Scheme of Evaluation
Internal Assessment Test 1 – May 2021

Sub:	Cloud Computing and its Application							**Code:**	18CS643
Date:	20/05/2020	Duration:	90mins	Max Marks:	50	**Sem:**	VI	**Branch:**	ISE

Note: **Answer Any five full questions.**

Question #		**Description**	**Marks Distribution**		**Max Marks**
1	a)	"Cloud Computing" what do you understand by this term? What is the Cloud Computing Reference model? Summarize with the neat diagram. Cloud Computing Definition-with explanation Cloud Reference model diagram Explanation	2M 4M 4M	10M	10M
2	a)	Discuss live migration and server consolidation with a neat diagram. Definition live migration Server Consolidation Diagram Explanation	1M 1M 2M 2M	6M	10M
2	b)	Discuss major milestones which have lead to cloud computing. Mainframes Cluster Computing Grid Computing	2M 1M 1M	4M	
3	a)	With neat diagram, explain Xen architecture and guest OS management. Xen Definition Xen Architecture Diagram Xen Architecture Explanation	2M 2M 2M	6M	10M

3	b)	Explain the major deployment models in cloud computing. Public Cloud Private Cloud Hybrid Cloud	2M 1M 1M	4M	
4	a)	Describe the characteristics of virtualized environment with neat required diagrams. Characteristics of virtualization Increased Security Managed Execution Sharing Aggregation Emulation Isolation Portability	2M 3M 1M	6M	10M
4	b)	List and explain pros and cons of virtualization. Advantages of virtualization Disadvantages Performance degradation Inefficiency and degraded user experience Security holes and new threats	1M 1M 1M 1M	4M	
5	a)	Analyze the open challenges in cloud computing. Cloud interoperability Scalability & Fault Tolerance Security, trust and privacy Organizational aspects	2M 2M 2M	6M	10M
5	b)	What are the benefits of community cloud? Openness Community Graceful failures Convenience and control Environmental sustainability (Any four)	1M*4	4M	

6	a)	The cloud service provider (CSP) wants to provide IaaS. Discuss the IaaS based solution for cloud computing which CSP will provide with example. IaaS- Explanation IaaS solution Diagram Explanation of the solution	1M 3M 3M	7M	10M
6	b)	Compare and contrast between IaaS, PaaS and SaaS? Any 3 differences	1M*3	3M	

Scheme Of Evaluation
Internal Assessment Test 1 – May 2021

Sub:	Cloud Computing and its Application							Code:	18CS643
Date:	20/05/2021	Duration:	90mins	Max Marks:	50	Sem:	VI	Branch:	ISE

Note: **Answer Any full five questions**

Q. 1 Define Cloud Computing. Explain cloud computing reference model with neat diagram.

Cloud computing is the delivery of different services through the Internet. These resources include tools and applications like data storage, servers, databases, networking, and software.

1.1.4 The cloud computing reference model

A fundamental characteristic of cloud computing is the capability to deliver, on demand, a variety of IT services that are quite diverse from each other. This variety creates different perceptions of what cloud computing is among users. Despite this lack of uniformity, it is possible to classify cloud computing services offerings into three major categories: *Infrastructure-as-a-Service (IaaS)*, *Platform-as-a-Service (PaaS)*, and *Software-as-a-Service (SaaS)*. These categories are related to each other as described in Figure 1.5, which provides an organic view of cloud computing. We refer to this diagram as the *Cloud Computing Reference Model*, and we will use it throughout the

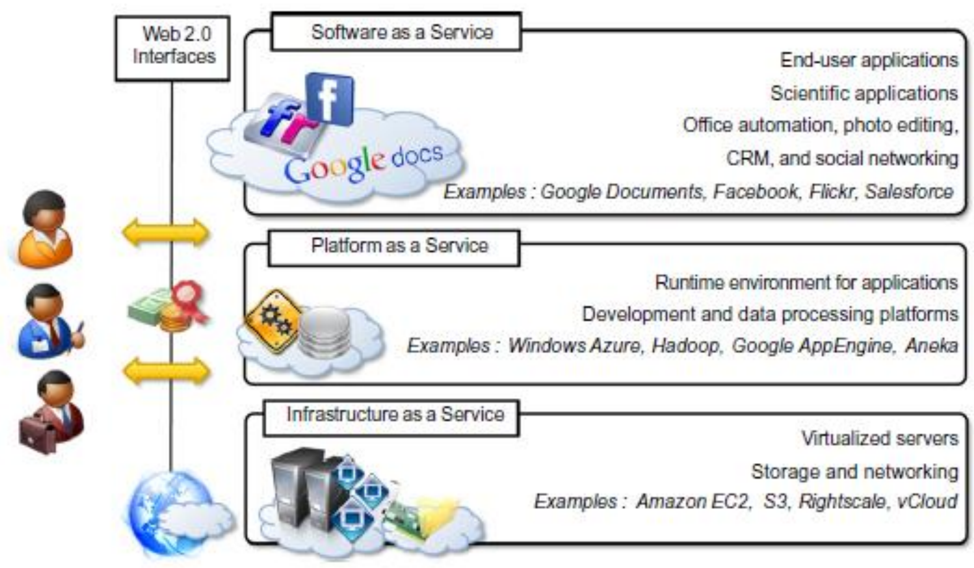

FIGURE 1.5

The Cloud Computing Reference Model.

At the base of the stack, *Infrastructure-as-a-Service* solutions deliver infrastructure on demand in the form of virtual *hardware*, *storage*, and *networking*. Virtual hardware is utilized to provide compute on demand in the form of virtual machine instances. These are created at users' request on the provider's infrastructure, and users are given tools and interfaces to configure the software stack installed in the virtual machine. The pricing model is usually defined in terms of dollars per hour, where the hourly cost is influenced by the characteristics of the virtual hardware. Virtual storage is delivered in the form of raw disk space or object store. The former complements a virtual hardware offering that requires persistent storage. The latter is a more high-level abstraction for storing entities rather than files. Virtual networking identifies the collection of services that manage the networking among virtual instances and their connectivity to the Internet or private networks.

Platform-as-a-Service solutions are the next step in the stack. They deliver scalable and elastic runtime environments on demand and host the execution of applications. These services are backed by a core middleware platform that is responsible for creating the abstract environment where applications are deployed and executed. It is the responsibility of the service provider to provide scalability and to manage fault tolerance, while users are requested to focus on the logic of the application developed by leveraging the provider's APIs and libraries. This approach increases the level of abstraction at which cloud computing is leveraged but also constrains the user in a more controlled environment.

At the top of the stack, *Software-as-a-Service* solutions provide applications and services on demand. Most of the common functionalities of desktop applications—such as office automation, document management, photo editing, and customer relationship management (CRM) software—are replicated on the provider's infrastructure and made more scalable and accessible through a browser on demand. These applications are shared across multiple users whose interaction is isolated from the other users. The SaaS layer is also the area of social networking Websites, which leverage cloud-based infrastructures to sustain the load generated by their popularity.

Each layer provides a different service to users. IaaS solutions are sought by users who want to leverage cloud computing from building dynamically scalable computing systems requiring a specific software stack. IaaS services are therefore used to develop scalable Websites or for background processing. PaaS solutions provide scalable programming platforms for developing applications and are more appropriate when new systems have to be developed. SaaS solutions target mostly end users who want to benefit from the elastic scalability of the cloud without doing any software development, installation, configuration, and maintenance. This solution is appropriate when there are existing SaaS services that fit users needs (such as email, document management, CRM, etc.) and a minimum level of customization is needed.

Q. 2 a) Discuss live migration and server consolidation with a neat diagram.

Virtualization plays an important role in cloud computing since it allows for the appropriate degree of customization, security, isolation, and manageability that are fundamental for delivering IT services on demand. Particularly important is the role of virtual computing environment and execution virtualization techniques. Among these, hardware and programming language virtualization are the techniques adopted in cloud computing systems.
Besides being an enabler for computation on demand, virtualization also gives the opportunity to design more efficient computing systems by means of consolidation, which is performed transparently to cloud computing service users.

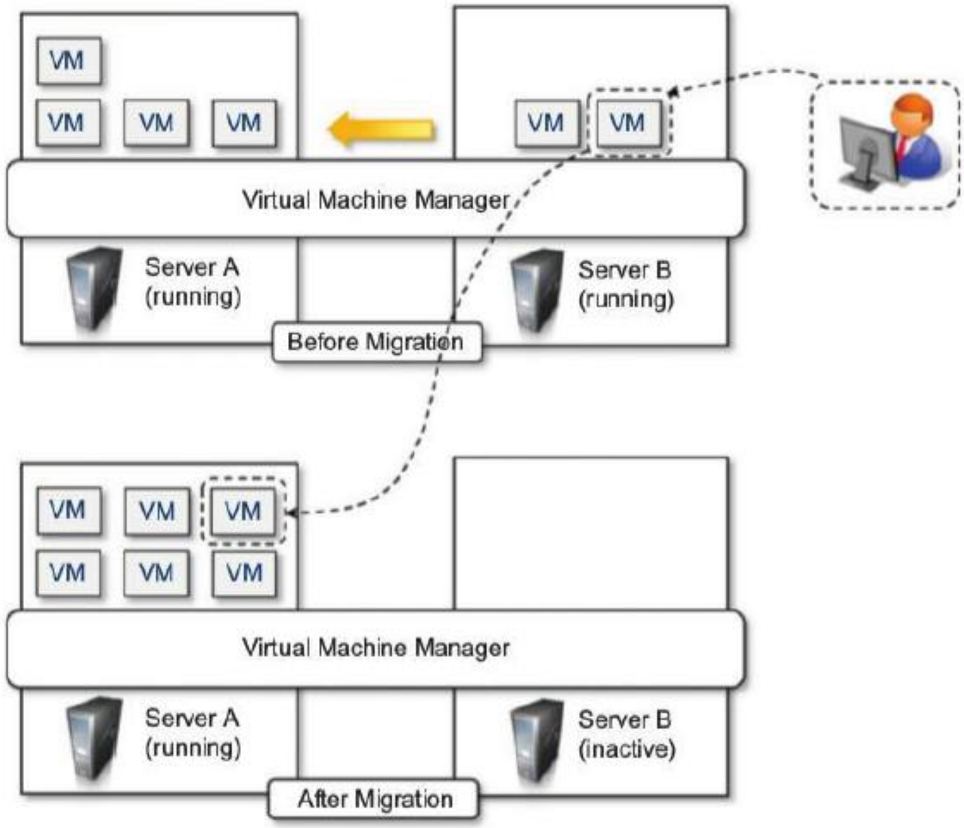

FIGURE 3.10

Live migration and server consolidation.

Since virtualization allows us to create isolated and controllable environments, it is possible to serve these environments with the same resource without them interfering with each other. This opportunity is particularly attractive when resources are underutilized, because it allows reducing the number of active resources by aggregating virtual machines over a smaller number of resources that become fully utilized. This practice is also known as server consolidation, while the movement of virtual machine instances is called virtual machine migration (see Figure 3.10). Because virtual machine instances are controllable environments, consolidation can be applied with a minimum impact, either by temporarily stopping its execution and moving its data to the new resources or by performing a finer control and moving the instance while it is running. This second techniques is known as live migration and in general is more complex to implement but more efficient since there is no disruption of the activity of the virtual machine instance.

Q. 2 b) Discuss major milestones which have lead to cloud computing.

Three major milestones have led to cloud computing: mainframe computing, cluster computing, and grid computing.

- **Mainframes:**These were the first examples of large computational facilities leveraging multipleprocessing units. Mainframes were powerful, highly reliable computers specialized

for large data movement and massive input/output (I/O) operations. They were mostly used by large organizations for bulk data processing tasks such as online transactions, enterprise resource planning, and other operations involving the processing of significant amounts of data.

- **Clusters:**Cluster computing started as a low-cost alternative to the use of mainframes andsupercomputers. The technology advancement that created faster and more powerful mainframes and supercomputers eventually generated an increased availability of cheap commodity machines as a side effect. These machines could then be connected by a high-bandwidth network and controlled by specific software tools that manage them as a single system. Starting in the 1980s.

  Cluster technology contributed considerably to the evolution of tools and frameworks for distributed computing, including Condor, Parallel Virtual Machine (PVM), and Message Passing Interface (MPI).

- **Grid computing** appeared in the early 1990s as an evolution of cluster computing. In an analogy tothe power grid, grid computing proposed a new approach to access large computational power, hugeBstorage facilities, and a variety of services.
  A computing grid was a dynamic aggregation of heterogeneous computing nodes, and its scale was nationwide or even worldwide. Several developments made possible the diffusion of computing grids:
  a. clusters became quite common resources;
  b. they were often underutilized;
  c. new problems were requiring computational power that went beyond the capability of single clusters and
  d. the improvements in networking and the diffusion of the Internet made possible long-distance, high-bandwidth connectivity.

Q. 3a) With neat diagram, explain Xen architecture and guest OS management.

Xen: paravirtualization

Xen is an open-source initiative implementing a virtualization platform based on paravirtualization. Initially developed by a group of researchers at the University of Cambridge in the United Kingdom, Xen now has a large open-source community backing it. . Xen-based technology is used for either desktop virtualization or server virtualization, and recently it has also been used to provide cloud computing solutions by means of Xen Cloud Platform (XCP).

Figure 2.11 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.

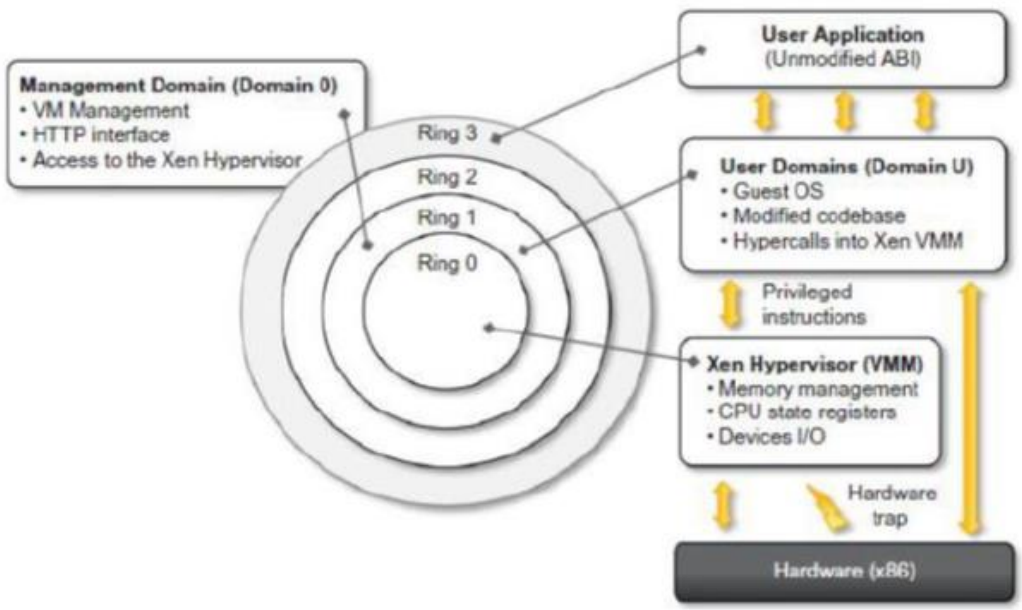

Figure 2.11 Xen architecture and guest OS management

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, con- figuration, 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.

Because of the structure of the x86 instruction set, some instructions allow code executing in Ring 3 to jump into Ring 0 (kernel mode). Such operation is performed at the hardware level and therefore within a virtualized environment will result in a trap or silent fault, thus preventing the normal operations of the guest operating system, since this is now running in Ring 1.This condition is generally triggered by a subset of the system calls. To avoid this situation, operating systems need to be changed in their implementation, and the sensitive system calls need to be reimplemented with hypercalls, which are specific calls exposed by the virtual machine interface of Xen. With the use of hypercalls, the Xen hypervisor is able to catch the execution of all the sensitive instructions, manage them, and return the control to the guest operating system by means of a supplied handler.

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. Open-source operating systems such as Linux can be easily modified, since their code is publicly available and Xen provides full support for their virtualization, whereas components of the Windows family are generally not supported by Xen unless hardware-assisted virtualization is available.

Q. 3 b) Explain the major deployment models in cloud computing.

The three major models for deploying and accessing cloud computing environments are public clouds, private/enterprise clouds, and hybrid clouds (see Figure 1.4). *Public clouds* are the most common deployment models in which necessary IT infrastructure (e.g., virtualized datacenters) is established by a third-party service provider that makes it available to any consumer on a subscription basis. Such clouds are appealing to users because they allow users to quickly leverage compute, storage, and application services. In this environment, users' data and applications are deployed on cloud datacenters on the vendor's premises.

Large organizations that own massive computing infrastructures can still benefit from cloud computing by replicating the cloud IT service delivery model in-house. This idea has given birth to the concept of *private clouds* as opposed to public clouds. In 2010, for example, the U.S. federal government, one of the world's largest consumers of IT spending (around $76 billion on more than

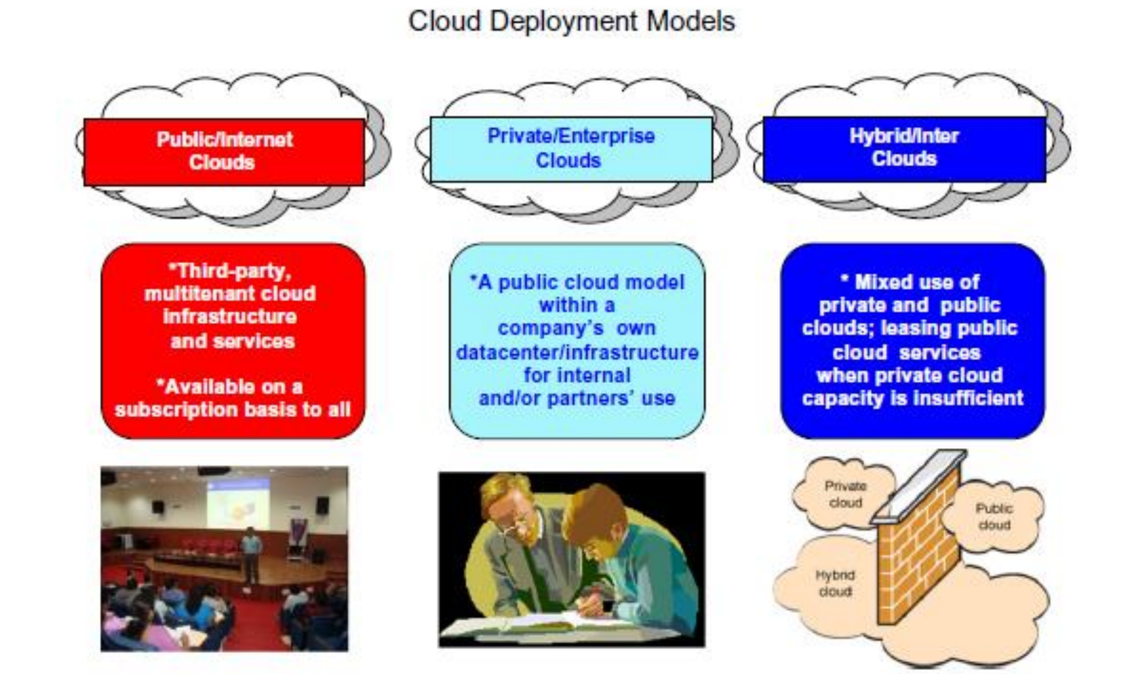

FIGURE 1.4

Major deployment models for cloud computing.

10,000 systems) started a cloud computing initiative aimed at providing government agencies with a more efficient use of their computing facilities. The use of cloud-based in-house solutions is also driven by the need to keep confidential information within an organization's premises. Institutions such as governments and banks that have high security, privacy, and regulatory concerns prefer to build and use their own private or enterprise clouds.

Whenever private cloud resources are unable to meet users' quality-of-service requirements, hybrid computing systems, partially composed of public cloud resources and privately owned infrastructures, are created to serve the organization's needs. These are often referred as *hybrid clouds*, which are becoming a common way for many stakeholders to start exploring the possibilities offered by cloud computing.

Q. 4a) Describe the characteristics of virtualized environment with neat required diagrams.

3.2 Characteristics of Virtualized Environments

Virtualization is a broad concept that refers to the creation of a virtual version of something, whether hardware, a software environment, storage, or a network. In a virtualized environment there are three major components: guest, host, and virtualization layer. The guest represents the system component that interacts with the virtualization layer rather than with the host, as would normally happen. The host represents the original environment where the guest is supposed to be managed. The virtualization layer is responsible for recreating the same or a different environment where the guest will operate (see Figure 3.1).

The characteristics of virtualized solutions are:

a. Increased security
b. Managed execution
c. Portability

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.

Managed Execution.

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 are the most relevant features (see Figure 3.2).

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

Aggregation: Not only is it possible to share physical resource among several guests, butvirtualization 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.

Emulation: Guest programs are executed within an environment that is controlled by thevirtualization 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.

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.

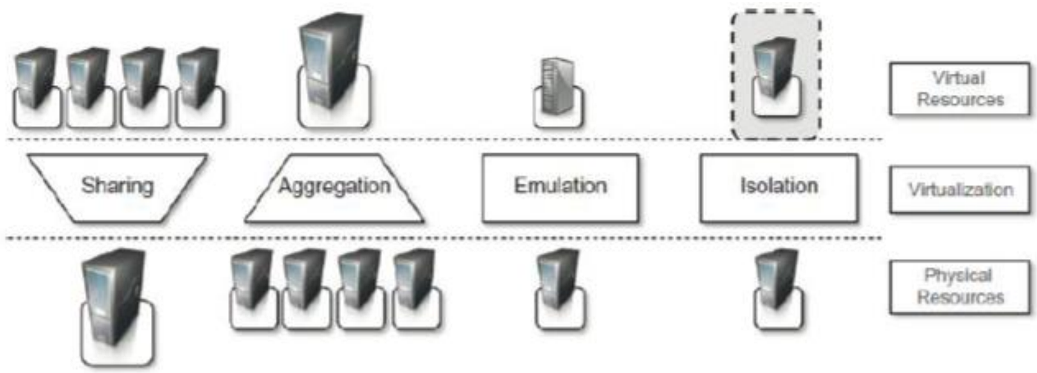

FIGURE 3.2

Functions enabled by managed execution.

Q. 4 b) List and explain pros and cons of virtualization.

3.5 Pros and Cons of Virtualization

Virtualization has now become extremely popular and widely used, especially in cloud computing. Today, the capillary diffusion of the Internet connection and the advancements in computing technology have made virtualization an interesting opportunity to deliver on-demand IT infrastructure and services.

3.5.1 Advantages of Virtualization

- 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.
- Portability is another advantage of virtualization, especially for execution virtualization techni-ques. Virtual machine instances are normally represented by one or more files that can be easily transported with respect to physical systems.
- 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.
- 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 inter-fering with each other.

3.5.2 The other side of the coin: disadvantages

- **Performance Degradation**

  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 (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

- **Inefficiency and Degraded User Experience**

  Virtualization can sometime 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 thecase 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 libraries are used.

- **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.

Q. 5 a) Analyze the open challenges in cloud computing.

4.4 Open Challenges

Cloud computing presents many challenges for industry and academia. There is a significant amount of work in academia focused on defining the challenges brought by this phenomenon.

- **Cloud interoperability and standards**
- **Scalability and fault tolerance**
- **Security, trust, and privacy**
- **Organizational aspects**

4.4.2 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.
The presence of standards that are actually implemented and adopted in the cloud computing community could give room for interoperability and then lessen the risks resulting from vendor lock-in.
The first steps toward a standardization process have been made, and a few organizations, such as the **Cloud Computing Interoperability Forum (CCIF)**, the Open Cloud Consortium, and the **DMTF** Cloud Standards Incubator, are leading the path.
Another interesting initiative is the **Open Cloud Manifesto**, which embodies the point of view of various stakeholders on the benefits of open standards in the field.
The **Open Virtualization Format (OVF)** is an attempt to provide a common format for storing the information and metadata describing a virtual machine image. Even though the OVF provides a full specification for packaging and distributing virtual machine images in completely platform-independent fashion, it is supported by few vendors that use it to import static virtual machine images.

4.4.3 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. 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 areeasy to manage and at the same time provide competitive performance.

4.4.4 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 sensi-tive information. The massive use of virtualization technologies exposes the existing system to new threats, which previously were not considered applicable.
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.
The lack of control over data and processes also poses severe problems for the trust we give to the cloud service provider and the level of privacy we want to have for our data.

4.4.5 Organizational aspects

More precisely, storage, compute power, network infrastructure, and applications are delivered as metered services over the Internet. This introduces a billing model that is new within typical enterprise IT departments, which requires a certain level of cultural and organizational process maturity.
In particular, the following questions have to beconsidered:

- What is the new role of the IT department in an enterprise that completely or significantly relies on the cloud?
- How will the compliance department perform its activity when there is a considerable lack of control over application workflows?
- What are the implications (political, legal, etc.) for organizations that lose control over some aspects of their services?
- What will be the perception of the end users of such services?

From an organizational point of view, the lack of control over the management of data and processes poses not only security threats but also new problems that previously did not exist.

Q. 5 b) What are the benefits of community cloud?

4.2.4 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 National Institute of Standards and Technologies (NIST) [43] characterizes community clouds as follows:

The benefits of these community clouds are the following:

- **Openness.** By removing the dependency on cloud vendors, community clouds are open systemsin which fair competition between different solutions can happen.
- **Community.** Being based on a collective that provides resources and services, the infrastructureturns out to be more scalable because the system can grow simply by expanding its user base.
- **Graceful failures.** Since there is no single provider or vendor in control of the infrastructure,there is no single point of failure.
- **Convenience and control.** Within a community cloud there is no conflict between convenience
- and control because the cloud is shared and owned by the community, which makes all the decisions through a collective democratic process.
- **Environmental sustainability.** The community cloud is supposed to have a smaller carbonfootprint because it harnesses underutilized resources. Moreover, these clouds tend to be more organic by growing and shrinking in a symbiotic relationship to support the demand of the community, which in turn sustains it.

Q. 6a) The cloud service provider (CSP) wants to provide IaaS. Discuss the IaaS based solution for cloud computing which CSP will provide with example.

4.1.2 Infrastructure- and hardware-as-a-service

Infrastructure- and Hardware-as-a-Service (IaaS/HaaS) solutions are the most popular and developed market segment of cloud computing. They deliver customizable infrastructure on demand. The available options within the IaaS offering umbrella range from single servers to entire infrastructures, including network devices, load balancers, and database and Web servers. The main technology used to deliver and implement these solutions is hardware virtualization:
one or more virtual machines opportunely configured and interconnected define the distributed system on top of which applications are installed and deployed. Virtual machines also constitute the atomic components that are deployed and priced according to the specific features of the virtual hardware: memory, number of processors, and disk storage.
From the perspective of the customer it reduces the administration and maintenance cost as well as the capital costs allocated to purchase hardware.

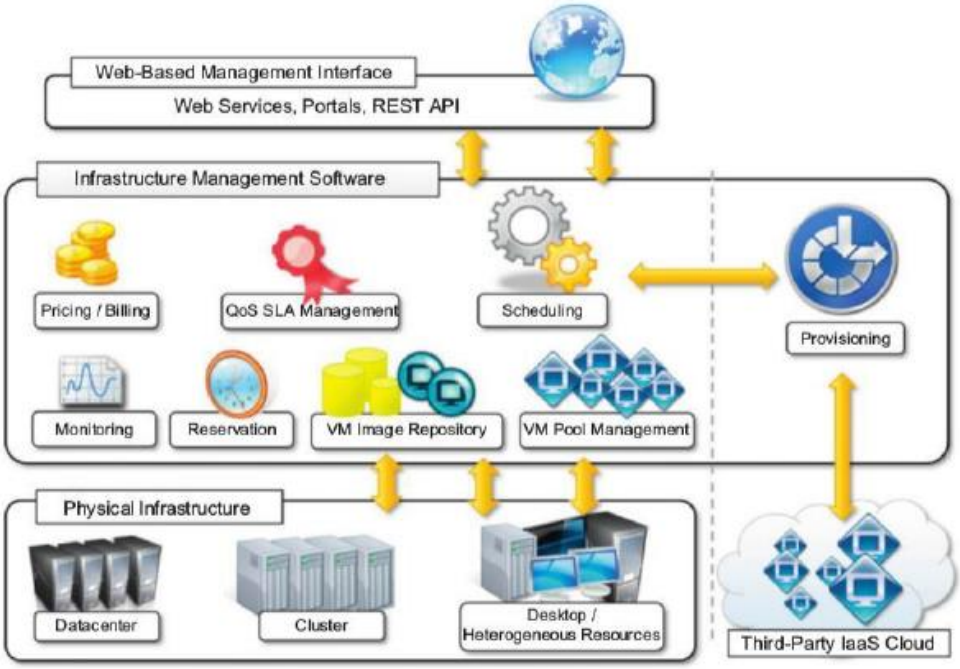

FIGURE 4.2

Infrastructure-as-a-Service reference implementation.

Figure 4.2 provides an overall view of the components forming an Infrastructure-as-a-Servicesolution. It is possible to distinguish three principal layers: the physical infrastructure, the software management infrastructure, and the user interface.

At the top layer the user interface provides access to the services exposed by the software management infrastructure. Such an interface is based on Web 2.0 technologies: Web services, RESTful APIs, and mash-ups. These technologies allow either applications or final users to access the services exposed by the underlying infrastructure. Web 2.0 applications allow developing full-featured management consoles completely hosted in a browser or a Web page. Web services and RESTful APIs allow programs to interact with the service without human intervention, thus providing complete integration within a software system.

The core features of an IaaS solution are implemented in the infrastructure management software layer. In particular, management of the virtual machines is the most important function performed by this layer. A central role is played by the scheduler, which is in charge of allocating the execution of virtual machine instances. The scheduler interacts with the other components that perform a variety of tasks:

- The pricing and billing component takes care of the cost of executing each virtual machine instance and maintains data that will be used to charge the user.
- The monitoring component tracks the execution of each virtual machine instance and maintains data required for reporting and analyzing the performance of the system.

- The reservation component stores the information of all the virtual machine instances thathave been executed or that will be executed in the future.

If support for QoS -based execution is provided, a QoS/SLA management component will maintain a repository of all the SLAs made with the users; together with the monitoring component, this component is used to ensure that a given virtual machine instance is executed with the desired quality of service.

- The VM repository component provides a catalog of virtual machine images that users can use to create virtual instances. Some implementations also allow users to upload their specific virtual machine images.
- A VM pool manager component is responsible for keeping track of all the live instances.

Finally, if the system supports the integration of additional resources belonging to a third-party IaaS provider, a provisioning component interacts with the scheduler to provide a virtual machine instance that is external to the local physical infrastructure directly managed by the pool.

The bottom layer is composed of the physical infrastructure, on top of which the management layer operates. As previously discussed, the infrastructure can be of different types; the specific infrastructure used depends on the specific use of the cloud. A cloud infrastructure developed in house, in a small or medium-sized enterprise or within a university department, will most likely rely on a cluster. At the bottom of the scale it is also possible to consider a heterogeneous environment where different types of resources—PCs, workstations, and clusters—can be aggregated.

Q. 6 b) Compare and contrast between IaaS, PaaS and SaaS?

Parameters	IaaS	PaaS	SaaS
Full Name	Infrastructure as a Service	Platform as a Service	Software as a Service
Who uses it ?	System administrators	Developers	End users
Which service users get ?	**Virtual data center** to store information and create platforms for services and app development, testing and deployment	**Virtual platform and tools** to create, test and deploy apps and services	**Web software and apps** to complete business tasks
Provider controls what ?	Servers, Storage, Networking, Virtualization	Servers, Storage, Networking, Virtualization, OS, Middleware, Runtime	Servers, Storage, Networking, Virtualization, OS, Middleware, Runtime, Applications, Data
User controls what ?	OS, Middleware, Runtime, Applications, Data	Applications, Data	-
Cost	Most expensive	Mid level cost	Cheapest
Flexibility	Very flexible	Flexible but with some limitations	Lowest modifications
Security	Most control over data, but need advanced knowledge in security	Secure but higher level of risk than SaaS	Secure but can be accessed by provider