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

Internal Assessment Test 2 – November 2024

Sub:	**CLOUD COMPUTING**					Sub Code:	**21CS72**	Branch:	ISE	
Date:	20/11/2024	Duration:	90 min's	Max Marks:	50	Sem/Sec:	VII / A, B, C		OBE	
	Answer any FIVE FULL Questions							MARKS	CO	RBT
1.	Explain Cloud Computing architecture in detail with an architectural diagram.							10M	CO3	L2
2.	Explain Infrastructure-and and hardware-as-a-service reference implementation in detail with an implementation diagram.							10M	CO3	L2
3.	Briefly explain different types of Clouds in detail.							10M	CO3	L2
4.	Explain Cloud security risks with Surfaces of attacks in a Cloud Computing environment in detail.							10M	CO4	L2
a	Demonstrate the Virtual machine security with the virtual security service provided.							8M	CO4	L2
5. b	What are the four widely accepted fair information practices?							2M	CO4	L1
6. a	How would you identify and mitigate the Security risks posed by shared images?							8M	CO4	L3
6. b	List the benefits of a pay-as-you-go model.							2M	CO3	L1

Faculty Signature CCI Signature HOD Signature

Internal Assessment Test 2 – November 2024

	21CS72 - CLOUD COMPUTING Scheme & Solutions	MARKS	CO	RBT
1.	Explain Cloud Computing architecture in detail with an architectural diagram. It is possible to organize all the concrete realizations of cloud computing into a layered view covering the entire stack (see Figure 4.1), from hardware appliances to software systems. **Cloud resources are harnessed to offer the "computing horsepower" required for providing services.** This layer is implemented using a **data center** in which hundreds and thousands of nodes are stacked together. Cloud infrastructure can be **heterogeneous in nature** because a variety of resources, such as clusters and even networked PCs, can be used to build it. Moreover, database systems and other storage services can also be part of the infrastructure. **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.	(8+2) 10 M	CO 3	L2

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)).**
- **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.

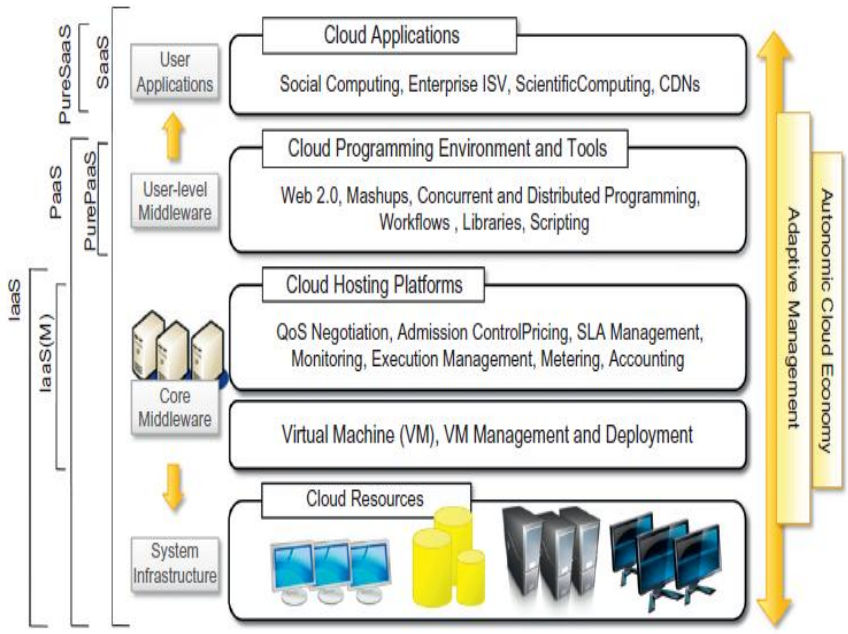

FIGURE 4.1

The cloud computing architecture.

The top layer of the reference model depicted in Figure 4.1 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. Other applications belonging to this layer are those that strongly leverage the Internet for their core functionalities that rely on the cloud to sustain a larger number of users; this is the case of gaming portals and, in general, social networking websites. The reference model described in Figure 4.1 also introduces the concept of everything as a Service (XaaS). This is one of the most important elements of cloud computing: **Cloud services from different providers can be combined to provide a completely integrated solution covering all the computing stack of a system.** IaaS providers can offer the bare metal in terms of virtual machines where PaaS solutions are deployed. When there is no need for a PaaS layer, it is possible to directly customize the virtual infrastructure with the software stack needed to run applications**.** **This is the case of virtual Web farms: a distributed system composed of Web servers, database servers, and load balancers on top of which prepackaged software is installed to run Web applications.** This possibility has made **cloud computing an interesting option for reducing startups' capital investment in IT, allowing them to quickly commercialize their ideas and grow their infrastructure according to their revenues.** **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 The above Table 4.1 summarizes the characteristics of the three major categories used to classify cloud computing solutions.			
2.	Explain Infrastructure-as-a-service reference implementation in detail with an implementation diagram. **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**.	(8+2) 10 M	CO3	L2

	The available options within the IaaS offering umbrella range from single servers to entire infra- structures, 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 sys- tem 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**. IaaS/HaaS solutions bring all the benefits of hardware virtualization: **workload partitioning, application isolation, sandboxing, and hard- ware tuning**. From the perspective of the service provider, IaaS/HaaS allows better exploiting the IT infrastructure and provides a more secure environment where executing third party applications. From the perspective of the customer, it reduces the **administration and maintenance cost** as well as the capital costs allocated to purchase hardware. At the same time, users can take advantage of the full customization offered by virtualization to deploy their infrastructure in the cloud; in most cases virtual machines come with only the selected operating system installed and the system can be configured with all the required packages and applications. Other solutions provide prepackaged system images that already system images that already contain the software stack required for the most common uses**: Web servers, database servers, or LAMP1 stacks**. Besides the basic virtual machine management capabilities, additional services can be provided, generally including the following: **SLA resource-based allocation, workload management, support for infrastructure design through advanced Web inter- faces, and the ability to integrate third-party IaaS solutions**			

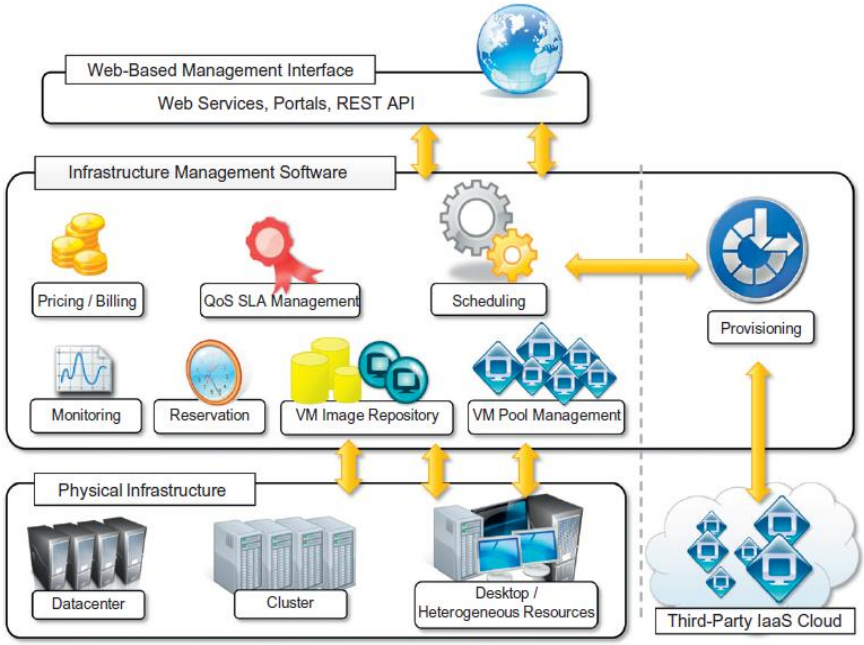

FIGURE 4.2

Infrastructure-as-a-Service reference implementation.

Figure 4.2 provides an overall view of the components forming an Infrastructure-as-a-Service solution.

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 generally based on Web 2.0 technologies: Web services, RESTful APIs, and mashups.

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 soft- ware 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 that have 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.
- **AVM pool manager component** is responsible for keeping track of all the live instances.
- If the system supports the integration of additional resources belonging to a third-party IaaS provider, https://nitw.irins.org/profile/154055#personal_information_panelinteracts 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 service provider will most likely use a **massive datacenter** containing hundreds or thousands of nodes.
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 environ- ment where different types of resources—PCs, workstations, and clusters—can be aggregated.
This case mostly represents an evolution of desktop grids where any available computing resource (such as PCs and workstations that are idle outside of working hours) is harnessed to provide a huge compute power.
From an architectural point of view, the physical layer also includes the virtual resources that are rented from external IaaS providers.
In the case of complete IaaS solutions, all three levels are offered as service.
This is generally the case with public clouds vendors such as **Amazon, GoGrid, Joyent, Rightscale, Terremark, Rackspace, ElasticHosts, and Flexiscale,** which own large datacenters and give access to their computing infrastructures using an IaaS approach.
Other solutions instead cover only the user interface and the infrastructure software management layers.

	They need to provide credentials to access third-party IaaS providers or to own a private infrastructure in which the management software is installed. This is the case with **Enomaly, Elastra, Eucalyptus, OpenNebula, and specific IaaS (M) solutions from VMware, IBM, and Microsoft.** The proposed architecture only represents a reference model for IaaS implementations. It has been used to provide general insight into the most common features of this approach for providing cloud computing services and the operations commonly implemented at this level. Different solutions can feature additional services or even not provide support for some of the features discussed here. Finally, the reference architecture applies to IaaS implementations that provide computing resources, especially for the scheduling component. If storage is the main service provided, it is still possible to distinguish these three layers. The role of infrastructure management software is not to keep track and manage the execution of virtual machines but to provide access to large infrastructures and implement storage virtualization solutions on top of the physical layer.			
3.	Briefly explain different types of Clouds in detail. Clouds constitute the primary outcome of cloud computing. **They are a type of parallel and distributed system harnessing physical and virtual computers presented as a unified computing resource.** **Clouds build the infrastructure on top of which services are implemented and delivered to customers. Such infrastructures can be of different types and provide useful information about the nature and the services offered by the cloud.** A more useful classification is given according to the administrative domain of a cloud: **It identifies the boundaries within which cloud computing services are implemented, provides hints on the underlying infrastructure adopted to support such services, and qualifies them.** It is then possible to differentiate four different types of cloud: • *Public clouds.* The cloud is open to the wider public. • *Private clouds.* The cloud is implemented within the private premises of an institution and generally made accessible to the members of the institution or a subset of them. • *Hybrid or heterogeneous clouds.* The cloud is a combination of the two previous solutions and most likely identifies a private cloud that has been augmented with resources or services hosted in a public cloud. • *Community clouds.* The cloud is characterized by a multi-administrative domain involving different deployment models (public, private, and hybrid), and it is specifically designed to address the needs of a specific industry. **<u>Public clouds:</u>** 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 centres connected together	(2+2+2+2+2 (diagram) 10 M	CO3	L2

Any customer can easily sign in with the cloud provider, enter her **credential and billing details**, and use the services offered public clouds were the first class of cloud that were implemented and offered.
They offer solutions for **minimizing IT infrastructure costs and serve as a viable option for handling peak loads on the local infrastructure.**
They have become an interesting option for small enterprises, which are able to start their businesses without large up-front investments by completely relying on public infrastructure for their IT needs.
What made attractive public clouds compared to the reshaping of the private premises and the purchase of hardware and software was the ability to grow or shrink according to the needs of the related business.
By renting the infrastructure or subscribing to application services, customers were able to dynamically upsize or downsize their IT according to the demands of their business.
Currently, public clouds are used both to completely replace the IT infrastructure of enterprises and to extend it when it is required.
A fundamental characteristic of public clouds is **multitenancy.**
A public cloud is meant to serve a **multitude of users, not a single customer.** Any customer requires a virtual computing environment that is separated, and most likely isolated, from other users.
This is a fundamental requirement to provide effective monitoring of user activities and guarantee the desired performance and the other QoS attributes negotiated with users. **QoS management** is a very important aspect of public clouds.
Hence, a significant portion of the software infrastructure is devoted to monitoring the cloud resources, to bill them according to the contract made with the user, and to keep a complete history of cloud usage for each customer.

These features are fundamental to public clouds because they help providers offer services to users with full accountability.
A public cloud can offer any kind of service:

- Infrastructure
- Platform
- Applications

For example,
Amazon EC2 is a public cloud that provides infrastructure as a service
Google AppEngine is a public cloud that provides an application development platform as a service.
SalesForce.com is a public cloud that provides software as a service.
What makes public clouds peculiar is the way they are consumed: They are available to everyone and are generally architected to support a large quantity of users.
From an architectural point of view there is no restriction concerning the type of distributed system implemented to support public clouds.
Most likely, one or more data centers constitute the physical infrastructure on top of which the services are implemented and delivered. Public clouds can be composed of geographically dispersed data centers to share the load of users and better serve them according to their locations.

For example, **Amazon Web Services has data centers installed in the United States, Europe, Singapore, and Australia; they allow their customers to choose between three different regions: us-west-1, us-east-1, or eu-west-1.**
Such regions are priced differently and are further divided into availability zones, which map to specific data centres. According to the specific class of services delivered by the cloud, a different software stack is installed to manage the infrastructure: virtual machine managers, distributed middleware, or distributed applications.

4.3.2 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.
Private clouds **are virtual distributed systems that rely on a private infrastructure and provide internal users with dynamic provisioning of computing resources.**
Instead of a pay-as-you-go model as in public clouds, there could be other schemes in place, taking into account the usage of the cloud and proportionally billing the different departments or sections of an enterprise.
Private clouds have the advantage of keeping the core business operations in-house by relying on the existing IT infrastructure and reducing the burden of maintaining it once the cloud has been set up.

In this scenario, **security concerns are less critical, since sensitive information does not flow out of the private infrastructure.**
Moreover, existing IT resources can be better utilized because **the private cloud can provide services to a different range of users.**
Another interesting opportunity that comes with private clouds is the **possibility of testing applications and systems at a comparatively lower price rather than public clouds before deploying them on the public virtual infrastructure.**
A Forrester report [34] on the benefits of delivering in-house cloud computing solutions for enter-prises highlighted some of the key advantages of using a private cloud computing infrastructure:

- *Customer information protection.* Despite assurances by the public cloud leaders about security, few provide satisfactory disclosure or have long enough histories with their cloud offerings to provide warranties about the specific level of security put in place on their systems. In-house security is easier to maintain and rely on.
- *Infrastructure ensuring SLAs.* Quality of service implies specific operations such as appropriate clustering and failover, data replication, system monitoring and maintenance, and disaster recovery, and other uptime services can be commensurate to the application needs. Although public cloud vendors provide some of these features, not all of them are available as needed.
- *Compliance with standard procedures and operations.* If organizations are subject to third-party compliance standards, specific procedures have to be put in place when deploying and executing applications. This could be not possible in the case of the virtual public infrastructure.

From an architectural point of view, private clouds can be implemented on more heterogeneous hardware: They generally rely on the existing IT infrastructure already deployed on the private pre- mises.
This could be a **datacenter, a cluster, an enterprise desktop grid, or a combination of them**.

	The physical layer is complemented with infrastructure management software (i.e., IaaS (M); see Section 4.2.2) or a PaaS solution, according to the service delivered to the users of the cloud. Different options can be adopted to implement private clouds. Figure 4.4 provides a comprehensive view of the solutions together with some reference to the most popular software used to deploy private clouds. At the bottom layer of the software stack, virtual machine technologies such as Xen, KVM, and VMware serve as the foundations of the cloud. **Virtual machine management technologies such as VMware vCloud, Eucalyptus, and OpenNebulacan be used to control the virtual infrastructure and provide an IaaS solution.** **VMware vCloud** is a proprietary solution, but Eucalyptus provides full compatibility with Amazon Web Services interfaces and supports different virtual machine technologies such as Xen, KVM, and VMware. Like Eucalyptus, **OpenNebula** is an open-source solution for virtual infrastructure man- agement that supports KVM, Xen, and VMware, which has been designed to easily integrate third-party IaaS providers. Its modular architecture allows extending the software with additional features such as the capability of reserving virtual machine instances by using Haizea [39] as scheduler. Solutions that rely on the previous virtual machine managers and provide added value are OpenPEX and InterGrid. **OpenPEX** is Web-based system that allows the reservation of virtual machine instances and is designed to support different back ends. **InterGrid** provides added value on top of OpenNebula and Amazon EC2 by allowing the reservation of virtual machine instances and managing multi- administrative domain clouds. PaaS solutions can provide an additional layer and deliver a high- level service for private clouds. Among the options available for private deployment of clouds we can consider DataSynapse, Zimory Pools, Elastra, and Aneka. DataSynapse is a global provider of application virtualization software. By relying on the VMware virtualization technology**, DataSynapse** provides a flexible environment for building private clouds on top of datacenters.			

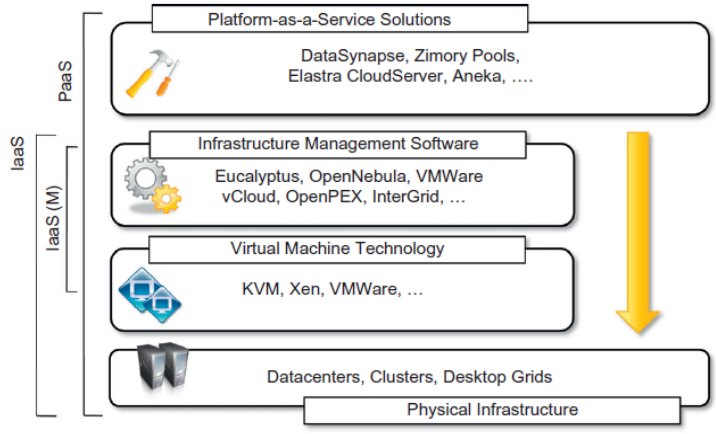

FIGURE 4.4

Private clouds hardware and software stack.

Elastra Cloud Server is a platform for easily configuring and deploying distributed application infrastructures on clouds.

Zimory provides a software infrastructure layer that automates the use of resource pools based on Xen, KVM, and VMware virtualization technologies.

It allows creating an internal cloud composed **of sparse private and public resources and provides facilities for migrating applications within the existing infrastructure**.

Aneka is a software development platform that can be used to deploy a cloud infrastructure on top of heterogeneous hardware: datacenters, clusters, and desktop grids.

It provides a pluggable service-oriented architecture that's mainly devoted to supporting the execution of distributed applications with different programming models: bag of tasks, MapReduce, and others.

Private clouds can provide in-house solutions for cloud computing, but if compared to public clouds they exhibit more limited capability to scale elastically on demand.

4.3.3 Hybrid Cloud

Public clouds are large software and hardware infrastructures that have a capability that is huge enough to serve the needs of multiple users, but they suffer from security threats and administrative pitfalls.

Private clouds are the perfect solution when it is necessary to keep the processing of information within an enterprise's premises or it is necessary to use the existing hardware and software infrastructure. One of the major drawbacks of private deployments is the inability to scale on demand and to efficiently address peak loads.

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.

Figure 4.5 provides a general overview of a hybrid cloud:

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 cloudbursting.
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. What is missing is then an advanced scheduling engine that's able to differentiate these resources and provide smart allocations by taking into account the budget available to extend the existing infrastructure. 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.

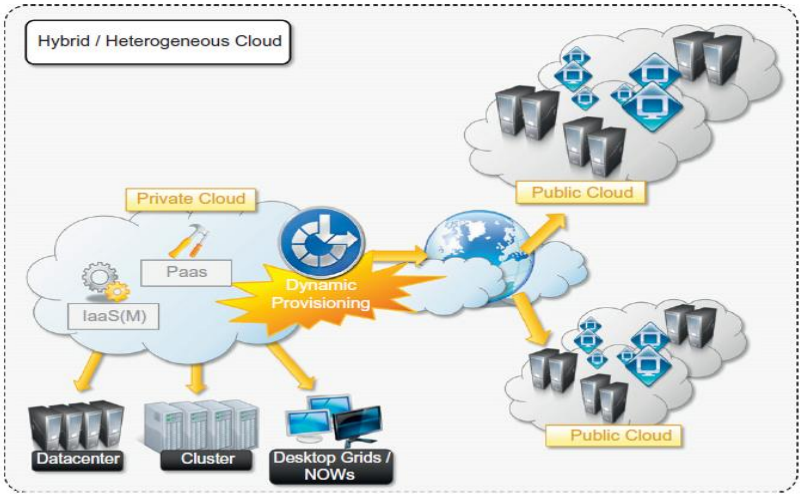

FIGURE 4.5
Hybrid/heterogeneous cloud overview.

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.

Dynamic provisioning is most commonly implemented in PaaS solutions that support hybrid clouds. As previously discussed, one of the fundamental components of PaaS middleware is the mapping of distributed applications onto the cloud infrastructure.

In this scenario, the role of dynamic provisioning becomes fundamental to ensuring the execution of applications under the QoS agreed on with the user. For example, **Aneka provides a provisioning service that leverages different IaaS providers for scaling the existing cloud infrastructure.**

The provisioning service cooperates with the scheduler, which is in charge of guaranteeing a specific QoS for applications. In particular, each user application has a budget attached, and the scheduler

uses that budget to optimize the execution of the application by renting virtual nodes if needed. Other PaaS implementations support the deployment of hybrid clouds and provide dynamic provisioning capabilities. Among those discussed for the implementation and management of private clouds we can cite Elastra CloudServer and Zimory Pools.

4.3.4 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) characterizes community clouds as follows:

The infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on premise or off premise.

Figure 4.6 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.

Candidate sectors for community clouds are as follows:

- Media industry
- Healthcare industry
- Energy and other core industries
- Public sector.
- Scientificresearch.

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.

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.

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.

Scientific research: Science clouds are an interesting example of community clouds. In this case, the common interest driving different organizations sharing a large distributed infrastructure is scientific computing.

The term **community cloud** can also identify a more specific type of cloud that arises from concern over the controls of vendors in cloud computing and that aspire to combine the principles of **digital ecosystems.**

The benefits of these community clouds are the following:

Openness: By removing the dependency on cloud vendors, community clouds are open systems in which fair competition between different solutions can happen.

Community: Being based on a collective that provides resources and services, the infrastructure turns 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 carbon footprint because it harnesses underutilized resources.			
4.	Explain Cloud security risks with Surfaces of attacks in a Cloud Computing environment in detail. Traditional threats are those experienced for some time by any system connected to the Internet, but with some cloud-specific twists. The impact of traditional threats is amplified due to the vast amount of cloud resources and the large user population that can be affected. The traditional threats begin at the user site. The user must protect the infrastructure used to connect to the cloud and to interact with the application running on the cloud. This task is more difficult because some components of this infrastructure are outside the firewall protecting the user. The next threat is related to the authentication and authorization process. The procedures in place for one individual do not extend to an enterprise. In this case the cloud access of the members of an organization must be nuanced; individuals should be assigned distinct levels of privilege based on their roles in the organization. It is also nontrivial to merge or adapt the internal policies and security metrics of an organization with the ones of the cloud. Moving from the user to the cloud, we see that the traditional types of attack have already affected cloud service providers. The favorite means of attack are distributed denial-of-service (DDoS) attacks, which prevent legitimate users accessing cloud services; phishing;2 SQL injection;3 or cross-site scripting.4 Availability of cloud services is another major concern. System failures, power outages, and other catastrophic events could shut down cloud services for extended periods of time. When such an event occurs, data lock-in, discussed in Section 3.5, could prevent a large organization whose business model depends on that data from functioning properly. Third-party control generates a spectrum of concerns caused by the lack of transparency and limited user control. For example, a cloud provider may subcontract some resources from an third party whose level of trust is questionable. There are examples when subcontractors failed to maintain the customer data. There are also examples when the third party was not a subcontractor but a hardware supplier and the loss of data was caused by poor-quality storage devices [83]. Insecure APIs may not protect users during a range of activities, starting with authentication and access control to monitoring and control of the application during runtime. The cloud service providers Data loss or leakage are two risks with devastating consequences for an individual or an organization using cloud services. Maintaining copies of the data outside the cloud is often unfeasible due to the sheer volume of data. If the only copy of the data is stored on the cloud, sensitive data is permanently lost when cloud data replication fails and is followed by a storage media	(8 +2) 10 M	CO 4	L2

	failure. Because some of the data often includes proprietary or sensitive data, access to such information by third parties could have severe consequences. Account or service hijacking is a significant threat, and cloud users must be aware of and guard against all methods of stealing credentials. Finally, unknown risk profile refers to exposure to the ignorance or underestimation of the risks of cloud computing. The three actors involved in the model considered are the user, the service, and the cloud infrastructure, and there are six types of attacks possible (see Figure 9.1). The user can be attacked from two directions: FIGURE Surfaces of attacks in a cloud computing environment. 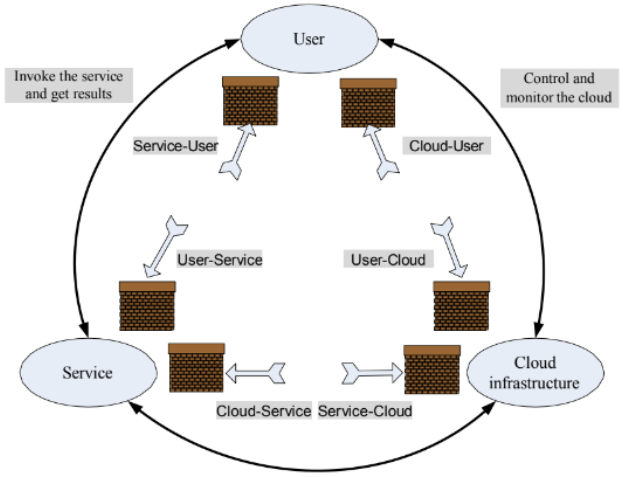 **FIGURE 9.1:** Surfaces of attacks in a cloud computing environment.			
5. a	Demonstrate the Virtual machine security with the virtual security service provided. VMM controls access to the hardware. 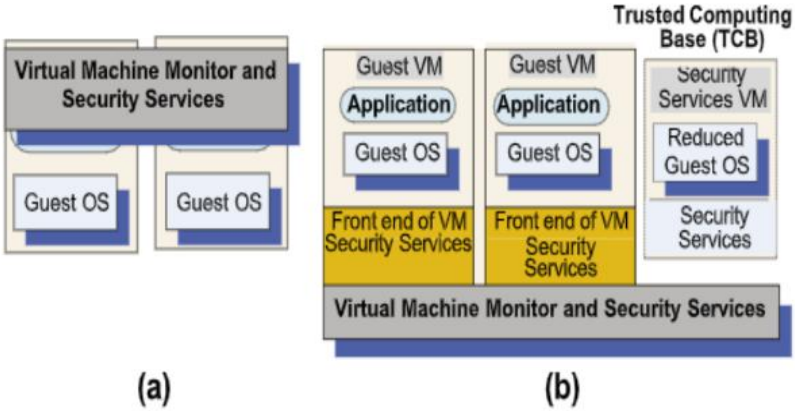 Virtual security services are typically provided by the VMM, as shown in Figure 9.2(a).	8(6+2)M	CO4	L2

	Another alternative is to have a dedicated security services VM, as shown in Figure 9.2(b). A FIGURE 9.2: Virtualsecurity services provided by the VMM. (b) A dedicated security VM. computing base (TCB) is a necessary condition for security in a virtual machine environment; if the TCB is compromised, the security of the entire system is affected. VM technology provides a stricter isolation of virtual machines from one another than the isolation of processes in a traditional operating system. Indeed, a VMM controls the execution of privileged operations and can thus enforce memory isolation as well as disk and network access. The VMMs are considerably less complex and better structured than traditional operating systems; thus, they are in a better position to respond to security attacks. A major challenge is that a VMM sees only raw data regarding the state of a guest operating system, whereas security services typically operate at a higher logical level, e.g., at the level of a file rather than a disk block. A guest OS runs on simulated hardware, and the VMM has access to the state of all virtual machines operating on the same hardware. The state of a guest virtual machine can be saved, restored, cloned, and encrypted by the VMM. Not only can replication ensure reliability, it can also support security, whereas cloning could be used to recognize a malicious application by testing it on a cloned system and observing whether it behaves normally. The security group involved with the NIST project has identified the following VMM- and VM-based threats: • VMM-based threats: 1. Starvation of resources and denial of service for some VMs. Probable causes: (a) badly configured resource limits for some VMs; (b) a rogue VM with the capability to bypass resource limits set in the VMM. 2. VM side-channel attacks. Malicious attacks on one or more VMs by a rogue VM under the same VMM. Probable causes: (a) lack of proper isolation of inter-VM traffic due to misconfiguration of the virtual network residing in the VMM; (b) limitation of packet inspection devices to handle high-speed traffic, e.g., video traffic; (c) presence of VM instances built from insecure VM images, e.g., a VM image having a guest OS without the latest patches. 3. Buffer overflow attacks. • VM-based threats: 1. Deployment of rogue or insecure VM. Unauthorized users may create insecure instances from images or may perform unauthorized administrative actions on existing VMs. Probable cause: improper configuration of access controls on VM administrative tasks such as instance creation, launching, suspension, reactivation, and so on. 2. Presence of insecure and tampered VM images in the VM image repository. Probable causes: (a) lack of access control to the VM image repository; (b) lack of mechanisms to verify the integrity of the images, e.g., digitally signed image.			
5. b	What are the four widely accepted fair information practices? Notice. Websites would be required to provide consumers clear and conspicuous notice of their information practices, including what information they collect, how they collect it (e.g., directly or through nonobvious means such as cookies), how they use it, how they provide Choice, Access, and Security to consumers, whether they disclose the information collected to other entities, and whether other entities are collecting information through the site.	2 M	CO 4	L1

	2. Choice. Web sites would be required to offer consumers choices as to how their personal identifying information is used beyond the use for which the information was provided (e.g., to consummate a transaction). Such choice would encompass both internal secondary uses (such as marketing back to consumers) and external secondary uses (such as disclosing data to other entities). 3. Access. Web sites would be required to offer consumers reasonable access to the information a Web site has collected about them, including a reasonable opportunity to review information and to correct inaccuracies or delete information. 4. Security. Web sites would be required to take reasonable steps to protect the security of the information they collect from consumers. The Commission recognizes that the implementation of these practices may vary with the nature of the information collected and the uses to which it is put, as well as with technological developments. For this reason, the Commission recommends that any legislation be phrased in general terms and be technologically neutral. Thus, the definitions of fair information practices set forth in the statute should be broad enough to provide flexibility to the implementing agency in promulgating its rules or regulations."			
6. a	How would you identify and mitigate the Security risks posed by shared images? Even when we assume that a cloud service provider is trustworthy, many users either ignore or under- estimate the danger posed by other sources of concern. One of them, especially critical to the IaaS cloud delivery model, is image sharing. For example, a user of AWS has the option to choose between Amazon Machine Images (AMIs), accessible through the Quick Start or the Community AMI menus of the EC2 service. The option of using one of these AMIs is especially tempting for a first-time or less sophisticated user. First, let's review the process to create an AMI. We can start from a running system, from another AMI, or from the image of a VM and copy the contents of the file system to the S3, the so-called bundling. The first of the three steps in bundling is to create an image, the second step is to compress and encrypt the image, and the last step is to split the image into several segments and then upload the segments to the S3. Two procedures for the creation of an image are available: ec2-bundle-image and ec2-bundle-volume. The first is used for images prepared as loopback files when the data is transferred to the image in blocks. To bundle a running system, the creator of the image can use the second procedure when bundling works at the level of the file system and files are copied recursively to the image. To use an image, a user has to specify the resources, provide the credentials for login, provide a firewall configuration, and specify the region. Once instantiated, the user is informed about the public DNS and the virtual machine is made available. A Linux system can be accessed using ssh at port 22, whereas the Remote Desktop at port 3389 is used for Windows. Many of the analyzed images allowed a user to undelete files and recover credentials, private keys, or other types of sensitive information with little effort and using standard tools. The results of this study were shared with	8 M	CO 4	L3

Amazon's Security Team, which acted promptly to reduce the threats posed to AWS users.

Three types of security risks were analyzed: (1) backdoors and leftover credentials, (2) unsolicited connections, and (3) malware. An astounding finding is that about 22% of the scanned Linux AMIs contained credentials allowing an intruder to remotely log into the system. Some 100 passwords, 995 sshkeys, and 90 cases in which both passwords and keys could be retrieved were identified.

To rent a Linux AMI, a user must provide the public part of the ssh key, and this key is stored in the authorized_keysin the home directory. This opens a backdoor for a malicious creator of an AMI who does not remove his own public key from the image and can remotely log into any instance of this AMI. Another backdoor is opened when the ssh server allows password-based authentication and the malicious creator of an AMI does not remove his own password. This backdoor is opened even wider as one can extract the password hashes and then crack the passwords using a tool such as John the Ripper.

Another threat is posed by the omission of the cloud-initscript that should be invoked when the image is booted. This script, provided by Amazon, regenerates the host key an ssh server uses to identify itself; the public part of this key is used to authenticate the server. When this key is shared among several systems, these systems become vulnerable to man-in-the middle11 attacks. When this In a man-in-the-middle an attacker impersonates the agents at both ends of a communication channel and makes them believe that they communicate through a secure channel.

For example, if B sends her public key to A, but C is able to intercept it, such an attack proceeds as follows: C sends a forged message to A claiming to be from B but instead includes C's public key. Then A encrypts his message with C's key, believing that he is using B's key, and sends the encrypted message to B. The intruder, C, intercepts, deciphers the message using her private key, possibly alters the message, and re-encrypts the public key B originally sent to A. When B receives the newly encrypted message, she believes it came from A.

Unsolicited connections pose a serious threat to a system. Outgoing connections allow an outside entity to receive privileged information, e.g., the IP address of an instance and events recorded by a syslog daemon to files in the var/log directory of a Linux system. Such information is available only to users with administrative privileges. The audit detected two Linux instances with modified syslog daemons, which forwarded to an outside agent information about events such as login and incoming requests to a Web server. Some of the unsolicited connections are legitimate – for example, connections to a software update site. It is next to impossible to distinguish legitimate from malicious connections.

Malware, including viruses, worms, spyware, and trojans, were identified using ClamAV, The creator of a shared AMI assumes some privacy risks; his private keys, IP addresses, browser history, shell history, and deleted files can be recovered from the published images. A malicious agent can recover the AWS API keys that are not password protected. Then the malicious agent can start AMIs and run cloud applications at no cost to herself, since the computing charges are passed on to the owner of the API key. The search can target files with names such as $pk - [0 - 9\ A - Z\ ] *.$ pem or $cert - [0 - 9\ A - Z\ ]$

	∗ . pem used to store API keys. Another avenue for a malicious agent is to recover sshkeys stored in files named id_dsa and id_rsa. Though sshkeys can be protected by a passphrase, the audit determined that the majority of sshkeys (54 out of 56) were not password protected. Recovery of IP addresses of other systems owned by the same user requires access to the lastlog or the lastb databases. The audit found 187 AMIs with a total of more than 66,000 entries in their lastb databases Users should be aware that when HTTP is used to transfer information from a user to a Web site, the GET requests are stored in the logs of the Web server. Passwords and credit card numbers communicated via a GET request can be exploited by a malicious agent with access to such logs. When remote credentials such as the DNS management password are available, a malicious agent can redirect traffic from its original destination to her own system. Recovery of deleted files containing sensitive information poses another risk for the provider of an image. When the sectors on the disk containing sensitive information are actually overwritten by another file, recovery of sensitive information is much harder.			
6. b	List the benefits of a pay-as-you-go model. Reducing the capital costs associated to the IT infrastructure Eliminating the depreciation or lifetime costs associated with IT capital assets Replacing software licensing with subscriptions Cutting the maintenance and administrative costs of IT resourc	2 M	C O 3	L1

Faculty Signature CCI Signature HOD Signature