community, and hybrid clouds. The view is in line with our discussion and shared by many IT practitioners and academics.

Even after many definitions, cloud can not be defined efficiently.

Cloud Interoperability and Standards

To fully realize this goal, introducing standards and allowing interoperability between solutions offered by different vendors are objectives of fundamental importance. Vendor lock-in constitutes one of the major strategic barriers against the seamless adoption of cloud computing at all stages. In particular there is major fear on the part of enterprises in which IT constitutes the significant part of their revenues. Vendor lock-in can prevent a customer from switching to another competitor's solution, or when this is possible, it happens at
considerable conversion cost and requires significant amounts of time.

Scalability and fault tolerance

The ability to scale on demand constitutes one of the most attractive features of cloud computing. Clouds allow scaling beyond the limits of the existing in-house IT resources, whether they are infrastructure (compute and storage) or applications services. To implement such a capability, the cloud middleware has to be designed with the principle of scalability along different dimensions in mind—for example, performance, size, and load. The cloud middleware manages a huge number of resource and users, which rely on the cloud to obtain the horsepower that they cannot obtain within the premises without bearing considerable administrative and maintenance costs. These costs are a reality for whomever develops, manages, and maintains the cloud middleware and offers the service to customers. In this scenario, the ability to tolerate failure becomes fundamental, sometimes even more important than providing an extremely efficient and optimized system. Hence, the challenge in this case is designing highly scalable and fault-tolerant systems that are easy to manage and at the same time provide competitive performance.

Security, trust and privacy

Security, trust, and privacy issues are major obstacles for massive adoption of cloud computing. The traditional cryptographic technologies are used to prevent data tampering and access to sensitive information. The massive use of virtualization technologies exposes the existing system to new threats, which previously were not considered applicable. For example, it might be possible that applications hosted in the cloud can process sensitive information; such information can be stored within a cloud storage facility using the most advanced technology in cryptography to protect data and then be considered safe from any attempt to access it without the required permissions. Although these data are processed in memory, they must necessarily be decrypted by the legitimate application, but since the application is hosted in a managed virtual environment it becomes accessible to the virtual machine manager that by program is designed to access the memory pages of such an application. In this case, what is experienced is a lack of control over the environment in which the application is executed, which is made possible by leveraging the cloud.

---

3 b) Imagine yourself as a cloud service provider. Illustrate how do you design the virtual machine manager?

[04]	CO2	L3

The basic criteria to consider before designing virtual machine manager (VMM) would be to decide whether I want to keep the VMM directly on the hardware or whether I would like to keep at as part of Operating System. The first one is called as Type 1 Hypervisor and the second is referred as Type 2 Hypervisor.

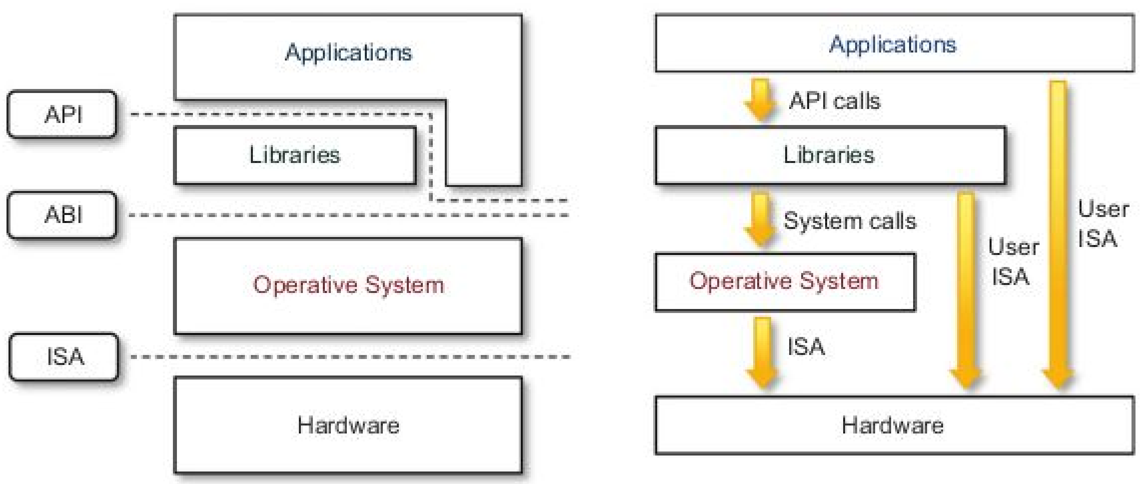

Virtualizing an execution environment at different levels of the computing stack requires a reference model that defines the interfaces between the levels of abstractions, which hide implementation details. From this perspective, virtualization techniques actually replace one of the layers and intercept the calls that are directed toward it. Therefore, a clear separation between layers simplifies their implementation, which only requires the emulation of the interfaces and a proper interaction with the underlying layer.

At the bottom layer, the model for the hardware is expressed in terms of the Instruction Set Architecture (ISA), which defines the instruction set for the processor, registers, memory, and interrupt management. ISA is the interface between hardware and software, and it is important to the operating system (OS) developer (System ISA) and developers of applications that directly manage the underlying hardware (User ISA). The application binary interface (ABI) separates the operating system layer from the applications and libraries, which are managed by the OS. ABI covers details such as low-level data types, alignment, and call conventions and defines a format for executable programs. System calls are defined at this level. This interface allows portability of applications and libraries across operating systems that implement the same ABI. The highest level of abstraction is represented by the application programming interface (API), which interfaces applications to libraries and/or the underlying operating system.

For any operation to be performed in the application level API, ABI and ISA are responsible for making it happen. The high-level abstraction is converted into machine-level instructions to perform the actual operations supported by the processor. The machine-level resources, such as processor registers and main memory capacities, are used to perform the operation at the hardware level of the central processing unit (CPU). This layered approach simplifies the development and implementation of computing systems and simplifies the implementation of multitasking and the coexistence of multiple executing environments. In fact, such a model not only requires limited knowledge of the entire computing stack, but it also provides ways to implement a minimal security model for managing and accessing shared resources.

---

4 (a)	**Classify various types of clouds.**	[05]	CO1	L2

Public clouds

Public clouds constitute the first expression of cloud computing. They are a realization of the canonical view of cloud computing in which the services offered are made available to anyone, from anywhere, and at any time through the Internet. From a structural point of view they are a distributed system, most likely composed of one or more data centers connected together, on top of which the specific services offered by the cloud are implemented. Any customer can easily sign in with the cloud provider, enter her credential and billing details, and use the services offered

Private clouds

Public clouds are appealing and provide a viable option to cut IT costs and reduce capital expenses, but they are not applicable in all scenarios. For example, a very common critique to the use of cloud computing in its canonical implementation is the loss of control. In the case of public clouds, the provider is in control of the infrastructure and, eventually, of the customers' core logic and sensitive data. Even though there could be regulatory procedure in place that guarantees fair management and respect of the customer's privacy, this condition can still be perceived as a threat or as an unacceptable risk that some organizations are not willing to take. In particular, institutions such as government and military agencies will not consider public clouds as an option for processing or storing their sensitive data. The risk of a breach in the security infrastructure of the provider could expose such information to others; this could simply be considered unacceptable.

Hybrid clouds

Hybrid clouds allow enterprises to exploit existing IT infrastructures, maintain sensitive information within the premises, and naturally grow and shrink by provisioning external resources and releasing them when they're no longer needed. Security concerns are then only limited to the public portion of the cloud that can be used to perform operations with less stringent constraints but that are still part of the system workload

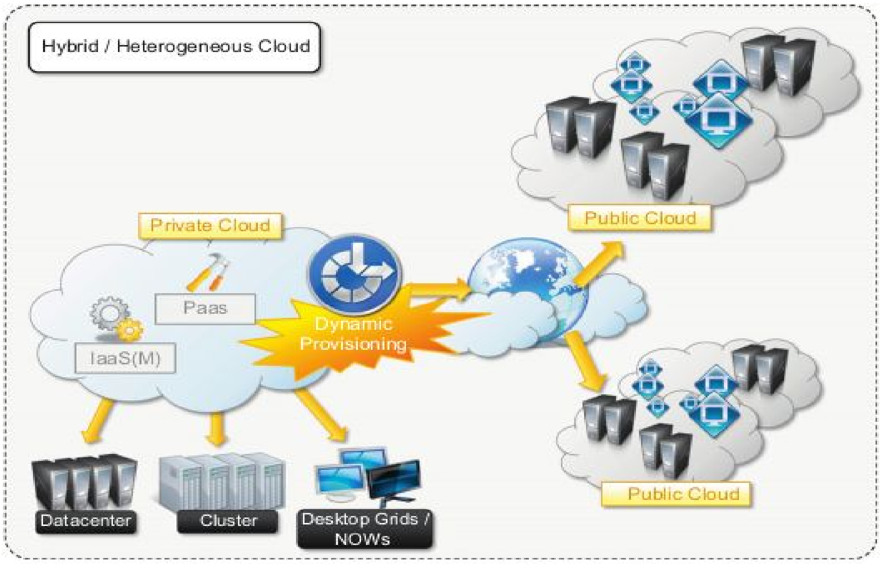

Community clouds

Community clouds are distributed systems created by integrating the services of different clouds to address the specific needs of an industry, a community, or a business sector. The users of a specific community cloud fall into a well-identified community, sharing the same concerns or needs; they can be government bodies, industries, or even simple users, but all of them focus on the same issues for their interaction with the cloud. This is a different scenario than public clouds, which serve a multitude of users with different needs. Community clouds are also different from private clouds, where the services are generally delivered within the
institution that owns the cloud.

Media industry: In the media industry, companies are looking for low-cost, agile, and simple solutions to improve the efficiency of content production. Most media productions involve an extended ecosystem of partners. In particular, the creation of digital content is the outcome of a collaborative process that includes movement of large data, massive compute-intensive rendering tasks, and complex workflow executions. Community clouds can provide a shared environment where services can facilitate business-to-business collaboration and offer the horsepower in terms of aggregate bandwidth, CPU, and storage required to efficiently support media production.

---

4 (b)	**Discuss the elements of virtualization in Xen**	[05]	CO2	L2

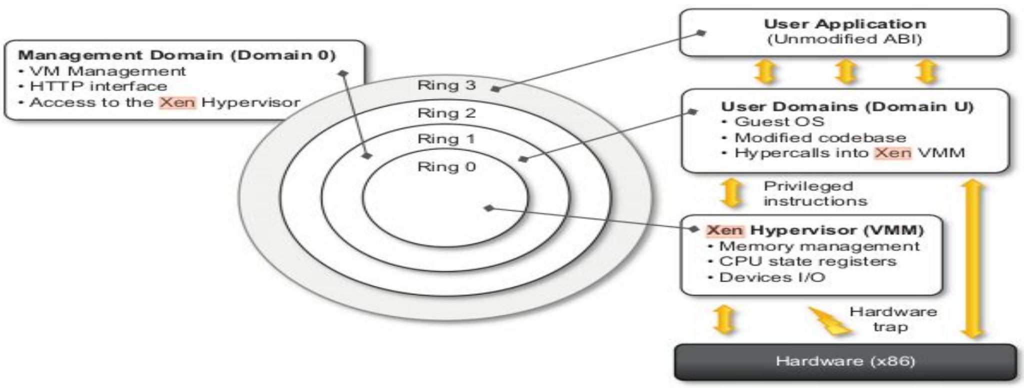

Xen is the most popular implementation of paravirtualization, which, in contrast with full virtualization, allows high-performance execution of guest operating systems. This is made possible by eliminating the performance loss while executing instructions that require special management. This is done by modifying portions of the guest operating systems run by Xen with reference to the execution of such instructions. Therefore it is not a transparent solution for implementing virtualization. This is particularly true for x86, which is the most popular architecture on commodity machines and servers.

Above figure describes the architecture of Xen and its mapping onto a classic x86 privilege model. A Xen-based system is managed by the Xen hypervisor, which runs in the highest privileged mode and controls the access of guest operating system to the underlying hardware. Guest operating systems are executed within domains, which represent virtual machine instances.

Moreover, specific control software, which has privileged access to the host and controls all the other guest operating systems, is executed in a special domain called Domain 0. This is the first one that is loaded once the virtual machine manager has completely booted, and it hosts a HyperText Transfer Protocol (HTTP) server that serves requests for virtual machine creation, configuration, and termination. This component constitutes the embryonic version of a distributed virtual machine manager, which is an essential component of cloud computing systems providing Infrastructure-as-a-Service (IaaS) solutions. Many of the x86 implementations support four different security levels, called rings, where Ring 0 represent the level with the highest privileges and Ring 3 the level with the lowest ones. Almost all the most popular operating systems, except OS/2, utilize only two levels: Ring 0 for the kernel code, and Ring 3 for user application and nonprivileged OS code. This provides the opportunity for Xen to implement virtualization by executing the hypervisor in Ring 0, Domain 0, and all the other domains running guest operating systems—generally referred to as Domain U—in Ring 1,
while the user applications are run in Ring 3.

Paravirtualization needs the operating system codebase to be modified, and hence not all operating systems can be used as guests in a Xen-based environment. More precisely, this condition holds in a scenario where it is not possible to leverage hardware-assisted virtualization, which allows running the hypervisor in Ring -1 and the guest operating system in Ring 0. Therefore, Xen exhibits some limitations in the case of legacy hardware and legacy operating systems.

5 (a)	**Explain Goldberg and Popek criteria and theorem**	[06]	CO2	L2

<u>Criteria:</u>

Equivalence: A guest running under the control of a virtual machine manager should exhibit the same behavior as when it is executed directly on the physical host.
Resource control: The virtual machine manager should be in complete control of virtualized resources.
Efficiency: A statistically dominant fraction of machine instructions should be executed without intervention from the virtual machine manager.

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 instructions.

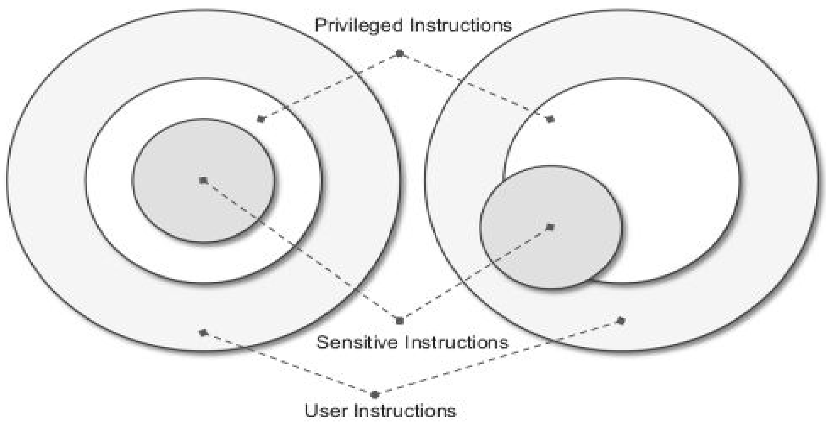

THEOREM 2
A conventional third-generation computer is recursively virtualizable if:

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

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.

5 (b)	**Describe about network, storage, desktop and OS virtualization**	[04]	CO1	L2

Storage virtualization is a system administration practice that allows decoupling the physical organization of the hardware from its logical representation. Using this technique, users do not have to be worried about the specific location of their data, which can be identified using a logical path.

Network virtualization combines hardware appliances and specific software for the creation and management of a virtual network. Network virtualization can aggregate different physical networks into a single logical network (external network virtualization) or provide network-like functionality to an operating system partition (internal network virtualization).

Desktop virtualization abstracts the desktop environment available on a personal computer in order to provide access to it using a client/server approach. Desktop virtualization provides the same outcome of hardware virtualization but serves a different purpose. Similarly to hardware virtualization, desktop virtualization makes accessible a different system as though it were natively installed on the host, but this system is remotely stored on a different host and accessed through a network connection. Moreover, desktop virtualization addresses the problem of making the same desktop environment accessible from everywhere. Although the term desktop virtualization strictly refers to the ability to remotely access a desktop environment, generally the desktop environment is stored in a remote server or a data center that provides a high-availability infrastructure and ensures the accessibility and persistence of the data.

Operating system-level virtualization offers the opportunity to create different and separated execution environments for applications that are managed concurrently. Differently from hardware virtualization, there is no virtual machine manager or hypervisor, and the virtualization is done within a single operating system, where the OS kernel allows for multiple isolated user space instances. The kernel is also responsible for sharing the system resources among instances and for limiting the impact of instances on each other. A user space instance in general contains a proper view of the file system, which is completely isolated, and separate IP addresses, software configurations, and access to devices. Operating systems supporting this type of virtualization are general-purpose, time shared operating systems with the capability to provide stronger namespace and resource isolation.

E.g FreeBSD Jails

6 (a)	**Name three cloud computing services**	[03]	CO1	L1

1. Infrastructure as a Service (IaaS)
2. Platform as a Service (PaaS)
3. Software as a Service (SaaS)

6 (b)	**Explain cloud computing architecture with the help of a diagram.**	[07]	CO2	L2

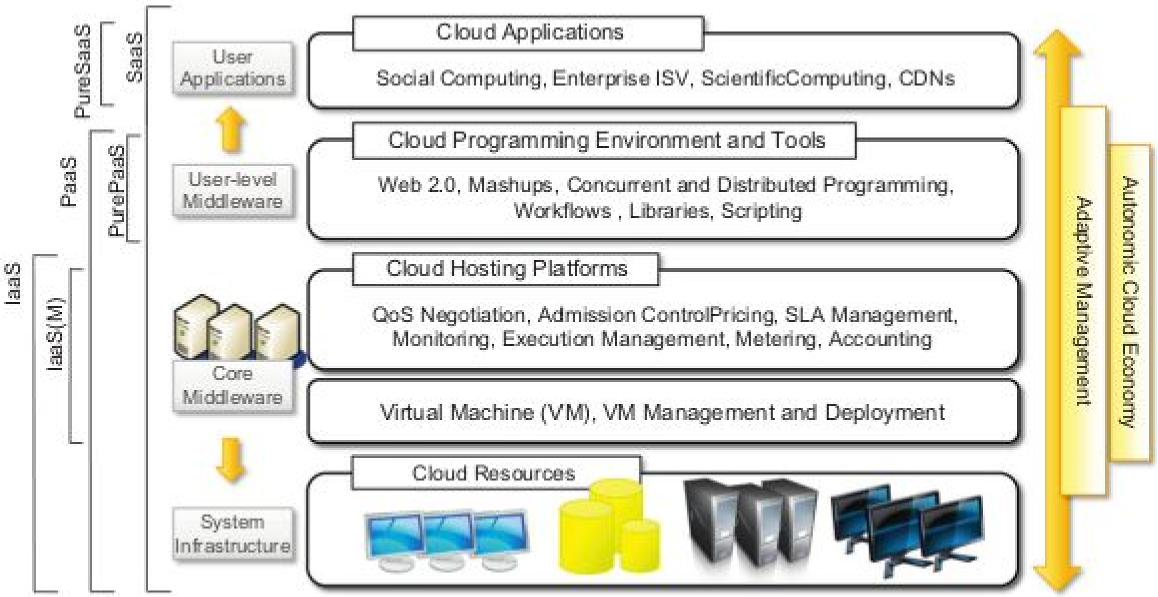

The physical infrastructure is managed by the core middleware, the objectives of which are to provide an appropriate runtime environment for applications and to best utilize resources. At the bottom of the stack, virtualization technologies are used to guarantee runtime environment customization, application isolation, sandboxing, and quality of service. Hardware virtualization is most commonly used at this level. Hypervisors manage the pool of resources and expose the distributed infrastructure as a collection of virtual machines. By using virtual machine technology it is possible to finely partition the hardware resources such as CPU and memory and to virtualize specific devices, thus meeting the requirements of users and applications. This solution is generally paired with storage and network virtualization strategies, which allow the infrastructure to be completely virtualized and controlled. According to the specific service offered to end users, other virtualization techniques can be used; for example, programming-level virtualization helps in creating a portable runtime environment where applications can be run and controlled. This scenario generally
implies that applications hosted in the cloud be developed with a specific technology or a programming language, such as Java, .NET, or Python. In this case, the user does not have to build its system from bare metal. Infrastructure management is the key function of core middleware, which supports capabilities such

as negotiation of the quality of service, admission control, execution management and monitoring, accounting, and billing.
The combination of cloud hosting platforms and resources is generally classified as a Infrastructure-as-a-Service (IaaS) solution. We can organize the different examples of IaaS into two categories: Some of them provide both the management layer and the physical infrastructure; others provide only the management layer (IaaS (M)). In this second case, the management layer is often integrated with other IaaS solutions that provide physical infrastructure and adds value to them.

IaaS solutions are suitable for designing the system infrastructure but provide limited services to build applications. Such service is provided by cloud programming environments and tools, which form a new layer for offering users a development platform for applications. The range of tools include Web-based interfaces, command-line tools, and frameworks for concurrent and distributed programming. In this scenario, users develop their applications specifically for the cloud by using the API exposed at the user-level middleware. For this reason, this approach is also known as Platform-as-a-Service (PaaS) because the service offered to the user is a development platform rather than an infrastructure. PaaS solutions generally include the infrastructure as well, which is bundled as part of the service provided to users. In the case of Pure PaaS, only the user-level middleware is offered, and it has to be complemented with a virtual or physical infrastructure.

Figure contains services delivered at the application level. These are mostly referred to as Software-as-a-Service (SaaS). In most cases these are Web-based applications that rely on the cloud to provide service to end users. The horsepower of the cloud provided by IaaS and PaaS solutions allows independent software vendors to deliver their application services over the Internet.

Table 4.1 Cloud Computing Services Classification

Category	Characteristics	Product Type	Vendors and Products
SaaS	Customers are provided with applications that are accessible anytime and from anywhere.	Web applications and services (Web 2.0)	SalesForce.com (CRM) Clarizen.com (project management) Google Apps
PaaS	Customers are provided with a platform for developing applications hosted in the cloud.	Programming APIs and frameworks Deployment systems	Google AppEngine Microsoft Azure Manjrasoft Aneka Data Synapse
IaaS/HaaS	Customers are provided with virtualized hardware and storage on top of which they can build their infrastructure.	Virtual machine management infrastructure Storage management Network management	Amazon EC2 and S3 GoGrid Nirvanix

---

7 (a)	**Illustrate with example how community clouds are significant in the current computational world?**	[05]	CO1	L3

Community clouds are distributed systems created by integrating the services of different clouds to address the specific needs of an industry, a community, or a business sector. Figure below provides a general view of the usage scenario of community clouds, together with reference architecture. The users of a specific community cloud fall into a well-identified community, sharing the same concerns or needs; they can be government bodies, industries, or even simple users, but all of them focus on the same issues for their interaction with the cloud. This is a different scenario than public clouds, which serve a multitude of users with different needs.

Community clouds are also different from private clouds, where the services are generally delivered within the
institution that owns the cloud. From an architectural point of view, a community cloud is most likely implemented over multiple administrative domains. This means that different organizations such as government bodies, private enterprises, research organizations, and even public virtual infrastructure providers contrib-
ute with their resources to build the cloud infrastructure.

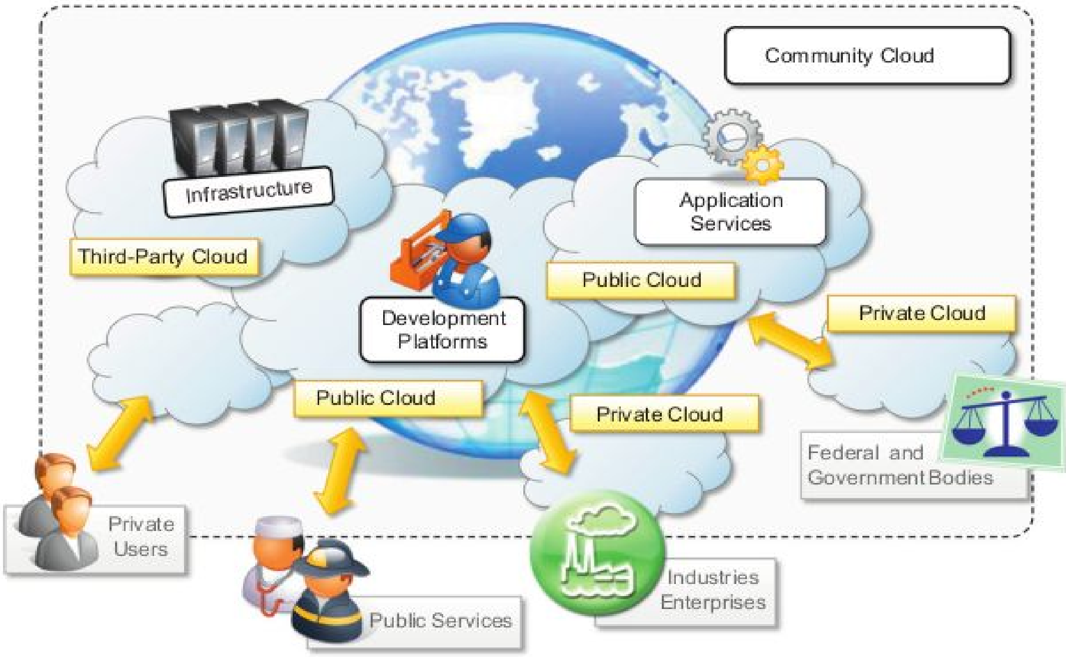

Healthcare industry: In the healthcare industry, there are different scenarios in which community clouds could be of use. In particular, community clouds can provide a global platform on which to share information and knowledge without revealing sensitive data maintained within the private infrastructure. The naturally hybrid deployment model of community clouds can easily support the storing of patient-related data in a private cloud while using the shared infrastructure for noncritical services and automating processes within
hospitals.

Energy and other core industries: In these sectors, community clouds can bundle the comprehensive set of solutions that together vertically address management, deployment, and orchestration of services and operations. Since these industries involve different providers, vendors, and organizations, a community cloud can provide the right type of infrastructure to create an open and fair market.

Public sector: Legal and political restrictions in the public sector can limit the adoption of public cloud offerings. Moreover, governmental processes involve several institutions and agencies and are aimed at providing strategic solutions at local, national, and international administrative levels. They involve business-to-administration, citizen-to-administration, and possibly business-to-business processes. Some examples include invoice approval, infrastructure planning, and public hearings. A community cloud can constitute the optimal venue to provide a distributed environment in which to create a communication platform for performing such operations.

7	(b)	**Explain with examples how hybrid clouds are used in industry?**	[05]	CO1	L2

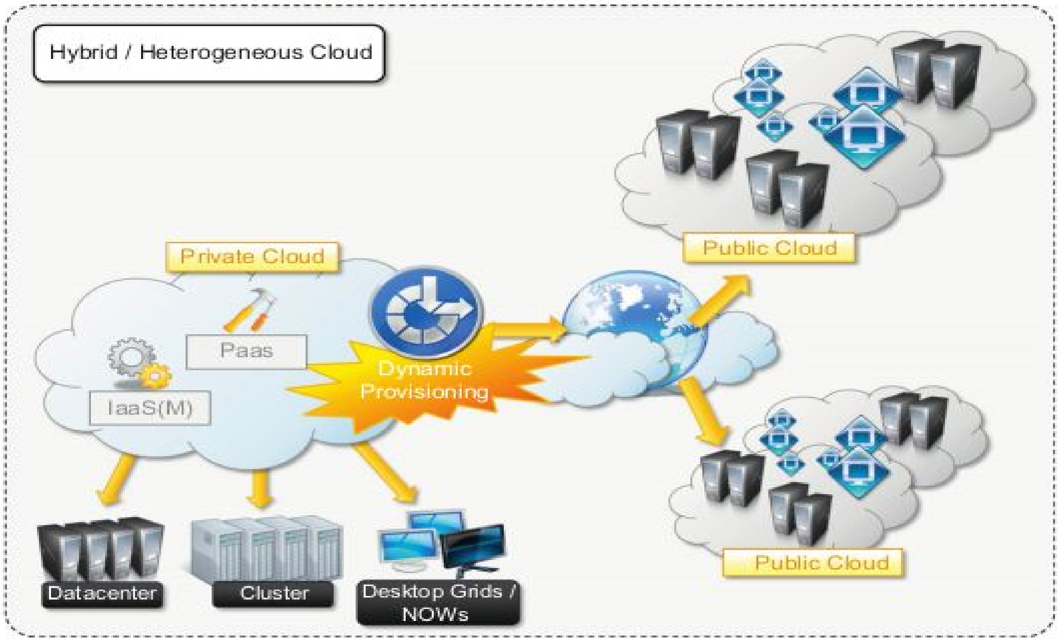

It is a heterogeneous distributed system resulting from a private cloud that integrates additional services or resources from one or more public clouds. For this reason they are also called heterogeneous clouds. As depicted in the diagram, dynamic provisioning is a fundamental component in this scenario. Hybrid clouds address scalability issues by leveraging external resources for exceeding capacity demand. These resources or services are temporarily leased for the time required and then released. This practice is also known as cloud bursting. Whereas the concept of hybrid cloud is general, it mostly applies to IT infrastructure rather than
software services. Service-oriented computing already introduces the concept of integration of paid software services with existing application deployed in the private premises. In an IaaS scenario, dynamic provisioning refers to the ability to acquire on demand virtual machines in order to increase the capability of the resulting distributed system and then release them. Infrastructure management software and PaaS solutions are the building blocks for deploying and managing hybrid clouds. In particular, with respect to private clouds, dynamic provisioning introduces a more complex scheduling algorithm and policies, the goal of which is also to optimize the budget spent to rent public resources.

Infrastructure management software such as OpenNebula already exposes the capability of integrating resources from public clouds such as Amazon EC2. In this case the virtual machine obtained from the public infrastructure is managed as all the other virtual machine instances maintained locally.

In the case of OpenNebula, advanced schedulers such as Haizea can be integrated to provide cost-based scheduling. A different approach is taken by InterGrid. This is essentially a distributed scheduling engine that manages the allocation of virtual machines in a collection of peer networks. Such networks can be represented by a local cluster, a gateway to a public cloud, or a combination of the two. Once a request is submitted to one of the InterGrid gateways, it is served by possibly allocating virtual instances in all the peered networks, and the allocation of requests is performed by taking into account the user budget and the peering arrangements between networks.