Cloud Computing

AS PER NEW SYLLABUS - GTU - SEM - VII (CE/CSE/ICT) Professional Elective - IV

Sub. Code : 3170717

- Simplified & Conceptual Approach
- Chapterwise Solved GTU Questions Summer 2017 to Winter 2018
- Multiple Choice Questions with Answers

TECHNICAL PUBLICATIONS®
SINCE 1993 An Up-Thrust for Knowledge

I. A. Dhotre

first edition : sept. 2021

Price : ₹ 195/-
ISBN 978-93-332-2419-2

TABLE OF CONTENTS

1 Introduction

Syllabus

Cloud Computing, Layers and Types of Clouds, Cloud Infrastructure Management, Challenges and Applications. Virtualization : Virtualization of Computing, Storage and Resources. Cloud Services : Introduction to Cloud Services IaaS, PaaS and SaaS.

Contents

1.1 Introduction to Cloud Computing

GTU : Summer-17, 18, Winter-17, 18

- Cloud computing refer to a variety of services available over the Internet that deliver compute functionality on the service provider's infrastructure.
- Its environment (infrastructure) may actually be hosted on either a grid or utility computing environment, but that doesn't matter to a service user.
- Cloud computing is a general term used to describe a new class of network based computing that takes place over the Internet, basically a step on from Utility Computing.
- In other words, this is a collection/group of integrated and networked hardware, software and Internet infrastructure (called a platform).
- Cloud computing refers to applications and services that run on a distributed network using virtualized resources and accessed by common Internet protocols and networking standards.
- Fig. 1.1.1 shows cloud symbol. It denotes cloud boundary.

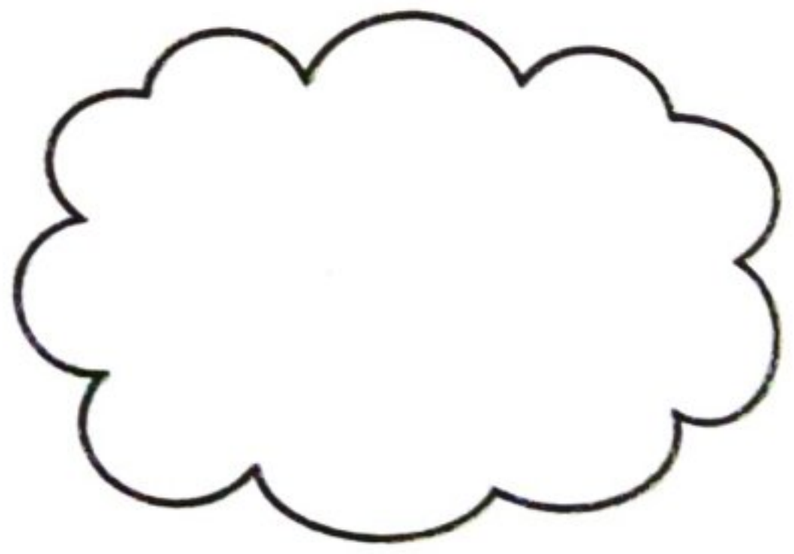

Fig. 1.1.1 : Cloud symbol

- Using the Internet for communication and transport provides hardware, software and networking services to clients.
- These platforms hide the complexity and details of the underlying infrastructure from users and applications by providing very simple graphical interface or API.
- In addition, the platform provides on demand services that are always on anywhere, anytime and anyplace. Pay for use and as needed.
- The hardware and software services are available to the general public, enterprises, corporations and business markets.

IT resources :

- IT resources are of two types : Software based and hardware based.
- Software based resources are virtual server, custom software program and hardware based means physical server and networking devices.
- IT resources include server, virtual server, storage device, networking device, services and software programs.
- An on-premise IT resource can access and interact with a cloud-based IT resource.

- An on-premise IT resource can be moved to a cloud, thereby changing it to a cloud-based IT resource.
- **Cloud provider** : A person, organization, or entity responsible for making a service available to interested parties. When assuming the role of cloud provider, an organization is responsible for making cloud services available to cloud consumers, as per agreed upon Service Level Agreement (SLA) guarantees. Cloud provider have their own IT resources.
- **Cloud consumer** : A person or organization that maintains a business relationship with, and uses service from, Cloud Providers. The cloud consumer uses a cloud service consumer to access a cloud service.
- **Cloud service owner** : The person or organization that legally owns a cloud service is called a cloud service owner. The cloud service owner can be the cloud consumer, or the cloud provider that owns the cloud within which the cloud service resides.
- **Resource administrator** : Cloud resource administrator is the person or organization responsible for administering a cloud-based IT resource. The cloud consumer or cloud provider, or even third-party organization could be a cloud resource administrator

Cloud types :

- Most people separate cloud computing into two distinct sets of models :
  1. **Deployment models** : This refers to the location and management of the cloud's infrastructure.
  2. **Service models** : This consists of the particular types of services that you can access on a cloud computing platform.
- Fig. 1.1.2 shows NIST cloud computing definitions.

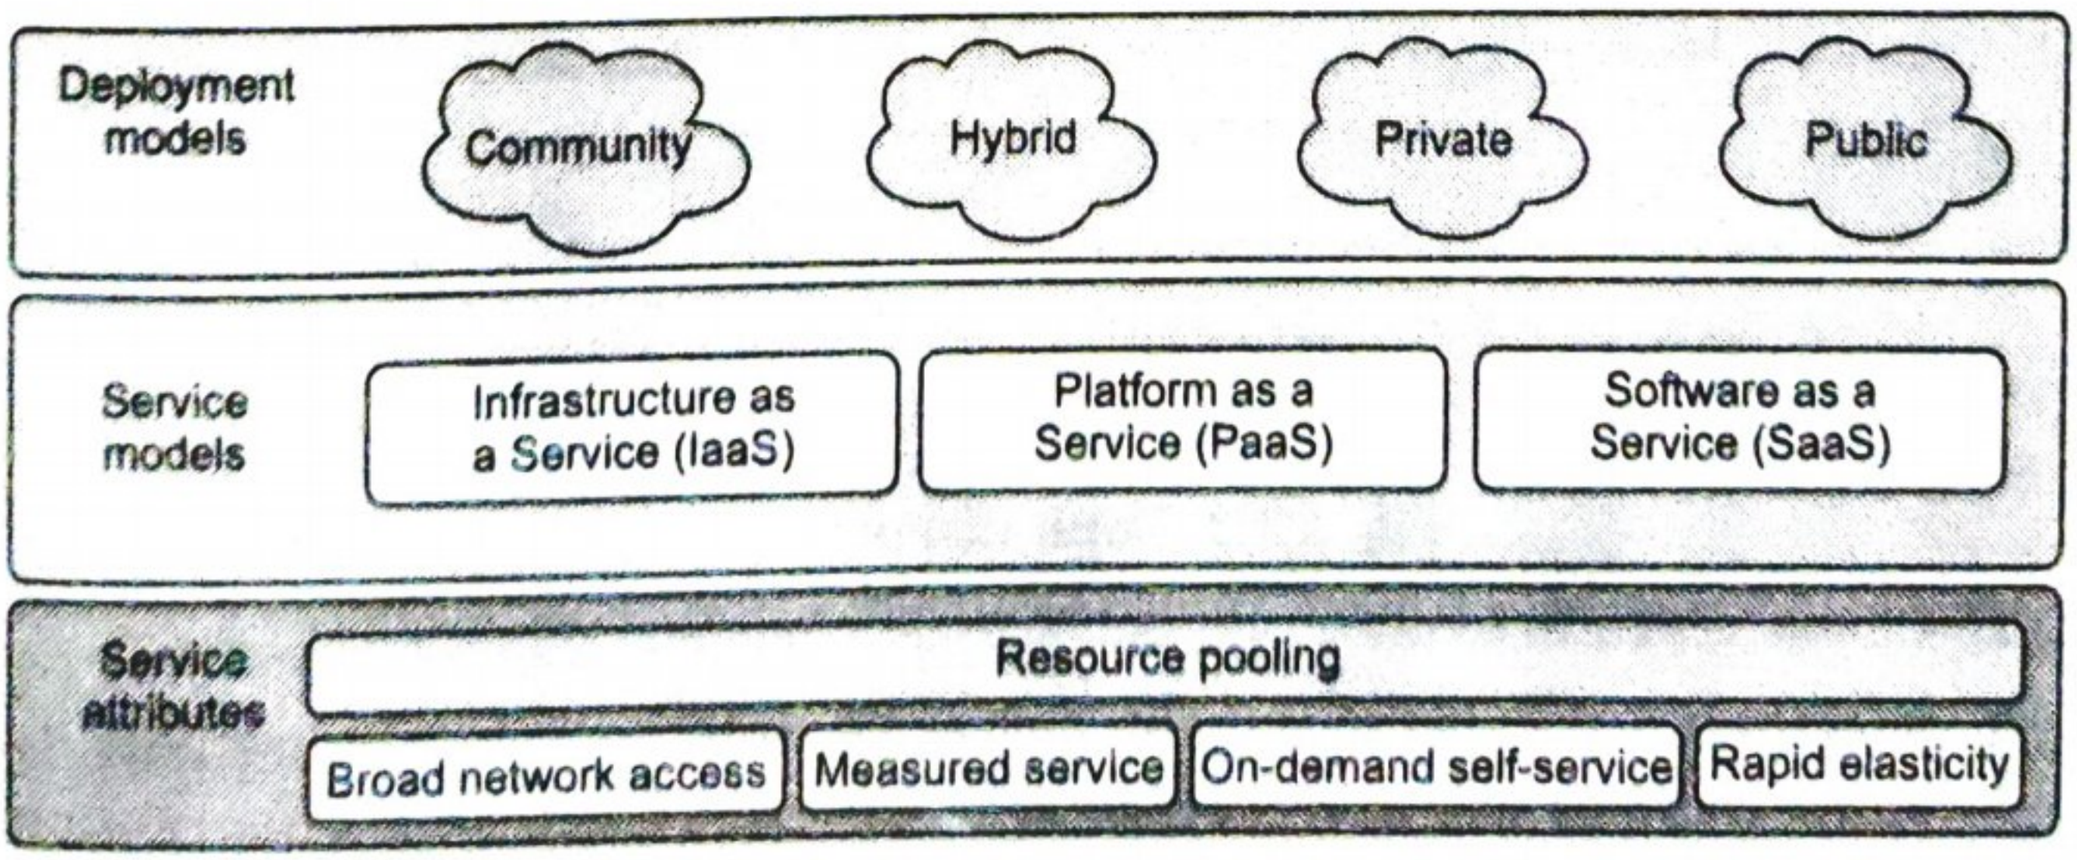

Fig. 1.1.2

- **On-demand self-service :** A client can provision computer resources without the need for interaction with cloud service provider personnel.
- **Broad network access :** Access to resources in the cloud is available over the network using standard methods in a manner that provides platform-independent access to clients of all types. This includes a mixture of heterogeneous operating systems, and thick and thin platforms such as laptops, mobile phones, and PDA.
- **Resource pooling :** A cloud service provider creates resources that are pooled together in a system that supports multi-tenant usage. Physical and virtual systems are dynamically allocated or reallocated as needed.
- **Rapid elasticity :** Resources can be rapidly and elastically provisioned
- **Measured service :** The use of cloud system resources is measured, audited, and reported to the customer based on a metered system.

1.1.1 Cloud Components

- Cloud computing solutions are made up of several elements. Fig. 1.1.3 shows cloud components.

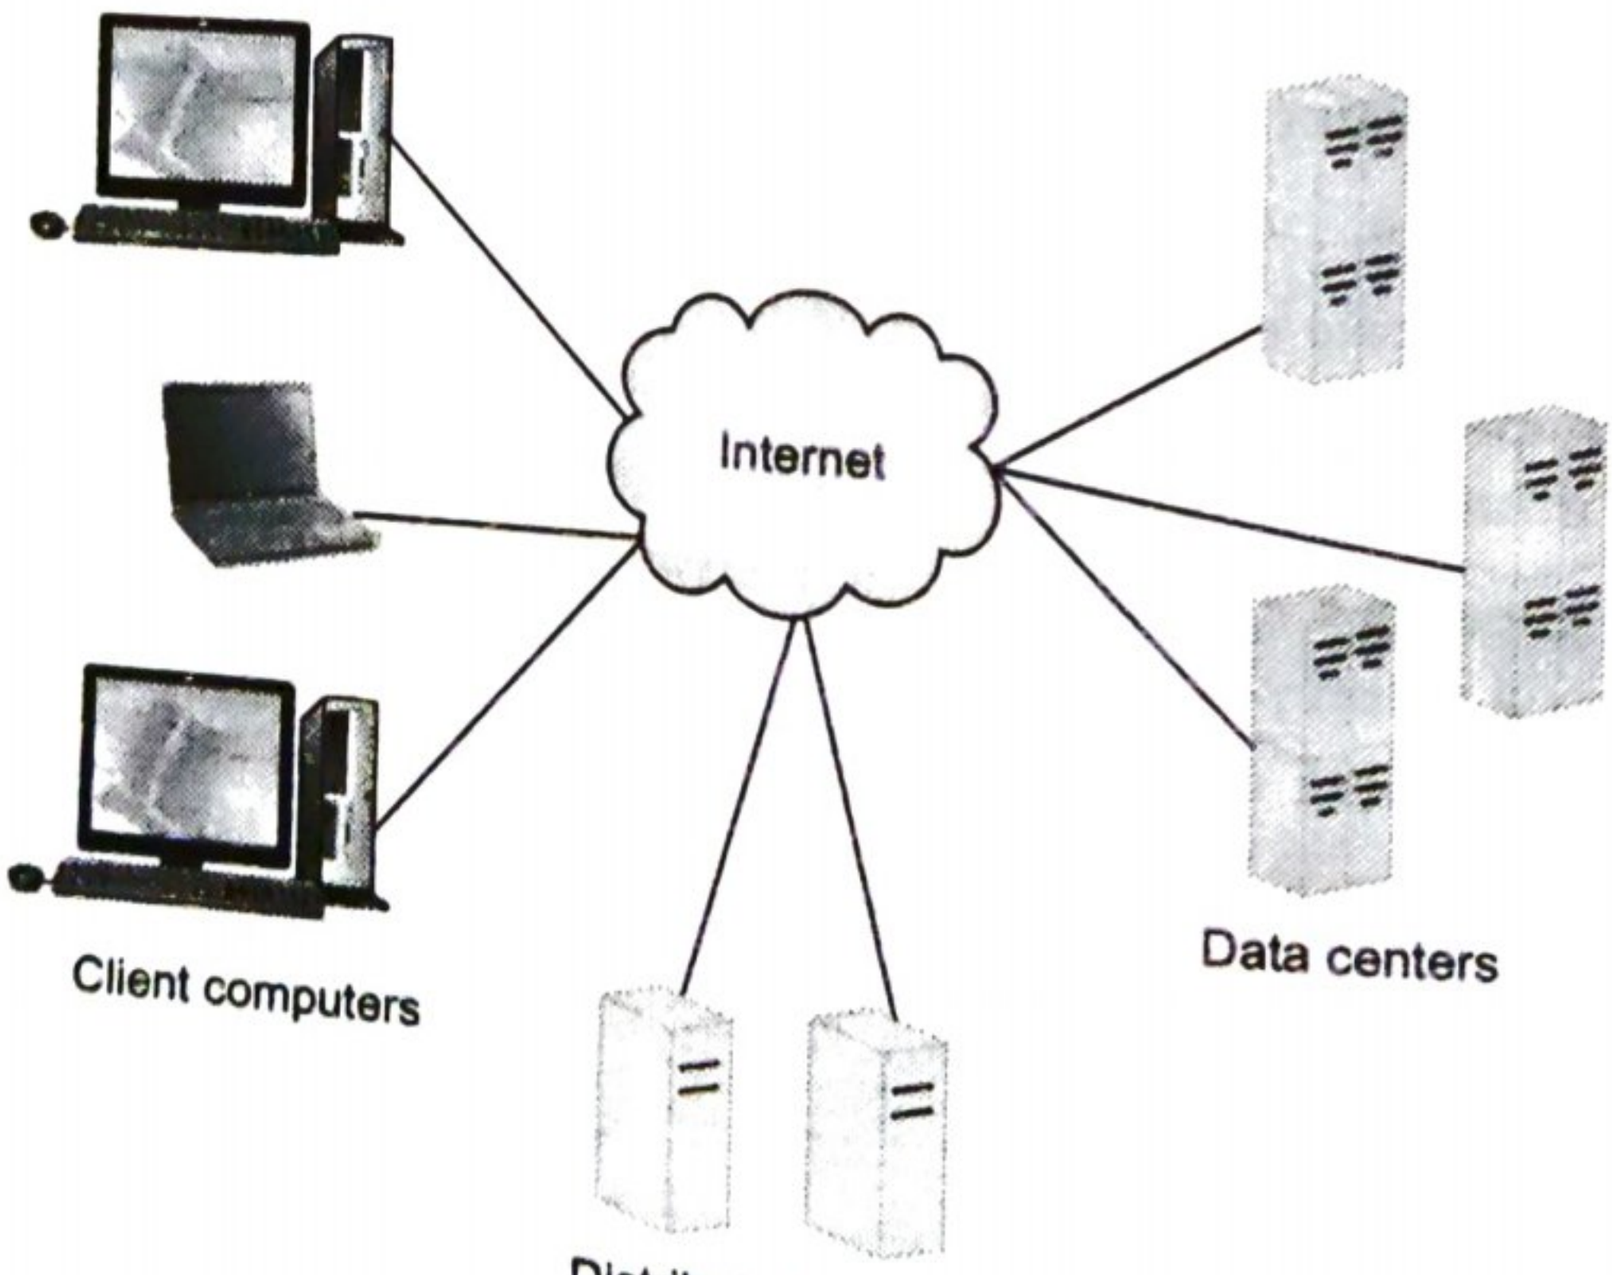

Fig. 1.1.3 Cloud components

1. **Clients :** Mobile, terminals or regular computers.
2. **Benefits :** Lower hardware costs, lower IT costs, security, data security, less power consumption, ease of repair or replacement, less noise.

3. **Data centers** : Collection of servers where the application to subscribe is housed. It could be a large room in the basement of your building or a room full of servers on the other side of the world
4. **Virtualizing servers** : Software can be installed allowing multiple instances of virtual servers to be used and a dozen virtual servers can run on one physical server.
5. **Distributed servers** : Servers don't all have to be housed in the same location. It can be in geographically disparate locations. If something were to happen at one site, causing a failure, the service would still be accessed through another site. If the cloud needs more hardware, they can add them at another site.

1.1.2 Cloud Computing Architecture

- Fig. 1.1.4 shows architectural framework of cloud computing.

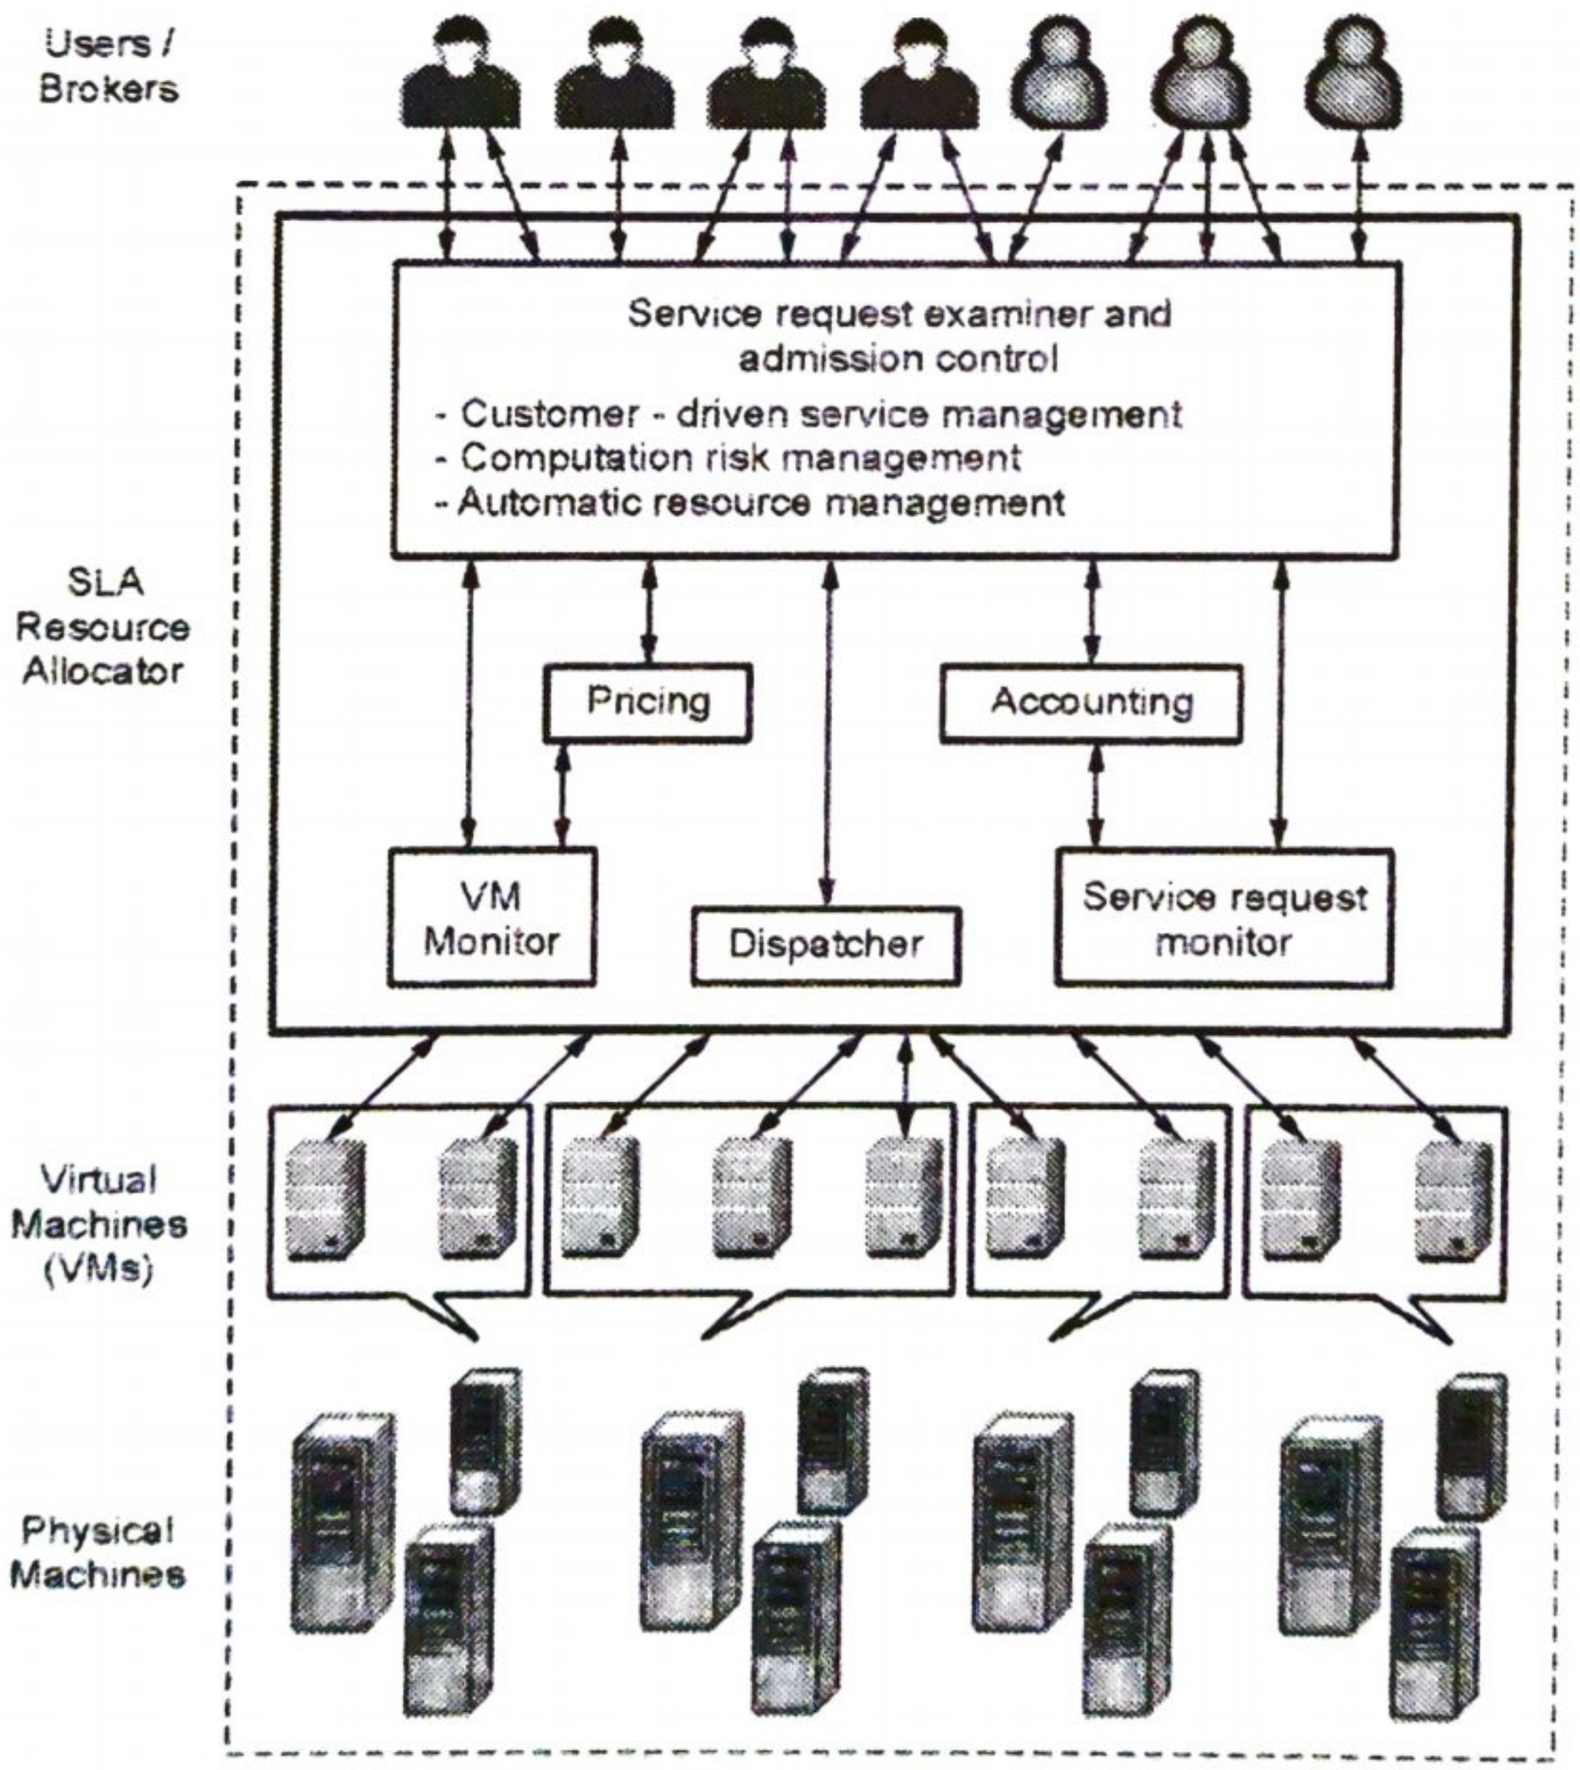

Fig. 1.1.4 Architectural framework

1. **Users/Brokers** : They submit their service requests from anywhere in the world to the cloud.
2. **SLA resource allocator** : It is a kind of interface between users and cloud service provider which enable the SLA-oriented resource management.
3. **Service request examiner and admission control** : It interprets the submitted request for QoS requirements before determining whether to accept or reject the request. Based on resource availability in the cloud and other parameters decide.
4. **Pricing** : It is in charge of billing based on the resource utilization and some factors. Some factors are request time, type etc.
5. **Accounting** : Maintains the actual usage of resources by request so that the final cost can be charged to the users.
6. **VM monitor** : Keeps tracks on the availability of VMs and their resources.
7. **Dispatcher** : The dispatcher mechanism starts the execution of admitted requests on allocated VMs.
8. **Service request monitor** : The request monitor mechanism keeps track on execution of request in order to be in tune with SLA.

Cloud computing service layers :

Parameters	Services	Description
Application Focused	Services	Services - Complete business services such as PayPal, OpenID, OAuth, Google Maps, Alexa
	Application	Application - Cloud based software that eliminates the need for local installation such as Google Apps, Microsoft Online
	Development	Development - Software development platforms used to build custom cloud based applications (PAAS and SAAS) such as SalesForce
Infrastructure Focused	Platform	Platform - Cloud based platforms, typically provided using virtualization, such as Amazon ECC, Sun Grid
	Storing	Storage - Data storage or cloud based NAS such as CTERA, iDisk, CloudNAS
	Hosting	Hosting - Physical data centers such as those run by IBM, HP, NaviSite, etc.

Cloud components :

- Cloud computing solutions are made up of several elements.
  1. **Clients** : Mobile, terminals or regular computers.

2. **Benefits** : Lower hardware costs, lower IT costs, security, data security, less power consumption, ease of repair or replacement, less noise.
3. **Data centers** : Collection of servers where the application to subscribe is housed. It could be a large room in the basement of your building or a room full of servers on the other side of the world
4. **Virtualizing servers** : Software can be installed allowing multiple instances of virtual servers to be used and a dozen virtual servers can run on one physical server.
5. **Distributed servers** : Servers don't all have to be housed in the same location. It can be in geographically disparate locations. If something were to happen at one site, causing a failure, the service would still be accessed through another site. If the cloud needs more hardware, they can add them at another site.

1.1.3 Characteristics of Cloud Computing

1. **On-demand self-service** : A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed without requiring human interaction with each service's provider.
2. **Ubiquitous network access** : Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms.
3. **Location-independent resource pooling** : The provider's computing resources are pooled to serve all consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand.
4. **Rapid elasticity** : Capabilities can be rapidly and elastically provisioned to quickly scale up, and rapidly released to quickly scale down.
5. **Pay per use** : Capabilities are charged using a metered, fee-for-service, or advertising-based billing model to promote optimization of resource use.

1.1.4 Cloud Applications

1. Through cloud cost flexibility, online marketplace gains access to more powerful analytics online. Cloud takes away the need to fund the building of hardware, installing software or paying dedicated software license fees.
2. Greater business scalability enables online video retailer to meet spikes in demand : Cloud enables businesses not just IT operations to add or provision computing resources just at the time they're needed.
3. Greater market adaptability provides online entertainment platform the ability to reach any type of customer device. A third of the executives we surveyed believe

cloud can help them adapt to diverse user groups with a diverse assortment of devices.

4. Masked complexity enables access to services, no matter how intricate the technology they're built on.
5. With context-driven variability, "intelligent assistants" are possible. "Because of its expanded computing power and capacity, cloud can store information about user preferences, which can enable product or service customization," the report states.
6. Ecosystem connectivity enables information exchange across business partners.

1.1.5 Pros and Cons of Cloud Computing

Pros of cloud computing :

1. **Lower computer costs :** Since applications run in the cloud, not on the desktop PC, your desktop PC does not need the processing power or hard disk space demanded by traditional desktop software.
2. **Improved performance :** Computers in a cloud computing system boot and run faster because they have fewer programs and processes loaded into memory.
3. **Reduced software costs :** Instead of purchasing expensive software applications, you can get most of what you need for free.
4. **Instant software updates :** When you access a web-based application, you get the latest version - without needing to pay for or download an upgrade.
5. **Improved document format compatibility :** You do not have to worry about the documents you create on your machine being compatible with other user's applications or operating systems.
6. **Unlimited storage capacity :** Cloud computing offers virtually limitless storage.
7. **Increased data reliability :** Unlike desktop computing, in which if a hard disk crashes and destroy all your valuable data, a computer crashing in the cloud should not affect the storage of your data.
8. **Universal document access :** All your documents are instantly available from wherever you are.
9. **Latest version availability :** The cloud always hosts the latest version of your documents; as long as you are connected, you are not in danger of having an outdated version.
10. **Easier group collaboration :** Sharing documents leads directly to better collaboration.
11. **Device independence :** Move to a portable device and your applications and documents are still available.

Cons of cloud computing :

1. It requires a constant Internet connection : Cloud computing is impossible if you cannot connect to the Internet.
2. Features might be limited.
3. Stored data might not be secure : With cloud computing, all your data is stored on the cloud.
4. Does not work well with low-speed connections.

University Questions

1. *Define cloud, on demand self service, resource pooling, broad network access, rapid elasticity and measured services.* **GTU : Summer-17, Marks 7**
2. *What are the characteristics of cloud computing ?* **GTU : Winter-17, Marks 3**
3. *Explain pros and cons of cloud computing.* **GTU : Summer-18,Winter-18, Marks 4**
4. *Define cloud computing and state its desirable features.* **GTU : Summer-18, Marks 7**
5. *Define cloud computing with its features.* **GTU : Winter-18, Marks 3**

1.2 Layers and Types of Clouds

GTU : Summer-17, 18, Winter-17, 18

- Cloud computing services are divided into three classes :
  1) Infrastructure as a service,
  2) Platform as a service
  3) Software as a service
- Fig. 1.2.1 shows the layered organization of the cloud stack from physical infrastructure to applications.

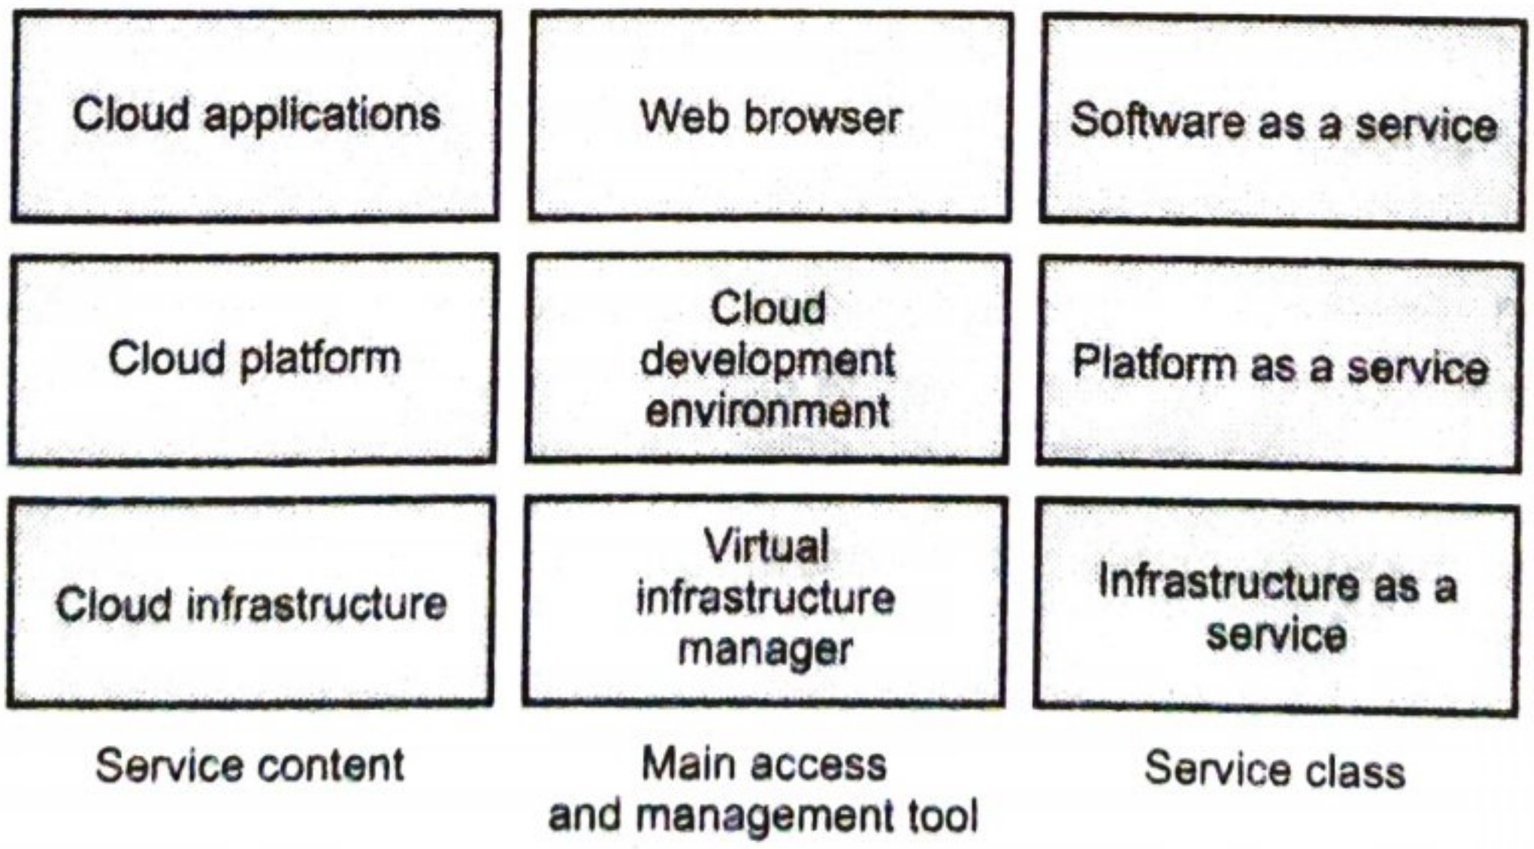

Fig. 1.2.1 Cloud stack

- Cloud services are designed to provide easy, scalable access to applications, resources and services and are fully managed by a cloud services provider.
- A cloud service can exist as a simple web-based software program with a technical interface invoked via the use of a messaging protocol or as a remote access point for administrative tools or larger environments and other IT resources.
- The organization that provides cloud-based IT resources is the cloud provider. Cloud providers normally own the IT resources for lease by cloud consumers, and could also resell IT resources leased from other providers.
- Cloud computing, often described as a stack, has a broad range of services built on top of one another under the name cloud.

1. Software as a service

- SaaS utilizes the internet to deliver applications, which are managed by a third-party vendor, to its users. A majority of SaaS applications run directly through your web browser, which means they do not require any downloads or installations on the client side.
- Salesforce.com uses SaaS model. SaaS applications are designed for end-users, delivered over the web.

2. Infrastructure as a service

- IaaS is a cloud computing service where enterprises rent or lease servers for compute and storage in the cloud. Users can run any operating system or applications on the rented servers without the maintenance and operating costs of those servers.
- Amazon Web Services mainly offers IaaS.
- IaaS is the hardware and software that powers it all - servers, storage, networks, operating systems.

3. Platform as a service

- Platform as a Service (PaaS), provide cloud components to certain software while being used mainly for applications. PaaS delivers a framework for developers that they can build upon and use to create customized applications.
- All servers, storage, and networking can be managed by the enterprise or a third-party provider while the developers can maintain management of the applications.
- Google AppEngine, an example of Platform as a Service.
- PaaS is the set of tools and services designed to make coding and deploying those applications quick and efficient.

1.2.1 Deployment Models

- Cloud deployment models are refers to the location and management of the cloud's infrastructure.
- Deployment models are defined by the ownership and control of architectural design and the degree of available customization. Cloud deployment models are private public and community clouds.
- Fig. 1.2.2 shows cloud deployment model.

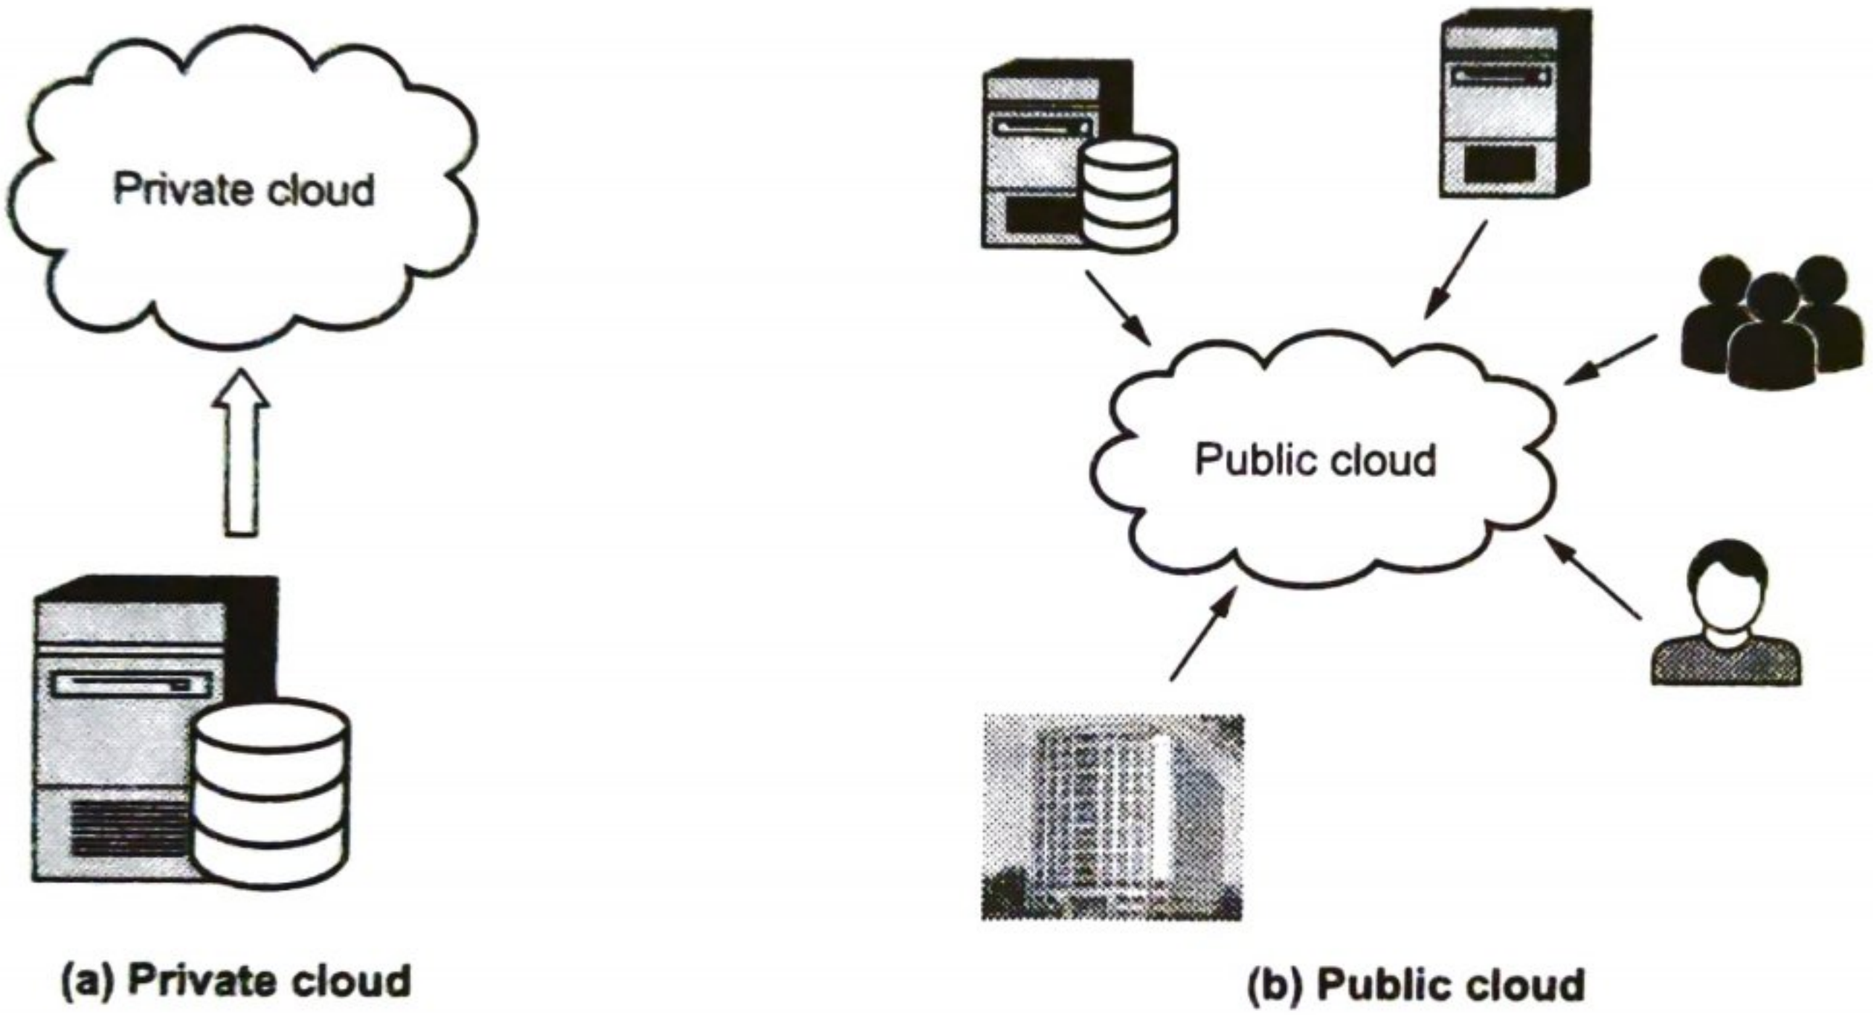

(a) Private cloud (b) Public cloud

Fig. 1.2.2

1. Public cloud :

- The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.
- Public cloud is a huge data centre that offers the same services to all its users. The services are accessible for everyone and much used for the consumer segment.
- Examples of public services are Facebook, Google and LinkedIn.
- **Public cloud benefits :**
  a) Low investment hurdle : Pay for what you use.
  b) Good test/development environment for applications that scale to many servers
- **Public cloud risks :**
  a) Security concerns : Multi-tenancy and transfers over the Internet.
  b) IT organization may react negatively to loss of control over data center function

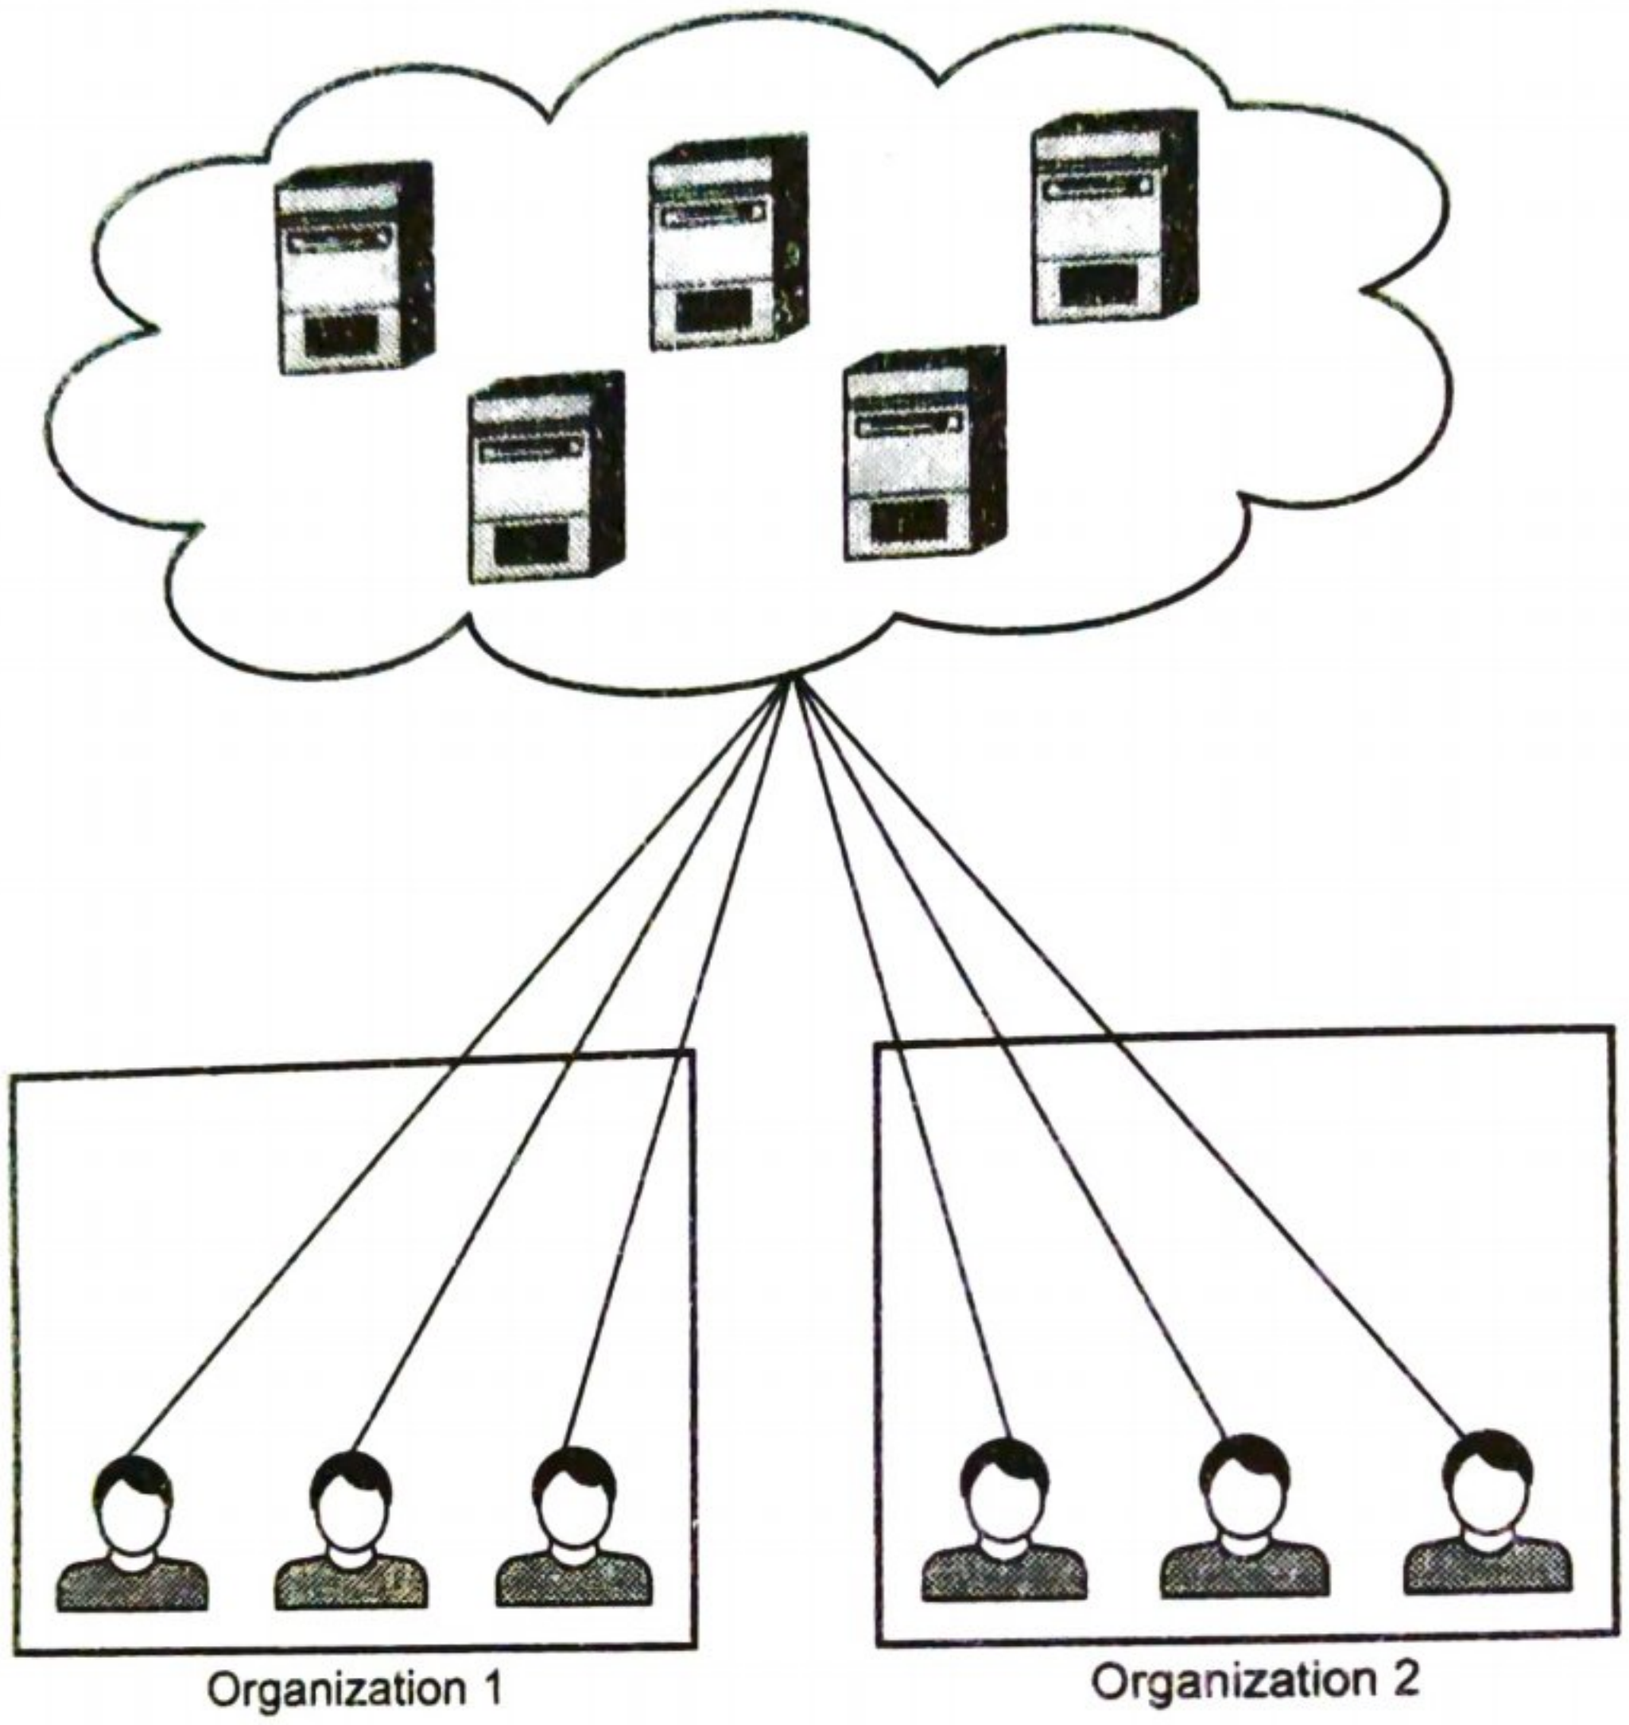

Fig. 1.2.3 : Community cloud

2. Private cloud :

- The cloud infrastructure is operated solely for a single organization. It may be managed by the organization or a third party and may exist on-premises or off-premises.
- **Private cloud benefits :**
  a) Fewer security concerns as existing data center security stays in place.
  b) IT organization retains control over data center.
- **Private cloud risks :**
  a) High investment hurdle in private cloud implementation, along with purchases of new hardware and software.
  b) New operational processes are required; old processes not all suitable for private cloud.

3. Community cloud :

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

4. Hybrid cloud :

- The cloud infrastructure is a composition of two or more clouds (private, community or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

- **Hybrid cloud benefits :**
  a) Operational flexibility : Run mission critical on private cloud, dev/test on public cloud
  b) Scalability : Run peak and bursty workloads on the public cloud
- **Hybrid cloud risks :**
  a) Hybrid clouds are still being developed; not many in real use
  b) Control of security between private and public clouds, some of same concerns as in public cloud

1.2.2 Difference between Public and Private Cloud

Public cloud	Private cloud
Public cloud infrastructure is offered via web applications and also as web services over Internet to the public.	Private cloud infrastructure is dedicated to a single organization.
Support multiple customer	Support dedicated customer
Full utilized of infrastructure.	Does not utilize shared infrastructure
Security is low as compared to private cloud	High level of security
Low cost	High cost
Azure, Amazon Web Services, Google App Engine and Force.com are a few examples of public clouds	An example of the Private Cloud is NIRIX's one Server with dedicated servers

University Questions

1. *Explain public, private, community and hybrid cloud.*

GTU : Summer-17, Winter-18, Marks 7

2. *List cloud deployment models and describe public cloud.* **GTU : Winter-17, Marks 3**
3. *List cloud deployment models and describe community cloud.* **GTU : Summer-18, Marks 4**
4. *Which cloud is better ? public or private. Justify your answer.* **GTU : Summer-18, Winter-18, Marks 4**
5. *List cloud deployment models and describe public cloud.* **GTU : Winter-18, Marks 4**

1.3 Cloud Infrastructure Management

- A key challenge IaaS providers face when building a cloud infrastructure is managing physical and virtual resources, namely servers, storage and networks, in a holistic fashion.
- The orchestration of resources must be performed in a way to rapidly and dynamically provision resources to applications.
- The software toolkit responsible for this orchestration is called a virtual infrastructure manager
- A multi-tenant cloud is a cloud computing architecture that allows customers to share computing resources in a public or private cloud. Each tenant's data is isolated and remains invisible to other tenants.
- It allows multiple users to work in a software environment at the same time, each with their own separate user interface, resources and services. The multitenant application design was created to enable multiple users (tenants) to access the same application logic simultaneously.
- Tenants can individually customize features of the application such as :
  1. **User interface :** Tenants can define a specialized look for their application interface.
  2. **Business process :** Tenants can customize the rules, logic and workflows of the business processes that are implemented in the applications.
  3. **Data model :** Tenants can extend the data schema of the application to include exclude or rename fields in the application data structures.
  4. **Access control :** Tenants can independently control the access rights for users and groups.

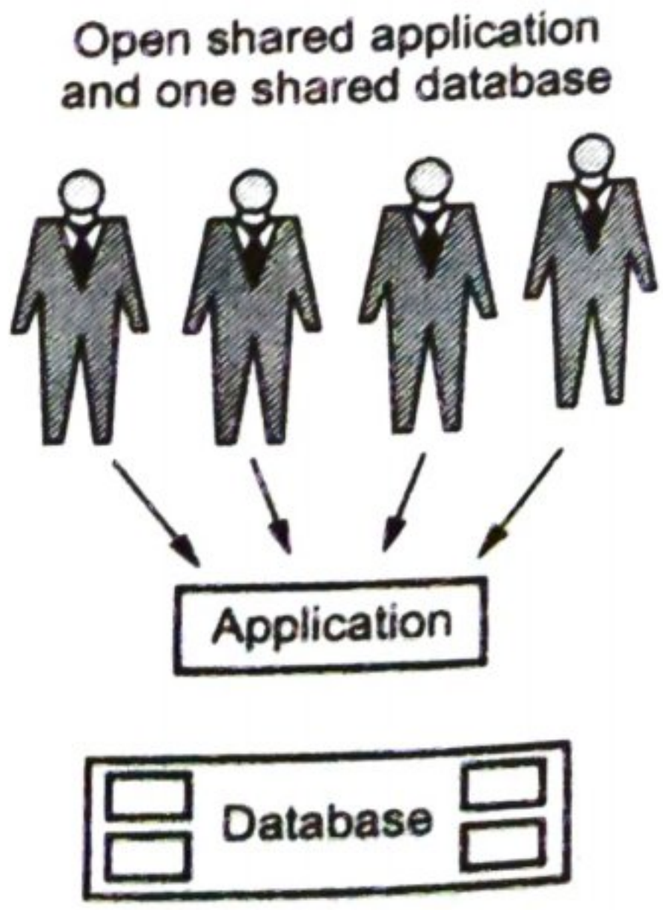

Fig. 1.3.1 Multi-tenant technology

- Benefits of a multitenancy technology :
  1. **Costs savings** : It yields tremendous economy of scale for the provider so he can offer the service at a lower cost to customers.
  2. **Improved quality, user satisfaction and customer retention** : A multitenant application is one large community hosted by the provider which can gather operational information from the collective user population and make frequent, incremental improvements to the service that benefit the entire user community at once.
  3. **Improved security** : Most current enterprise security models are perimeter-based, making them vulnerable to inside attacks.
- Self-service, on-demand resource provisioning. Self-service access to resources has been perceived as one the most attractive features of clouds. This feature enables users to directly obtain services from clouds, such as spawning the creation of a server and tailoring its software, configurations, and security policies, without interacting with a human system administrator.
- Storage virtualization : Virtualizing storage means abstracting logical storage from physical storage. Storage devices are commonly organized in a storage area network (SAN) and attached to servers via protocols such as Fibre Channel, iSCSI, and NFS; a storage controller provides the layer of abstraction between virtual and physical storage.
- Virtual networking : Virtual networks allow creating an isolated network on top of a physical infrastructure independently from physical topology and locations. A virtual LAN (VLAN) allows isolating traffic that shares a switched network, allowing VMs to be grouped into the same broadcast domain.

1.4 Challenges and Applications

1. Increased Security Vulnerabilities
2. Reduced Operational Governance Control
3. Limited Portability Between Cloud Providers
4. Multi-Regional Compliance and Legal Issues

- Use of cloud for business purpose means that the responsibility over data security becomes shared with the cloud provider. Organization extends their trust boundary to cloud consumer to external cloud.
- It is clear that the security issue has played the most important role in hindering cloud computing acceptance.
- Without doubt, putting your data, running your software on someone else's hard disk using someone else's CPU appears daunting to many.

- Well-known security issues such as data loss, phishing, pose serious threats to organization's data and software.

1.5 Virtualization

GTU : Summer-17, Winter-18

- Virtualization is a broad term that refers to the abstraction of resources across many aspects of computing. For our purposes : One physical machine to support multiple virtual machines that run in parallel.
- Virtualization is a frame work or methodology of dividing the resources of computer into multiple execution environments.
- Virtualization is an abstraction layer that decouples the physical hardware from the operating system to deliver greater IT resource utilization and flexibility.
- It allows multiple virtual machines, with heterogeneous operating systems to run in isolation, side-by-side on the same physical machine.
- Fig. 1.5.1 shows before and after virtualization.
- Virtualization means running multiple machines on a single hardware. The "Real" hardware invisible to operating system. OS only sees an abstracted out picture. Only Virtual Machine Monitor (VMM) talks to hardware.

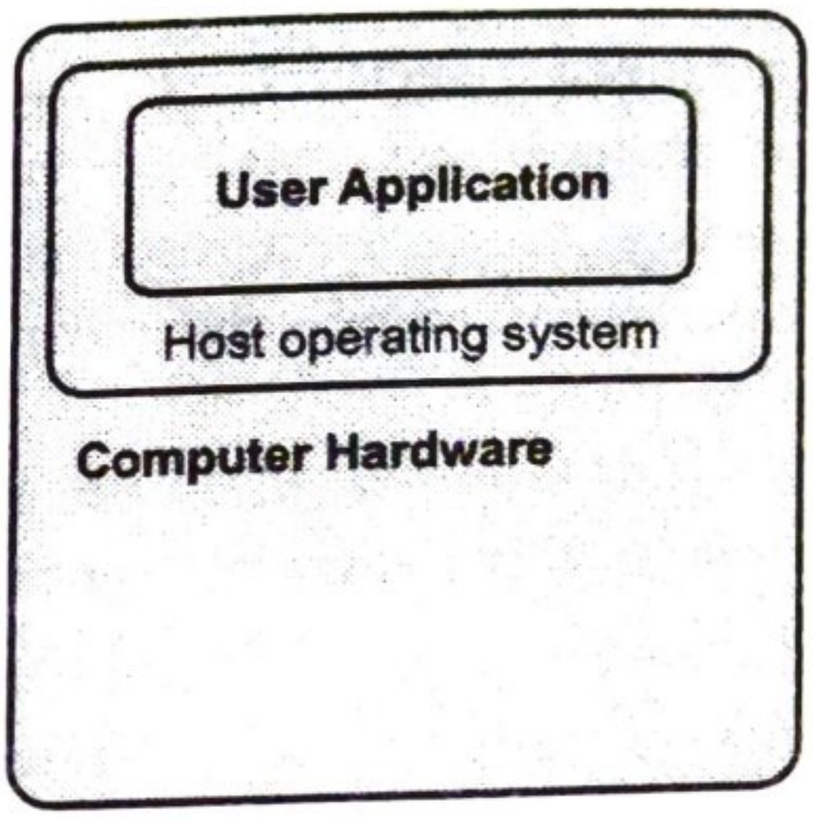

(a) : Before virtualization

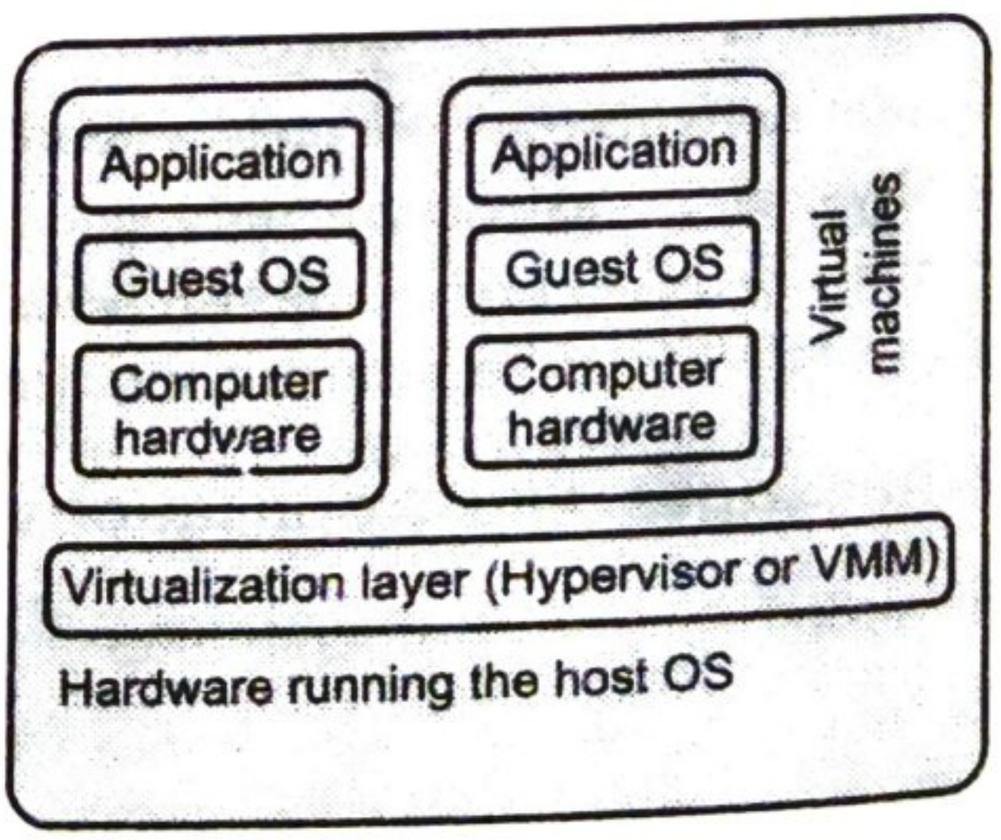

(b) After virtualization

Fig. 1.5.1

- It is "a technique for hiding the physical characteristics of computing resources from the way in which other systems, applications, or end users interact with those resources.
- This includes making a single physical resource appear to function as multiple logical resources; or it can include making multiple physical resources appear as a single logical resource."

- It is divided into two main categories :
  1. Platform virtualization involves the simulation of virtual machines.
  2. Resource virtualization involves the simulation of combined, fragmented or simplified resources.
- Fig. 1.5.2 shows taxonomy of virtualization.

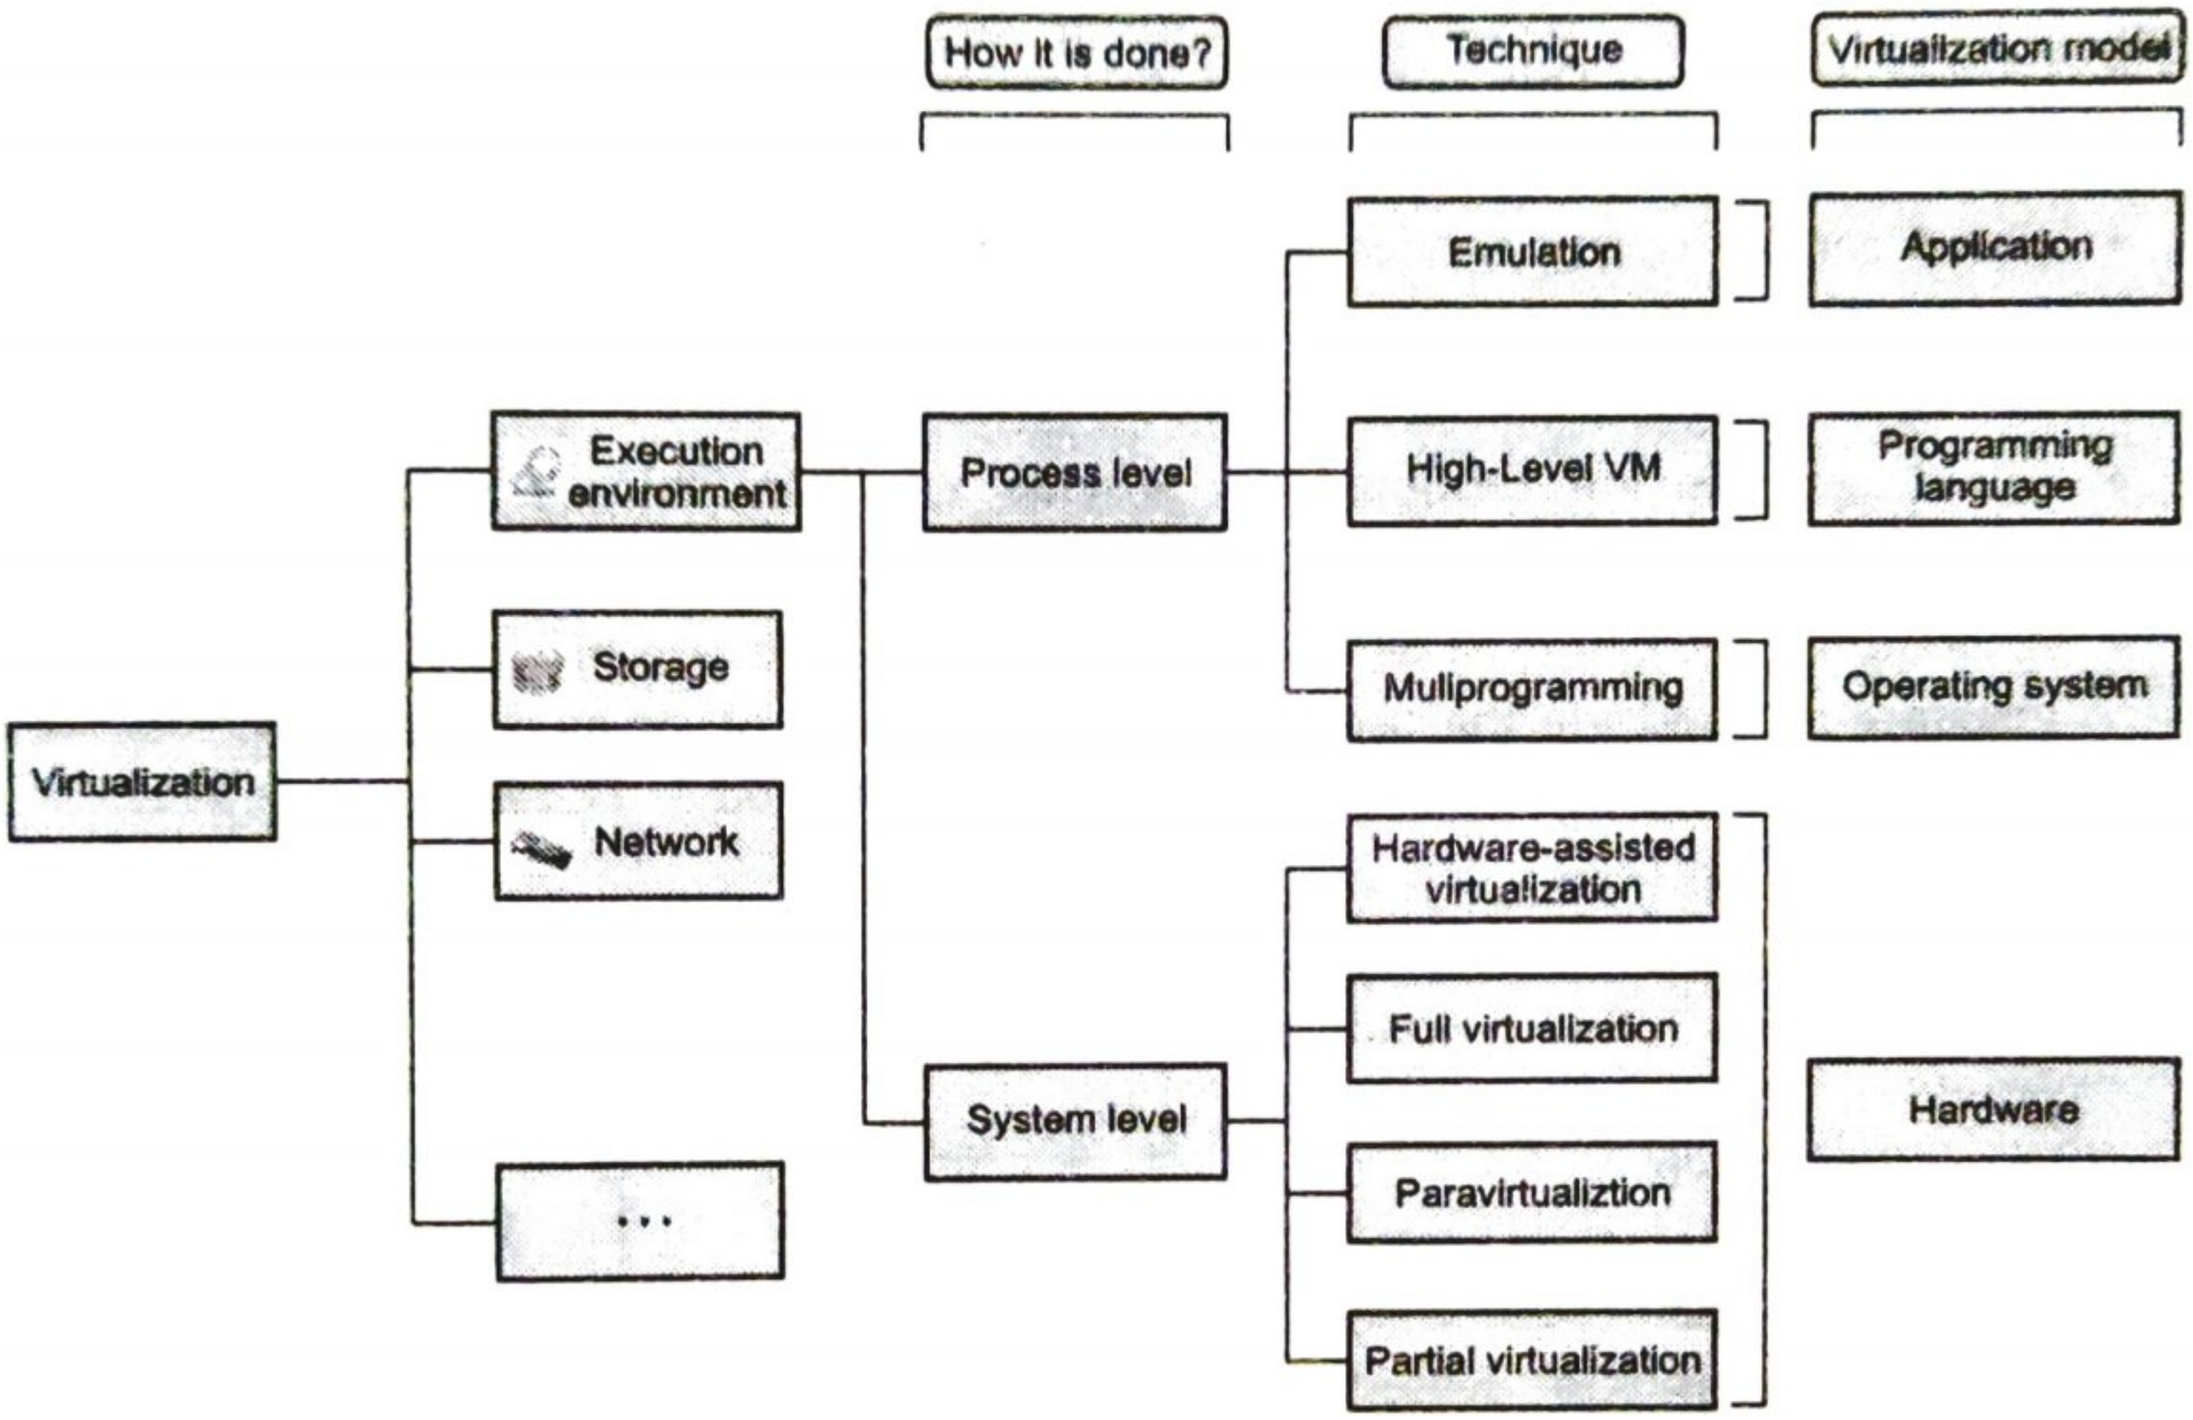

Fig. 1.5.2 Taxonomy of virtualization

- Virtualization is mainly used to emulate execution environment, storage and network. Execution environment classified into two types : process level and system level.
- Process level is implemented on top of an existing operating system.
- System level is implemented directly on hardware and do not or minimum requirement of existing operating system.

University Question

1. *Explain virtualization and hypervisor.* **GTU : Summer-17, Winter-18, Marks 7**

1.6 Cloud Services

GTU : Summer-17, 18, Winter-17, 18

- Service models describe the type of service that the service provider is offering. The best-known service models are Software as a Service, Platform as a Service, and Infrastructure as a Service.
- The service models build on one another and define what a vendor must manage and what the client's responsibility is.
- Service models : This consists of the particular types of services that you can access on a cloud computing platform.
- Cloud service is any service made available to users on demand via the Internet from a cloud computing provider's servers as opposed to being provided from a company's own on-premises servers.
- Cloud services are designed to provide easy, scalable access to applications, resources and services, and are fully managed by a cloud services provider.
- A cloud service can exist as a simple web-based software program with a technical interface invoked via the use of a messaging protocol, or as a remote access point for administrative tools or larger environments and other IT resources.
- The organization that provides cloud-based IT resources is the cloud provider. Cloud providers normally own the IT resources for lease by cloud consumers, and could also resell IT resources leased from other providers.
- Cloud computing, often described as a stack, has a broad range of services built on top of one another under the name cloud.
- Fig. 1.6.1 shows cloud computing stack.
- Flavors of Cloud Computing is as follows;

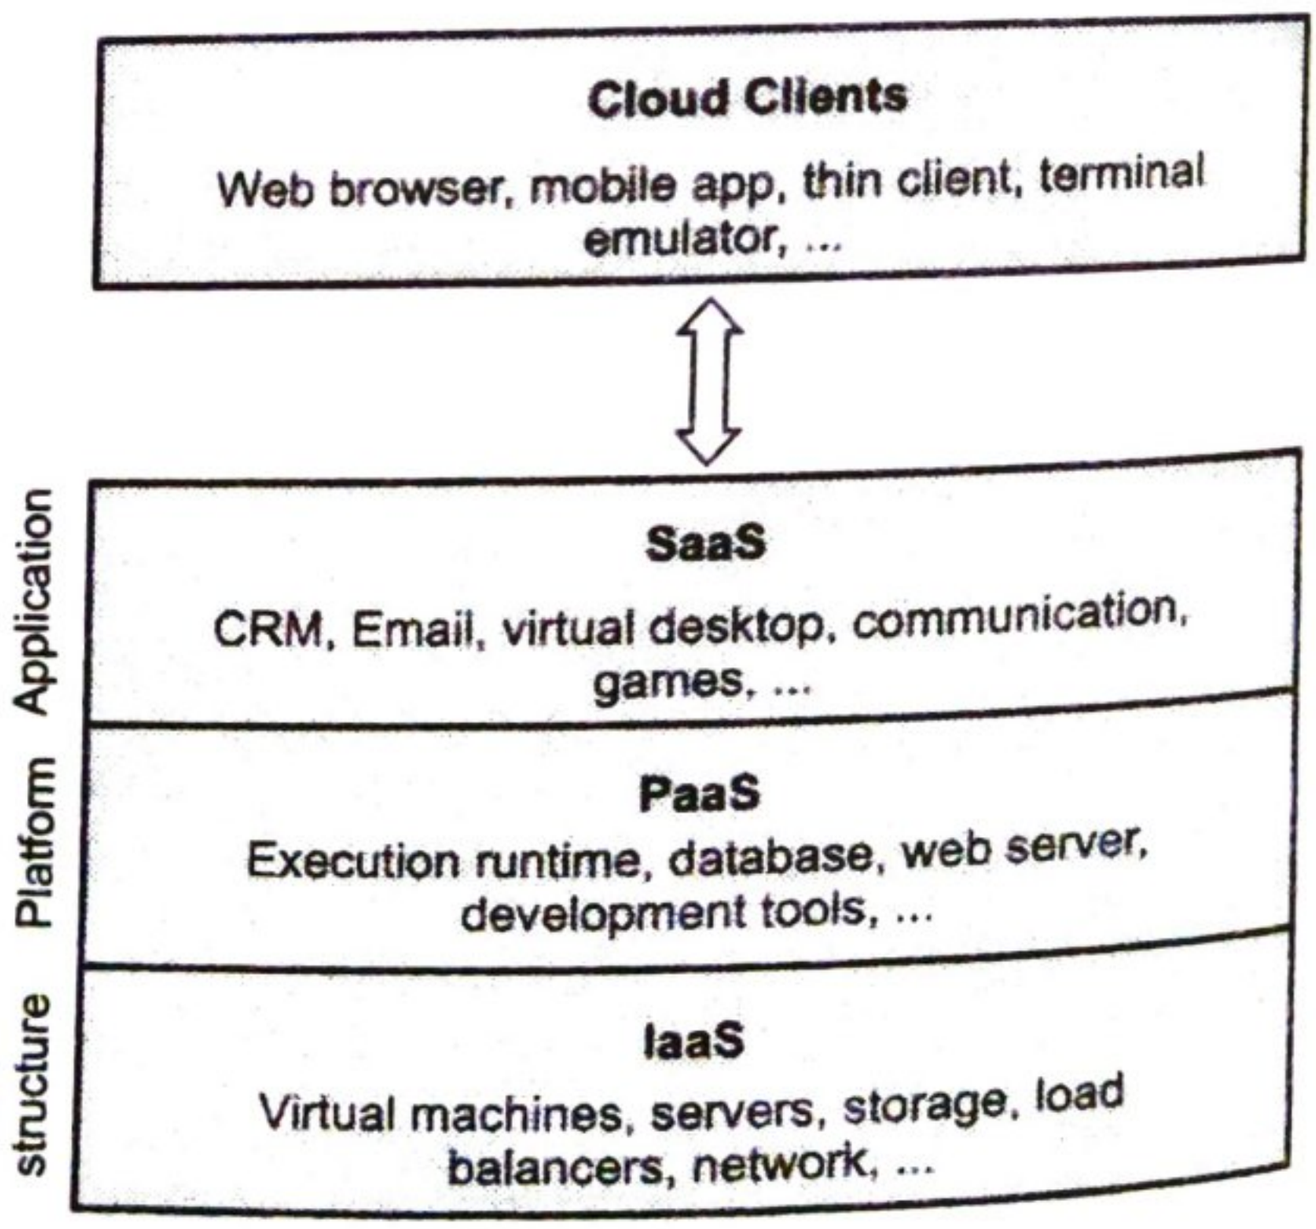

Fig. 1.6.1 : Cloud computing stack

1. SaaS applications are designed for end-users, delivered over the web
2. PaaS is the set of tools and services designed to make coding and deploying those applications quick and efficient

3. IaaS is the hardware and software that powers it all - servers, storage, networks, operating systems

1.6.1 Software as a Service (SaaS)

- Model in which an application is hosted as a service to customers who access it via the Internet.
- The provider does all the patching and upgrades as well as keeping the infrastructure running.
- The traditional model of software distribution, in which software is purchased for and installed on personal computers, is referred to as product.
- In this model, the user, client or consumer runs an application from a cloud infrastructure. Through an interface such as a web browser, the client or user may access this application from a variety of devices.
- The complete application is offered as on demand service. This saves the client from having to invest in any software licenses or servers up front, and can save the provider money since they are maintaining and providing only a single application.
- In this model, the client does not manage cloud infrastructure, networks or servers, storage, or operating systems. Even, Microsoft, Google, and Zoho offer SaaS.
- The SaaS concept can be defined as providing robust "web-based, on-demand software, storage and various applications" to organizations.
- The SaaS model has emerged as an alternative to traditional one-time licensing for providing and maintaining the software needed by knowledge workers within organizations.
- Fig. 1.6.2 shows SaaS.

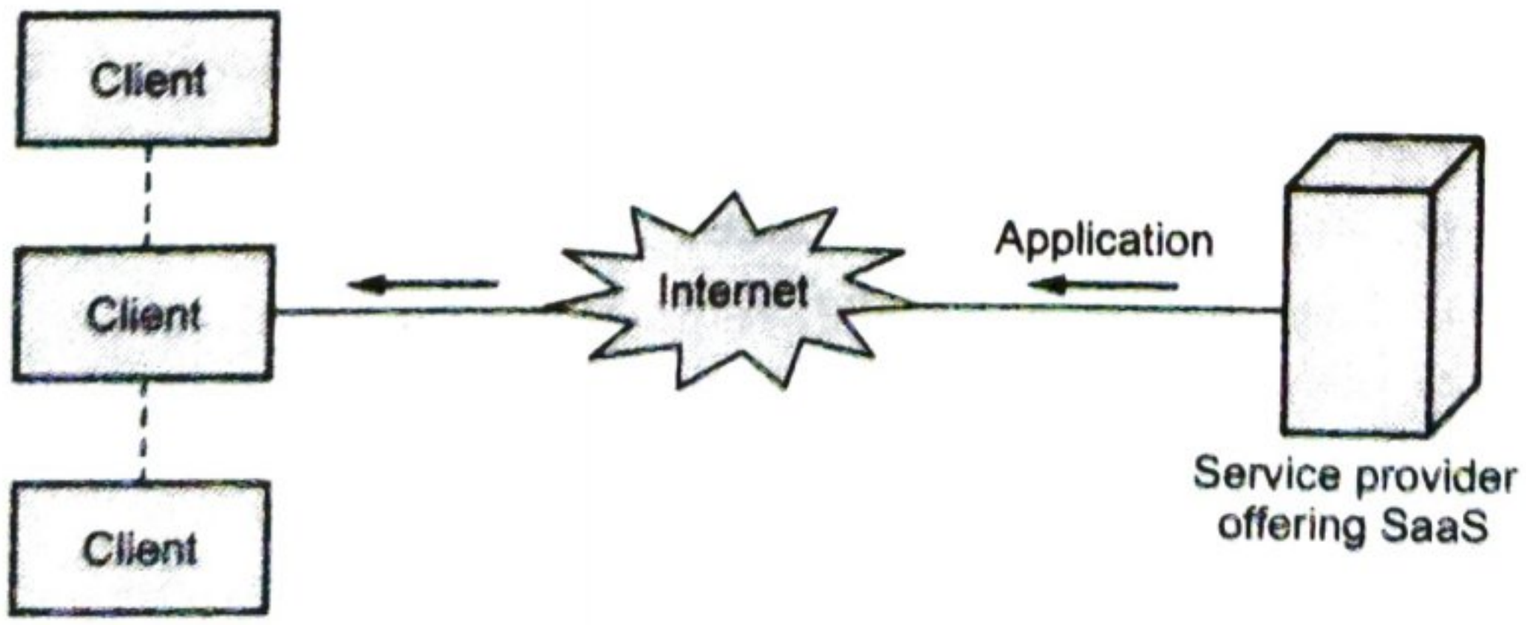

Fig 1.6.2 SaaS

Characteristics of SaaS :

1. Software applications or services are stored remotely.
2. A user can then access these services or software applications via the Internet.
3. In most cases, a user does not have to install anything onto their host machine, all they require is a web browser to access these services and in some cases, a browser may require additional plug-in/add-on for certain services.
4. Network-based management and access to commercially available software from central locations rather than at each customer's site, enabling customers to access applications remotely via the Internet.
5. Application delivery from a one-to-many model, as opposed to a traditional one-to- one model.

Benefits of SaaS :

1. You only pay for what you use
2. Easier administration and invoicing
3. Automatic updates and patch management
4. Compatibility: all users have access to the same version of software
5. Easier collaboration
6. It support automated update and patch management services

1.6.2 Platform as a Service (PaaS)

- Platform as a service is another application delivery model and also known as cloud-ware. Supplies all the resources required to build applications and services completely from the Internet, without having to download or install software.
- Services include : Application design, development, testing, deployment, and hosting, team collaboration, web service integration, database integration, security, scalability, storage, state management, and versioning.
- PaaS is closely related to SaaS but delivers a platform from which to work rather than an application to work with.
- This model involves software encapsulated and offered as a service, from which higher levels of service may then be built. The user, customer, or client in this model is the one building applications which then run on the provider's infrastructure.
- This in turn provides customers and clients with the capability to deploy applications onto the cloud infrastructure using programming tools and languages, which the provider supports.

- The customer still does not manage the framework, network, servers or operating system, but has control over deployed applications and sometimes over the hosting environment itself.
- Some examples of Platform as a Service include Google's App Engine or Force.com
- PaaS consists of following components :
  1. Browser based development studio
  2. Pay contrary to billing
  3. Management and supervising tools
  4. Seamless deployment to host run time environment.
- **Characteristics of PaaS :**
  1. It support multi-tenant architecture.
  2. It support for development of group collaboration.
  3. PaaS systems can be deployed as public cloud services or as private cloud services.
  4. Provision of runtime environments. Typically each runtime environment supports either one or a small set of programming languages and frameworks
  5. Support for custom applications. Support for the development, deployment and operation of custom applications.
  6. Preconfigured capabilities. Many PaaS systems are characterized by capabilities that are preconfigured by the provider, with a minimum of configuration available to developers and customer operations staff.
  7. Support for porting existing applications. While many PaaS systems are primarily designed to support "born on the cloud" applications.
  8. Security is an important characteristic in PaaS. It needs to provide authentication and authorization to differentiate the access rights of different users

Benefits of Paas :

1. Scalability including rapid allocation and deallocation of resources with a pay-as-you-use model
2. Reduced capital expenditure
3. Reduced lead times with on-demand availability of resources
4. Self-service with reduced administration costs
5. Reduced skill requirements
6. Support of team collaboration
7. Ability to add new users quickly

1.6.3 Infrastructure as a Service (IaaS)

- IaaS gives the storage room likeness to the in-house datacenter stood out from various organizations sorts.
- Center datacenter framework segments are capacity, servers (registering units), the system itself, and administration apparatuses for foundation upkeep and checking.
- Each of these parts has made a different market specialty. While some little organizations have practical experience in just a single of these IaaS cloud specialties, vast cloud suppliers like Amazon or Right Scale have offerings over all IaaS territories.
- Fig. 1.6.3 shows IaaS.

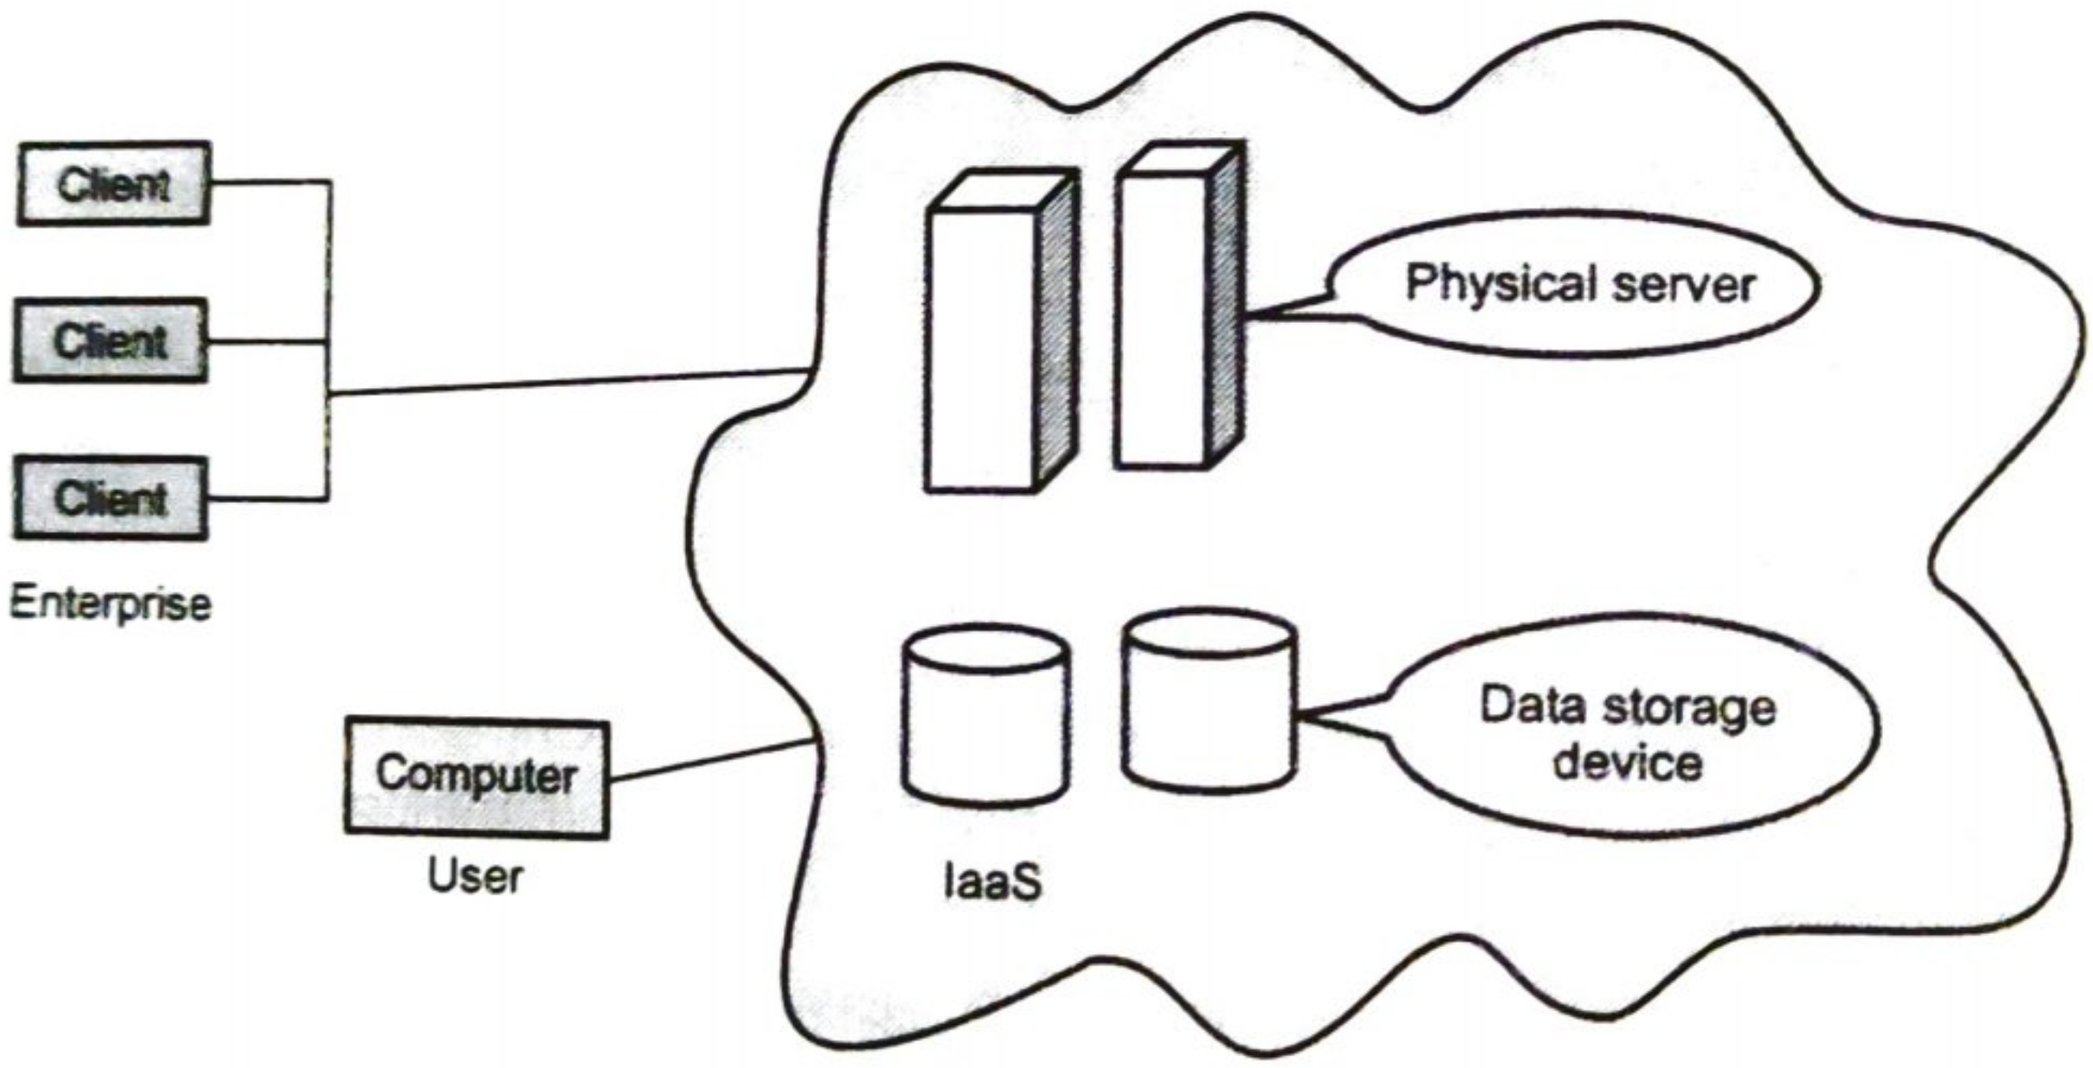

Fig. 1.6.3 IaaS

- It offers the hardware so that your organization can put whatever they want onto it. Rather than purchase servers, software, racks, and having to pay for the datacenter space for them, the service provider rents those resources :
  1. Server space 2. Network equipment
  3. Memory 4. CPU cycles 5. Storage space
- Again, the customer is not managing cloud infrastructure, but in this case, the customer does control operating systems, deployed applications, storage, and sometimes-certain networking components
- Examples : Amazon EC2, Rackspace Mosso, GoGrid
- IaaS server types :
  1. **Physical server :** Actual hardware is allocated for the customer's dedicated use.

2. **Dedicated virtual server :** The customer is allocated a virtual server, which runs on a physical server that may or may not have other virtual servers.
3. **Shared virtual server :** The customer can access a virtual server on a device that may be shared with other customers.

Advantages of IaaS :

1. Elimination of an expensive and staff-intensive data center
2. Ease of hardware scalability
3. Reduced hardware cost
4. On-demand, pay as you go scalability
5. Reduction of IT staff
6. Suitability for ad hoc test environments
7. Allows complete system administration and management
8. Support multiple tenants

1.6.4 Difference between IaaS, PaaS and SaaS

IaaS	PaaS	SaaS
IaaS gives users automated and scalable environments	PaaS provides a framework for quickly developing and deploying applications	SaaS makes applications available through the internet.
Amazon Web Services, for example, offers IaaS through the Elastic Compute Cloud, or EC2	Google Cloud Platform provides another PaaS option in App Engine	SaaS applications such as Gmail, Dropbox, Salesforce, or Netflix
In Iaas, infrastructure as a service.	In Paas, platform as a service	In SaaS, software as a service
Virtual platform on which required operating environment and application deployed	Operating environment included	Operating environment largely irrelevant, fully functional application provided
IaaS is a cloud service that provides basic computing infrastructure: servers, storage, and networking resources. In other words, IaaS is a virtual data center	PaaS refers to cloud platforms that provide runtime environments for developing, testing, and managing applications	SaaS allows people to use cloud-based web applications.

Major IaaS providers include Amazon Web Services, Microsoft Azure, and Google Compute Engine.	Examples of PaaS services are Heroku and Google App Engine.	email services such as Gmail and Hotmail are examples of cloud-based SaaS services.
IaaS services are available on a pay-for-what-you-use model	PaaS solutions are available with a pay-as-you-go pricing model.	SaaS services are usually available with a pay-as-you-go pricing model
Used by IT administrator	Used by software developers	Used by end user

1.6.5 Identity as a Service

- Identity as a Service (IDaaS) is cloud-based authentication operated by a third-party provider.
- Identity as a service (IDaaS) are SaaS-based Identity And Access Management (IAM) offerings that allow organizations to use Single Sign-on (SSO using SAML or OIDC), authentication and access controls to provide secure access to their growing number of software and SaaS applications.
- Five key capabilities are required to make enterprise IDaaS solutions possible :
  1. **Single Sign-on (SSO) :** With single sign-on employees, partners and customers obtain easy, fast and secure access to all SaaS, mobile and enterprise applications with a single authentication using corporate credentials.
  2. **Multi-factor Authentication (MFA) :** MFA typically includes adaptive authentication methods-options to step up as risk increases based on situational changes, user behavior or application sensitivity.
  3. **Access Security :** Access security is policy-based access management for applications and APIs to enhance security beyond SSO.
  4. **Directory :** While most enterprises prefer to integrate IDaaS with their existing user stores, they may use a cloud directory, especially to support customers and/or partners.
  5. **Provisioning :** Through SCIM support and integration with on-premises provisioning, user data is synced with web and enterprise applications.
- IDaaS supplies cloud-based authentication or identity management to enterprises who subscribe. The goal is to ensure users are who they claim to be, and to give them the right kinds of access to software applications, files, or other resources at the right times. If the infrastructure to make this happen is built on site, then the company has to figure out what to do every time a problem comes up.

Advantages of IDaaS :

1. Deliver access services efficiently and cost-effectively.
2. Protect against internal and external security threats
3. With IDaaS, costs drop to the subscription fee and the administration work
4. Your team has to keep up servers; purchase, upgrade, and install software; back up data regularly; pay hosting fees

1.6.6 Pods, Aggregation, Silos

- Workloads support a certain number of users. When the workload reaches the limit of largest virtual machine instance possible, a copy or clone of the instance is required. A group of users within a particular instance is called a **pod.**
- Pods are managed by a Cloud Control System (CCS). Fig. 1.6.4 shows pods, aggregation and failover in IaaS.

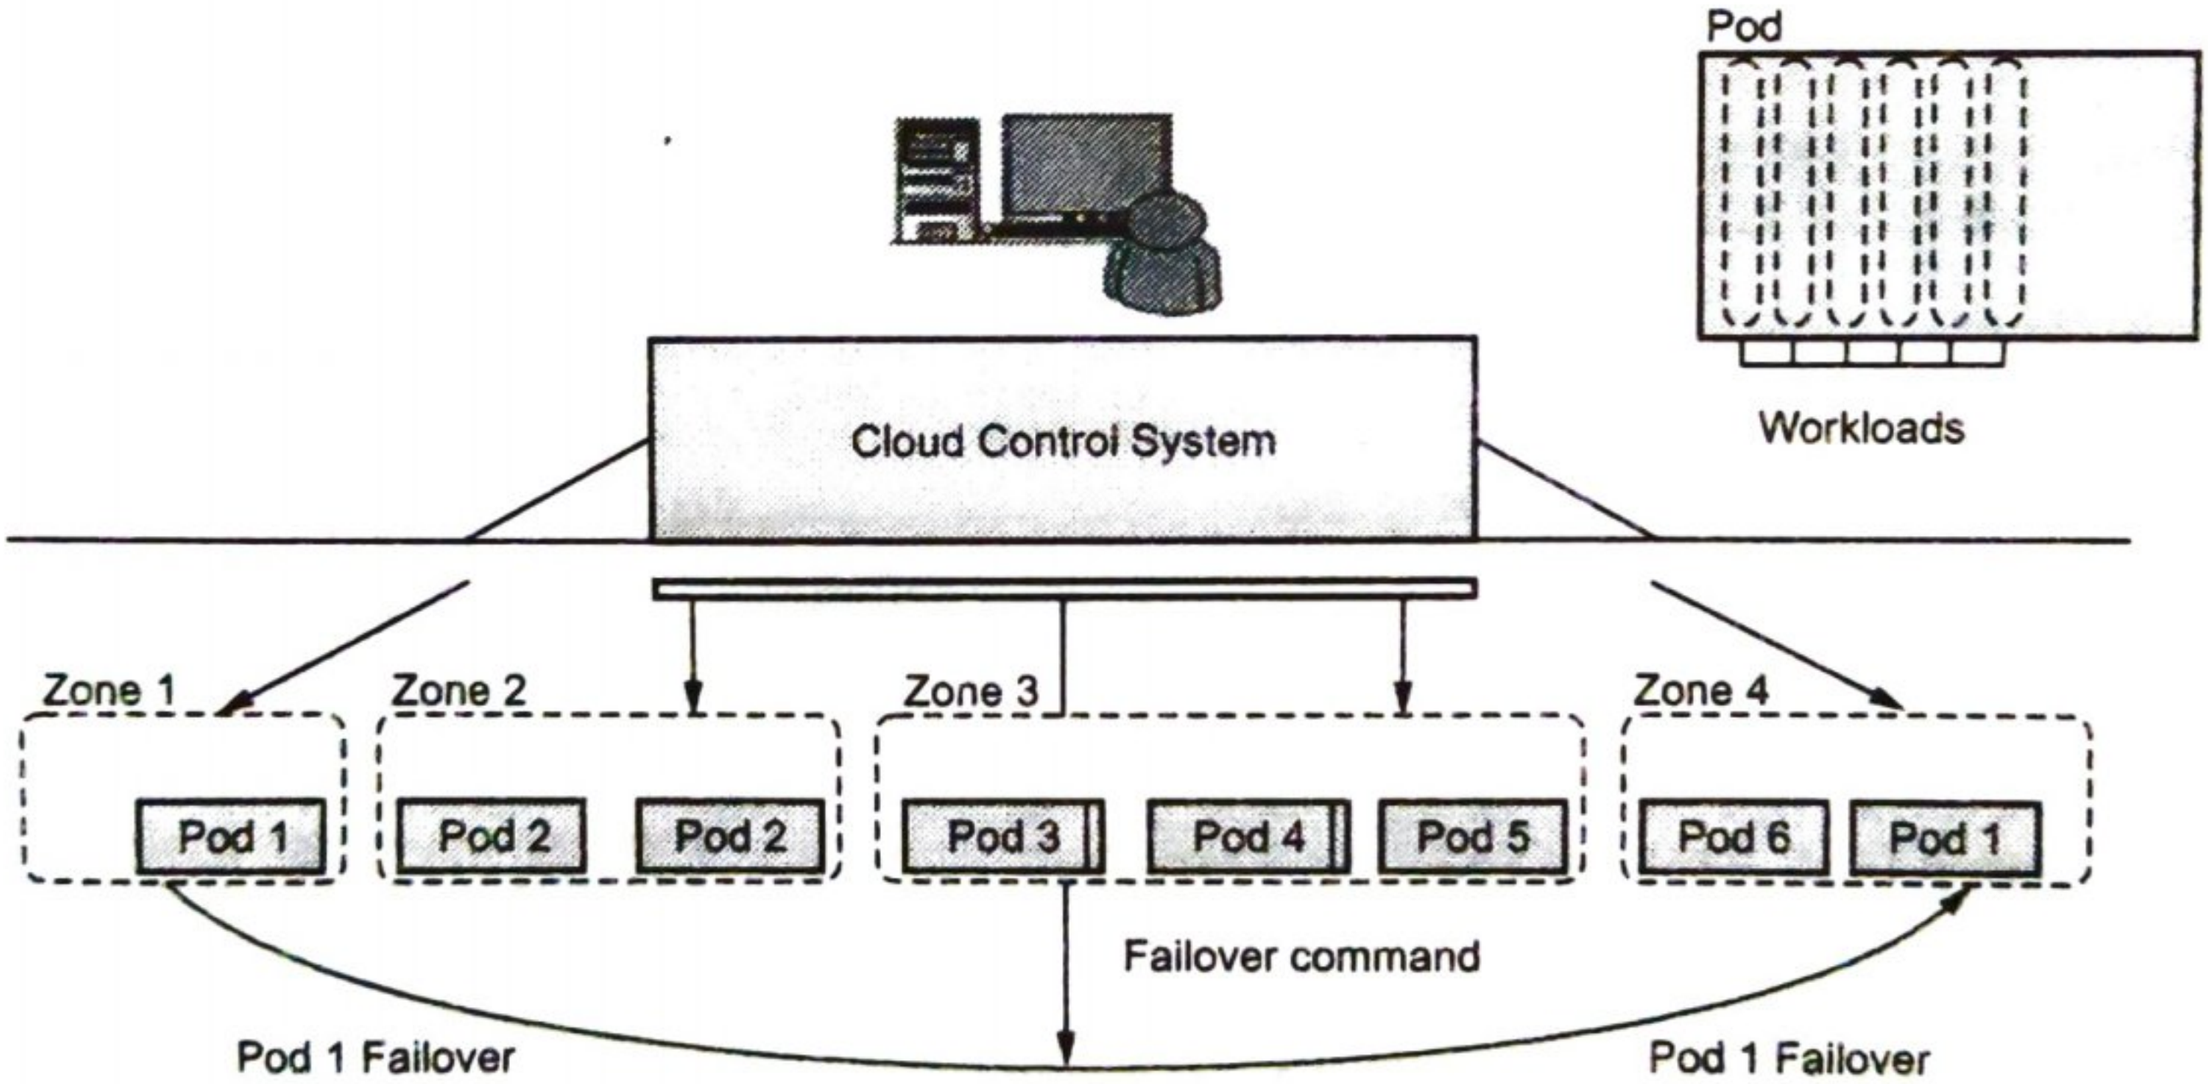

Fig. 1.6.4 Pods, aggregation and failover in IaaS

- Sizing limitation of pod need to be considered when building large cloud-based application. Pods are aggregated into pools within IaaS region or site called an **availability zone.**
- When the computing infrastructure isolates user clouds from one another, so that interoperating is impossible this creates an information silo, or simply **silo.**
- Silos are the cloud computing equivalent of compute islands : They are processing domains that are sealed off from the outside.

- When a cloud computing infrastructure isolates user clouds from each other so the management system is incapable of interoperating with other private clouds, it creates an information silo, or simply a silo.
- Most often, the term silo is applied to PaaS offerings such as Force.com or QuickBase, but silos often are an expression of the manner in which a cloud computing infrastructure is architected.
- Silos are the cloud computing equivalent of compute islands : They are processing domains that are sealed off from the outside.
- When you create a private virtual network within an IaaS framework, the chances are high that you are creating a silo.
- Silos impose restrictions on interoperability that runs counter to the open nature of build-componentized service-oriented applications.
- However, that is not always a bad thing. A silo can be its own ecosystem; it can be protected and secured in ways that an open system can't be. Silos just aren't as flexible as open systems and are subject to vendor lock-in.

University Questions

1. *Explain saas wtih an example.* **GTU : Summer-17, Marks 7**
2. *Compare different cloud service provider ?* **GTU : Summer-17, 18, Winter-17, 18, Marks 7**
3. *Illustrate different service midel of cloud.* **GTU : Summer-17, Marks 7**
4. *Describe layers and types of cloud computing services.* **GTU : Winter-17, Winter-18, Marks 7**
5. *Write a note on Cloud Computing Stack.* **GTU : Summer-18, Marks 7**
6. *Explain Pod, Silos and aggregations in terms of cloud.* **GTU : Summer-18, Marks 3**
7. *Explain IAAS, PAAS and SAAS with a example.* **GTU : Winter-18, Marks 7**

1.7 Multiple Choice Questions

Q.1 Point out the wrong statement :

a Abstraction enables the key benefit of cloud computing : shared, ubiquitous access.

b Virtualization assigns a logical name for a physical resource and then provides a pointer to that physical resource when a request is made.

c All cloud computing applications combine their resources into pools that can be assigned on demand to users.

d All of the mentioned.

Q.2 Point out the wrong statement :

a The massive scale of cloud computing systems was enabled by the popularization of the Internet.

b Soft computing represents a real paradigm shift in the way in which systems are deployed.

c Cloud computing makes the long-held dream of utility computing possible with a pay-as-you-go, infinitely scalable, universally available system.

d All of the mentioned.

Q.3 Which of the following is essential concept related to Cloud ?

a Reliability

b Productivity

c Abstraction

d All of the mentioned.

Q.4 Point out the wrong statement :

a All applications benefit from deployment in the cloud.

b With cloud computing, you can start very small and become big very fast.

c Cloud computing is revolutionary, even if the technology it is built on is evolutionary.

d None of the mentioned.

Q.5 Which of the following cloud concept is related to pooling and sharing of resources ?

a Polymorphism

b Abstraction

c Virtualization

d None of the mentioned

Q.6 CDC Stands for _____.

a Cloud Data Computing

b Cloud Data Cluster

c Cloud Data Center

d Computing Data on Cloud

Q.7 In which year IoT was introduced ?

a 1999

b 1998

c 1996

d 1997

Q.8 _____ is mainly used to utilize idle resources in the nodes.

a Dedicated cluster

b Enterprise cluster

c Distributed Cluster

d Centralized cluster

Q.9 _____ is the entry point into the cloud for user and administrators.

a Cloud manager
b Group manager
c Instance manager
d VM manager

Q.10 The services that provides utility may directly correlate with their _____ ?

a expectations
b profit
c time
d satisfaction

Q.11 The promise of _____ has raised the IT expectations of small and medium enterprises beyond measures.

a cloud computing
b cloud computing
c cluster computing
d client-server computing

Q.12 Cloud supports applications _____ and resources elasticity.

a performance
b reliability
c scalability
d none of the above

Q.13 _____ cloud is shared by several organizations and supports a specific community that has shared concerns.

a Private
b Public
c Managed
d Community

Q.14 _____ of grids/clouds to offer standard interfaces for dynamically scalable services delivery in their products.

a Vendors
b Consumers
c Integrators
d providers

Q.15 VAN stands for _____ .

a Virtual Application Network
b Virtual Acceptable Network
c Virtual Admission Network
d Virtual Area Network

Q.16 In what type of cluster, the nodes are closely packaged in one or more racks Sitting in a room and the nodes are not attached to peripherals.

a Compact
b Slack
c Loosely coupled
d Tightly coupled

Q.17 Virtualization is a computer architecture technology by which multiple _____ are multiplexed in the same hardware machine.

a virtual memory
b virtual machines
c physical machines
d virtual machine monitor

Q.18 Which of the following is not VI manager ?

a Apache VCL
b App Logic
c Google VI
d Nimbuz 3

Q.19 According to Infosys how many steps are there in the migration model ?

a 4
b 5
c 6
d 7

Q.20 _____ is the process of transferring data between storage types, formats or systems.

a Data mediation
b Data integrity
c Data modification
d Data migration

Q.21 Which of the following is not principle of the cloud ?

a Federation
b Non independent
c Trust
d Isolation

Q.22 _____ is a directory on the cluster node where a VM is running.

a Virtualization
b KVM
c Virtual machine directory
d VMware

Q.23 Which of the following is NOT Cloud application features ?

a Multitenancy
b Elasticity
c Homogeneous cloud platform
d On-demand service

Q.24 Usually, when accessing the public or private cloud, the users require minimum _____, which is sometimes defined by the cloud providers ?

a frequency
b bandwidth
c Internet
d all of these

Q.25 Cloud architecture consists of a _____ set of components that collectively describe the way the cloud works.

a horizontal
b vertical
c hierarchical
d all of these

Q.26 Which of the following is NOT phases of cloud migration ?

a Migration strategy
b Prototyping
c Provisioning
d Debugging

Q.27 _____ are a set of agreements that are signed between the user and service providers.

a Service level agreement
b Service oriented architecture
c Service layer agreement
d Software level agreement

Q.28 Which of the following type of virtualization is also characteristic of cloud computing ?

a Storage
b Application
c CPU
d All of the mentioned

Q.29 Which of the following network resources can be load balanced ?

a Connection through intelligent switches
b DNS
c Storage resources
d All of the mentioned

Q.30 How many types of virtual private server instances are partitioned in an IaaS stack ?

a One
b Two
c Three
d All of the mentioned

Q.31 Which of the following is associated with considerable vendor lock-in ?

a PaaS
b IaaS
c CaaS
d SaaS

Q.32 _____ for both hosted on premises applications and data sources.

a IaaS
b PaaS
c SaaS
d DaaS

Answer Keys for Multiple Choice Questions :

Q.1	c	Q.2	b	Q.3	c	Q.4	a
Q.5	c	Q.6	c	Q.7	a	Q.8	b
Q.9	a	Q.10	a	Q.11	b	Q.12	c
Q.13	d	Q.14	a	Q.15	d	Q.16	a
Q.17	b	Q.18	b	Q.19	d	Q.20	d
Q.21	a	Q.22	b	Q.23	c	Q.24	b
Q.25	c	Q.26	d	Q.27	a	Q.28	c
Q.29	d	Q.30	c	Q.31	a	Q.32	a

❑❑❑

2 Software as a Service

Syllabus

Evolution of SaaS, Challenges of SaaS Paradigm, SaaS Integration Services, SaaS Integration of Products and Platforms. Infrastructure As a Services (IaaS): Introduction, Background & Related Work, Virtual Machines Provisioning and Manageability, Virtual Machine Migration Services, VM Provisioning and Migration in Action. Platform As a service (PaaS): Integration of Private and Public Cloud, Technologies and Tools for Cloud Computing, Resource Provisioning services.

Contents

2.1 Evolution of SaaS

- Software was initially sold out in the form of floppy discs. The airline management systems of earlier days used these kinds of floppy discs which was used to store data.
- Although software took the form of products, there happened a change, which could not be called as big as a revolution, when floppy discs were replaced by Compact Discs and DVDs. Software as a product had one feature, having only one application.
- Software after the internet era could be called as a revolution when the software-as-a-product-market, had to go through a huge downfall, since people started downloading software directly from the internet instead of buying them.
- Next came the era of cloud computing, which made the process even simpler, when there was no necessity for even the installation of the software on the computer.
- The theory of Software as a Service (SaaS), the area of Cloud Computing where Salesforce is located, is that you pay for a service and get a turnkey solution but are not supposed to dive into the complexity of running the IT stack. Running an IT stack includes the hardware (IaaS) and the development platform (PaaS), if any.
- In 1999, Salesforce launched their Customer Relationship Management (CRM) platform as the first SaaS solution built from scratch to achieve record growth.
- In the earliest days of the SaaS industry, it was assumed that subscription-based software would not be viable for enterprise business. And, in fact, in those days the enterprise typically chose end-to-end software suites to manage their complex organizations.
- Today, exponential growth of SaaS and continued improvements to functionality make it a valid option even for enterprise-level businesses. It's also much cheaper and easier to use. SaaS customers frequently cite cost savings as one of its primary benefits. You can find SaaS products for almost any business applications you can think of.
- IT as a Service (ITaaS) is the most recent and efficient delivery method in the decisive IT landscape. Integration as a Service (IaaS) is the budding and distinctive capability of clouds in fulfilling the business integration requirements. Increasingly business applications are deployed in clouds to reap the business and technical benefits.
- IaaS overcomes challenges by smartly utilizing the time-tested business-to-business (B2B) integration technology as the value-added bridge between SaaS solutions and in-house business applications.

- B2B systems are capable of driving this new on-demand integration model because they are traditionally employed to automate business processes between manufacturers and their trading partner.
- Social media is nowadays becoming the tool where companies are interacting with their customers. Instagram is becoming the new tool for customer relationship management. This makes the process much more simpler wherein there is no need for maintaining customer database.
  - Instead, it becomes even more exciting since the companies can directly interact with their customers with the help of social media. Most of the companies working on software as a service, are nowadays switching over to social media analytics and integration.
  - Hub and spoke architecture, helps for simple implementation of SaaS. It also avoids placing an excessive processing burden on the customer sides. The hub is installed at the SaaS provider's cloud center to do the heavy lifting such as reformatting files. A spoke unit at each user site typically acts as basic data transfer utility.
  - SaaS applications are mostly delivered through a web browser or a thin client terminal. The subscribers pay for SaaS services, which are priced on different usage parameters such as the number of transactions or the number of users accessing the app.

2.1.1 Challenges of SaaS Paradigm

- Lack of the following features prevents the massive adoption of clouds :
  1. Controllability
  2. Visibility & flexibility
  3. Security and Privacy
  4. High Performance and Availability
  5. Integration and Composition
  6. Standards
- Challenges :
  1. **Integration Conundrum** : Organization without a method of synchronizing data between multiple lines of businesses are at a serious disadvantage in terms of maintaining accurate data, forecasting, and automating key business processes. Real-time data and functionality sharing is an essential ingredient for clouds.
  2. Application programming interfaces are not proper and insufficient.

3. **Security for transmission of data :** Data integrity, confidentiality, quality and value have to be preserved as services and applications are interlinked and saddled to work together.

2.1.2 SaaS Integration Services

- Integration-as-a-Service delivers an integration solution that provides connectivity to backend systems, sources, files, and operational applications through the implementation of well-defined interfaces, web services, and calls between applications and data sources.
- SaaS integration involves connecting a SaaS application with another cloud-based app or an on-premise software via Application Programming Interfaces (APIs). Once connected, the app can request and share data freely with the other app or on-premise system.
- Cloud middleware will be made available as a service. Due to varying integration requirements and scenarios, there are a number of middleware technologies and products such as JMS-compliant message queues and integration backbones such as enterprise service bus, Enterprise Data Bus (EDB), complex event processing etc.
- Enterprise Service Bus (ESB) is an architectural pattern whereby a centralized software component performs integrations between applications. It performs transformations of data models, handles connectivity, performs message routing, converts communication protocols and potentially manages the composition of multiple requests.
- Events are coming up fast and there are Complex Event Processing (CEP) engines that receive a stream of diverse events from diverse sources, process them at real-time to extract and figure out the encapsulated knowledge, and accordingly select and activate one or more target applications thereby a kind of lighter connectivity and integration occurs between the initiating and the destination applications.
- Amazon's Simple Queue Service (SQS) provides a straightforward way for applications to exchange messages via queues in the cloud. SQS is a classic example for understanding what happens when a familiar on-premise service is recast as a cloud service.
- The constraining attributes of SaaS applications are :
  a) Dynamic nature of the SaaS interfaces that constantly change.
  b) Dynamic nature of the metadata native to a SaaS provider such as Salesforce.com.
  c) Managing assets that exist outside of the firewall.

d) Massive amounts of information that need to move between SaaS and on-premise systems daily and the need to maintain data quality and integrity.

- **Limited Access :** Access to cloud resources is more limited than local applications. Accessing local applications is quite simple and faster. Imbedding integration points in local as well as custom applications is easier.
- Once applications move to the cloud, custom applications must be designed to support integration because there is no longer that low level of access.
- **Dynamic resources :** Cloud resources are virtualized and service-oriented.
- **Performance :** Clouds support application scalability and resource elasticity.

2.2 SaaS Integration of Products and Platforms

2.2.1 Jitterbit

- Jitterbit cloud integration enables organizations to replicate, cleanse, and synchronize their cloud-based data seamlessly and securely with their on-premise enterprise applications and systems.
- Beside user-friendly interfaces and wizard tools, Jitterbit supports not only XML but also Web services. JitterBit focuses on data integration in the context of point-to-point application integration, ETL and SOA.
- Jitterbit supports SOA, event-driven architectures, and traditional data integration methods, and can easily scale to fit any cloud integration initiative.

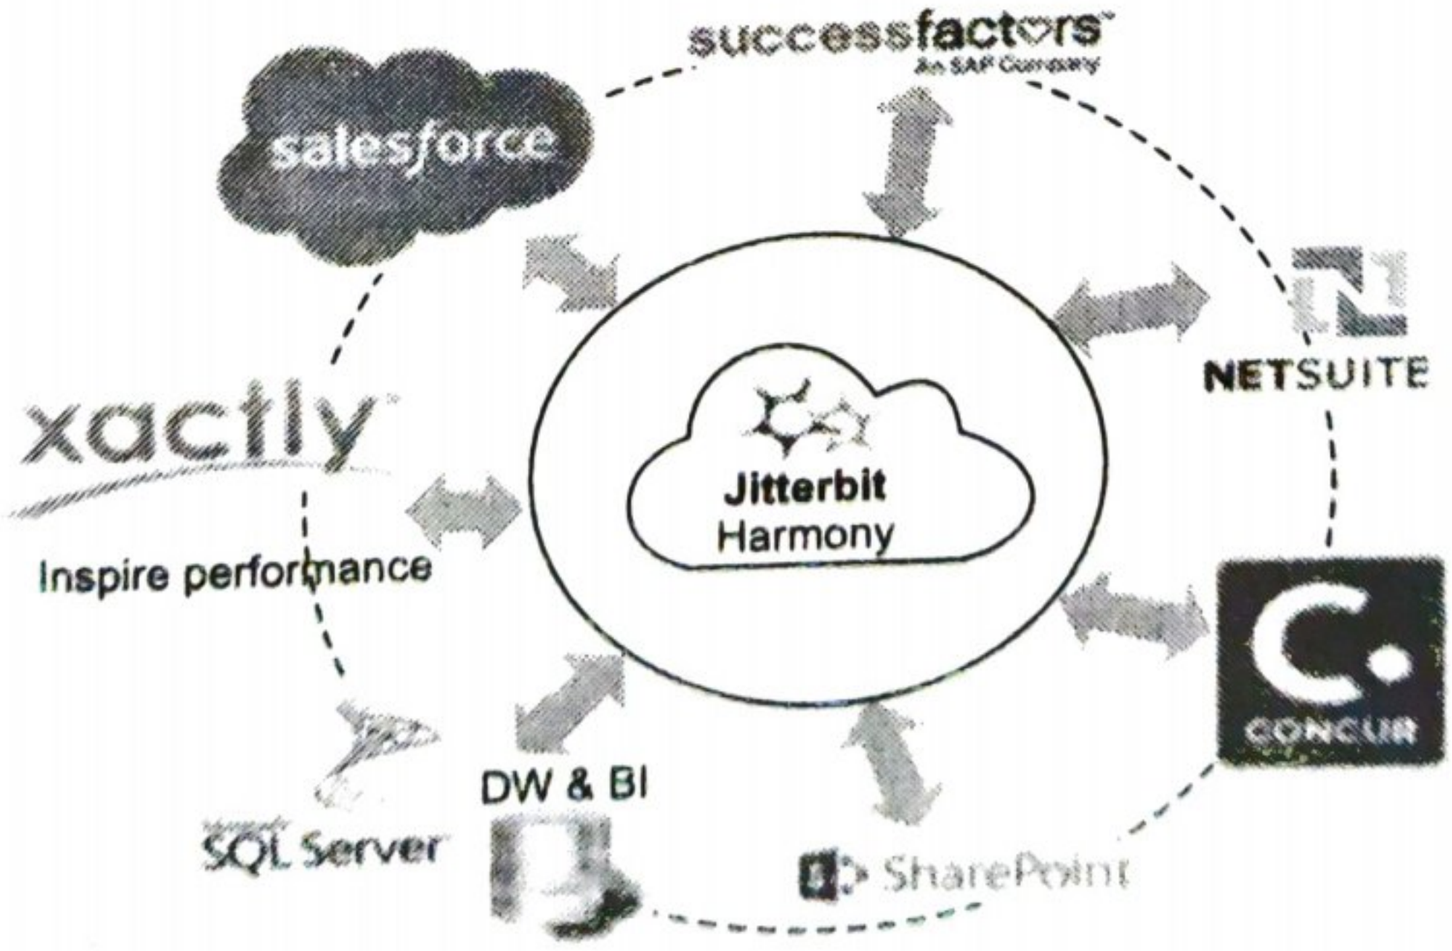

Fig. 2.2.1

- Jitterbit is a fully graphical integration solution that provides users a versatile platform and a suite of productivity tools to reduce the integration efforts sharply.
- Jitterbit can be used standalone or with existing EAI infrastructures, enabling users to create new projects or consume and modify existing ones offered by the open-source community or service provider.
- Jitterbit helps businesses make faster, more effective decisions by enabling them to unify and exploit data from all sources.
- Using the Jitterbit API integration platform companies can rapidly connect SaaS, on-premises and cloud applications and instantly infuse artificial intelligence into any business process.
- Jitterbit consists of two major components : Integration Server and Integration Environment.
  1. **Integration Environment :** An intuitive point-and-click graphical user interface that enables to quickly configure, test, deploy and manage integration projects on the Jitterbit server.
  2. **Integration Server :** A powerful and scalable run-time engine that processes all the integration operations, fully configurable and manageable from the Jitterbit application.

2.2.2 Boomi Software

- Boomi AtomSphere is an integration service that is completely on-demand and connects any combination of SaaS, PaaS, cloud, and on-premise applications without the burden of installing and maintaining software packages or appliances.
- Boomi offers the "pure SaaS" integration solution that enables to quickly develop and deploy connections between applications, regardless of the delivery model.
- Boomi Integration supports the B2B interactions with trading partners using industry-leading messaging standards like EDI X12, EDIFACT, and RosettaNet

2.2.3 Bungee Connect

- Bungee Connect enables cloud computing by offering an application development and deployment platform that enables highly interactive applications integrating multiple data sources and facilitating instant deployment.
- Bungee Connect reduces the efforts to integrate multiple web services into a single application. Applications built with Bungee Connect run at native speeds on each platform. An application built in Java with Bungee Connect will run natively on all targeted platforms.

- Bungee's unique "Connection" technology offers developers the ability to maintain a single set of application-specific logic where all multiple platform-specific assemblies are connected and configured through a simple design process.

2.2.4 OpSource Connect

- OpSource Connect also addresses the problems of SaaS integration by unifying different SaaS applications in the "cloud" as well as legacy applications running behind a corporate firewall.
- OpSource Connect is made up of key features including
  a) OpSource Services Bus
  b) OpSource Service Connectors
  c) OpSource Connect Certified Integrator Program
  d) OpSource Connect ServiceXchange
  e) OpSource Web Services Enablement Program

2.2.5 SnapLogic

- SnapLogic is a platform to integrate applications and data, allowing you to quickly connect apps and data sources. The company is also branching out into connecting and integrating data from IoT devices.
- SnapLogic offers a solution that provides flexibility for today's data integration challenges.
  1. Changing data sources : SaaS and on-premise applications, Web APIs, and RSS feeds
  2. Changing deployment options : On-premise, hosted, private and public cloud platforms
  3. Changing delivery needs : Databases, files, and data services
- **Advantages** : Includes many built-in integrations, and easy tracking of feeds into a system
- **Disadvantages** : Can take time to understand how the platform works; error handling not built-in.
- Pervasive data cloud is the first multi-tenant platform. Pervasive data cloud is a platform to deploy applications that are Scalable, Flexible, secure and Robust.

2.2.6 Online MQ

- It is an Internet-based queuing system. It support secure online messaging solution for sending and receiving messages over any network.

- It is a cloud messaging queuing service. In the integration space, messaging middleware as a service is the emerging trend.
- **Advantages :**
  1. Simple to use
  2. Maintenance free
  3. Support load balancing
  4. High availability

2.3 SaaS Integration Services

- The SaaS data integration allows you to distribute applications to those within your organization safely, securely, and without configuration, allowing you to easily share applications as well. It also allows you to work from that ambiguous cloud that everyone is raving about, without the public domain.

2.3.1 Informatica On-Demand

- Informatica Cloud is a data integration solution and platform that works like Software as a Service. It integrates cloud-based data with the data residing in on-premise databases and systems or between cloud applications.
- The informatica on-demand service is a subscription-based integration service that provides all the relevant features and functions, using an on-demand or an as-a-service delivery model.
- Informatica offers a variety of data integration software and services that are able to integrate with Salesforce. Informatica's cloud Salesforce connectivity lets organizations integrate on-premise systems, SaaS applications, and enterprise databases.
- For organizations using Service Cloud, Informatica helps unify your customer data and integrate it across Service Cloud, mobile, and on-premise environments.
- **Key benefits :**
  1. Rapid development and deployment with zero maintenance of the integration technology.
  2. Automatically upgraded and continuously enhanced by vendor.
  3. Proven SaaS integration solutions
  4. Proven data transfer and translation technology

2.3.2 Microsoft Internet Service Bus

- Azure is cloud operating system from Microsoft. Microsoft Azure Service Bus is a fully managed message broker with message queues and publish-subscribe topics. Service Bus is used to decouple applications and services from each other.
- Microsoft .NET services is a set of Microsoft-built and hosted cloud infrastructure services for building Internet-enabled applications.
- The .NET service bus acts as a perimeter network in the cloud, providing a single place to manage credentials of the client and services. The .NET service bus is the front end of the service; it encapsulates and isolates the service from malicious callers lurking on the internet and is responsible for repelling various attacks from denial-of-service to replay attacks, while obscuring the identity and true location of the actual service.
- The .NET Access Control Service is a hosted, secure, standards-based infrastructure for multiparty, federated authentication, rules-driven, and claims-based authorization.
- Fig. 2.3.1 shows how the relay service operates.

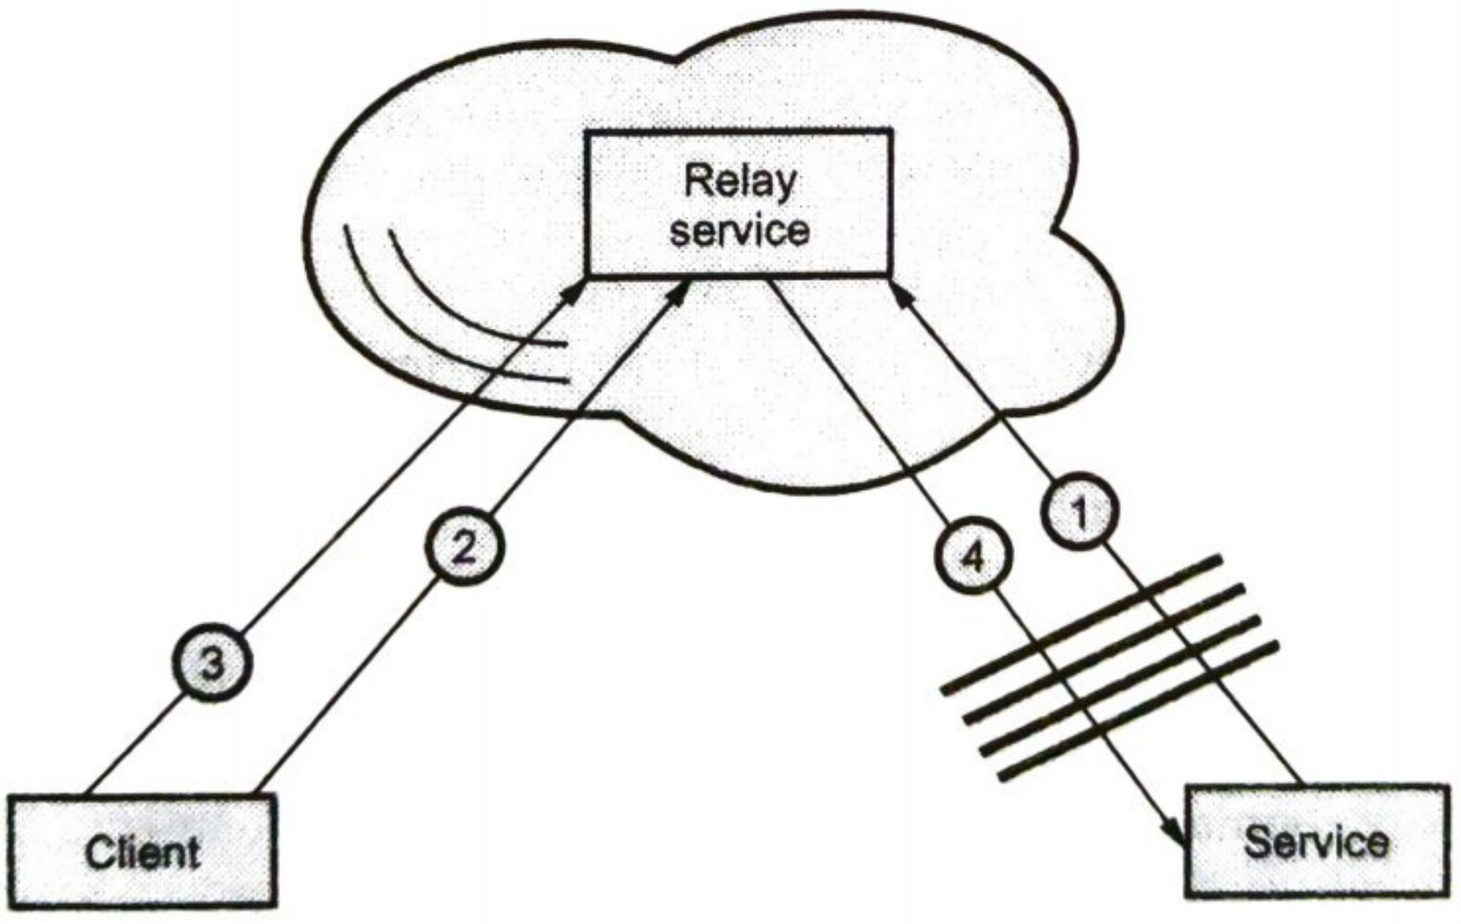

Fig. 2.3.1 Relay service operates

- First, both the service and the client must establish connections and authenticate against the relay service. At this point, the relay also records where the service is and how to best call back to it. When the client calls the relay service, the relay service forwards the call to the service.
- The .NET service bus acts as a perimeter network in the cloud, providing a single place to manage credentials of the client and services.

- The .NET service bus is the front end of the service; it encapsulates and isolates the service from malicious callers lurking on the internet and is responsible for repelling various attacks from denial-of-service to replay attacks, while obscuring the identity and true location of the actual service.
- The .NET service bus supports a WCF-friendly programming model by offering a set of dedicated bindings and behaviours.

2.4 Infrastructure as a Services

- Infrastructure-as-a-Service is a form of cloud computing that delivers fundamental compute, network, and storage resources to consumers on-demand, over the internet, and on a pay-as-you-go basis.
- In the IaaS model, the cloud provider owns and operates the hardware and software and also owns or leases the data center.
- IaaS includes virtual servers and cloud storage, cloud security, and access to data center resources

2.4.1 Background and Related Work

- IaaS is made up of a collection of physical and virtualized resources that provide consumers with the basic building blocks needed to run applications and workloads in the cloud.
- IaaS is typically understood as virtualized compute resources. Providers manage the hypervisors and end users can then programmatically provision virtual "instances" with desired amounts of compute and memory. The three primary types of cloud storage are block storage, file storage, and object storage.
- Virtualization can be defined as the abstraction of the four computing resources such as storage, processing power, memory, and network or I/O. Virtual machine's technology makes it very flexible and easy to manage resources in cloud computing environments.

2.4.2 Virtual Machines Provisioning and Manageability

- Virtual machine provisioning enables the cloud providers to make efficient utilization of available resources and make a good profit out of it. A cloud provider provisions their resources either statically or dynamically. In static Virtual machine provisioning the current demand of the user is not considered.
- Virtual Machine Lifecycle Management (VMLM) is a set of processes designed to help administrators oversee the implementation, delivery, operation, and maintenance of virtual machines (VMs) over the course of their existence.

- Fig. 2.4.1 shows virtual machine life cycle.

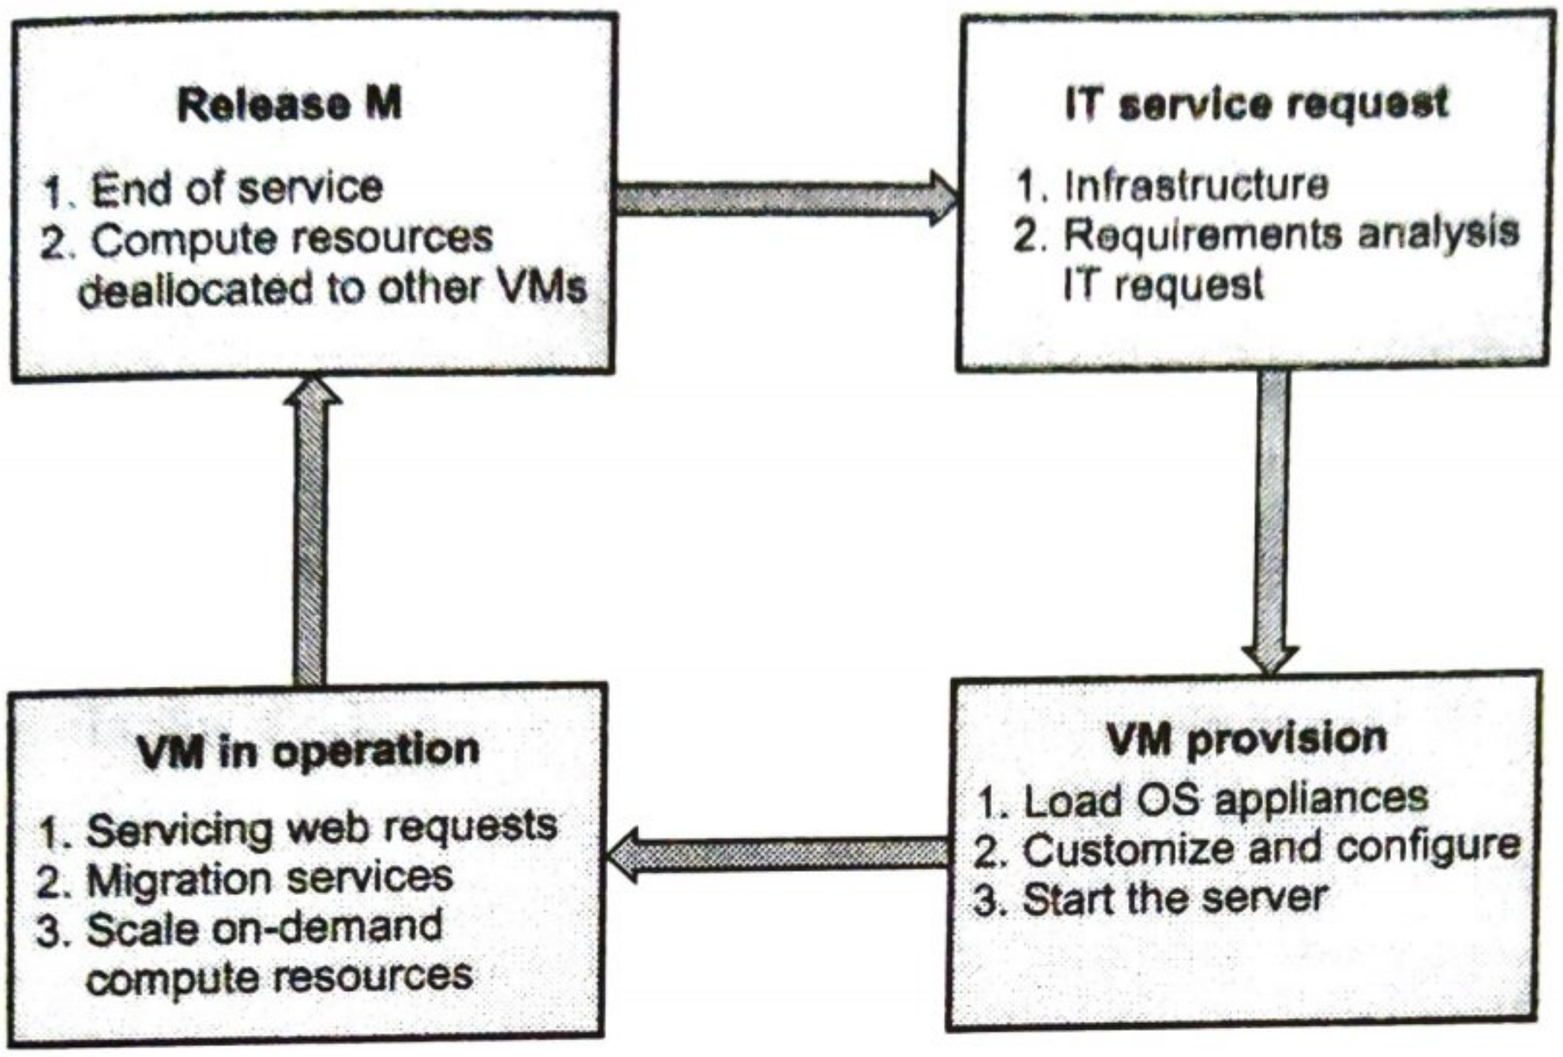

Fig. 2.4.1 Virtual machine life cycle

- Cycle starts by a request delivered to the organization information technology department, stating the requirement for creating a new server for a particular service.
- This request is being processed by the IT administration to start seeing the servers' resource pool, matching these resources with the requirements, and starting the provision of the needed virtual machine.
- Once it is provisioned and started, it is ready to provide the required service according to an service level agreement or a time period after which the virtual is being released; and free resources, in this case, won't be needed.
- **Steps to Provision VM :**
  1. Select a server from a pool of available servers along with the appropriate OS
  2. Load the required software
  3. Customize and configure the machine
  4. The virtual server is ready to start with its newly loaded software.

2.4.3 Virtual Machine Migration Services

- Migration service is the process of moving a virtual machine from one host server or storage location to another.

- Migration Time : Migration time refers to the total amount of time required to transfer a virtual machine at source to destination node without affecting its availability.
- It is used for load balancing and physical machine fault tolerant. It can also be used to reduce power consumption in cloud data centers.
- Virtual Machine Migration methods are divided into two types :
  1) **Hot (live) migration** - Virtual machine keeps running while migrating and does not lose its status.
  2) **Cold (non-live) migration** - The status of the VM loses and user can notice the service interruption
- Cold migration occurs when the VM is shut down. Live migration occurs while the VM is actually running.
- Migrations techniques are as follows :
  1. **Live migration and high availability :** Live migration is also called as hot or real-time migration. Live migration is the movement of a virtual machine from one physical host to another while being powered on. Live migration can also be used for load balancing in which work is shared among computers in order to optimize the utilization of available CPU resources.
- Fig. 2.4.2 shows live migration timelines.

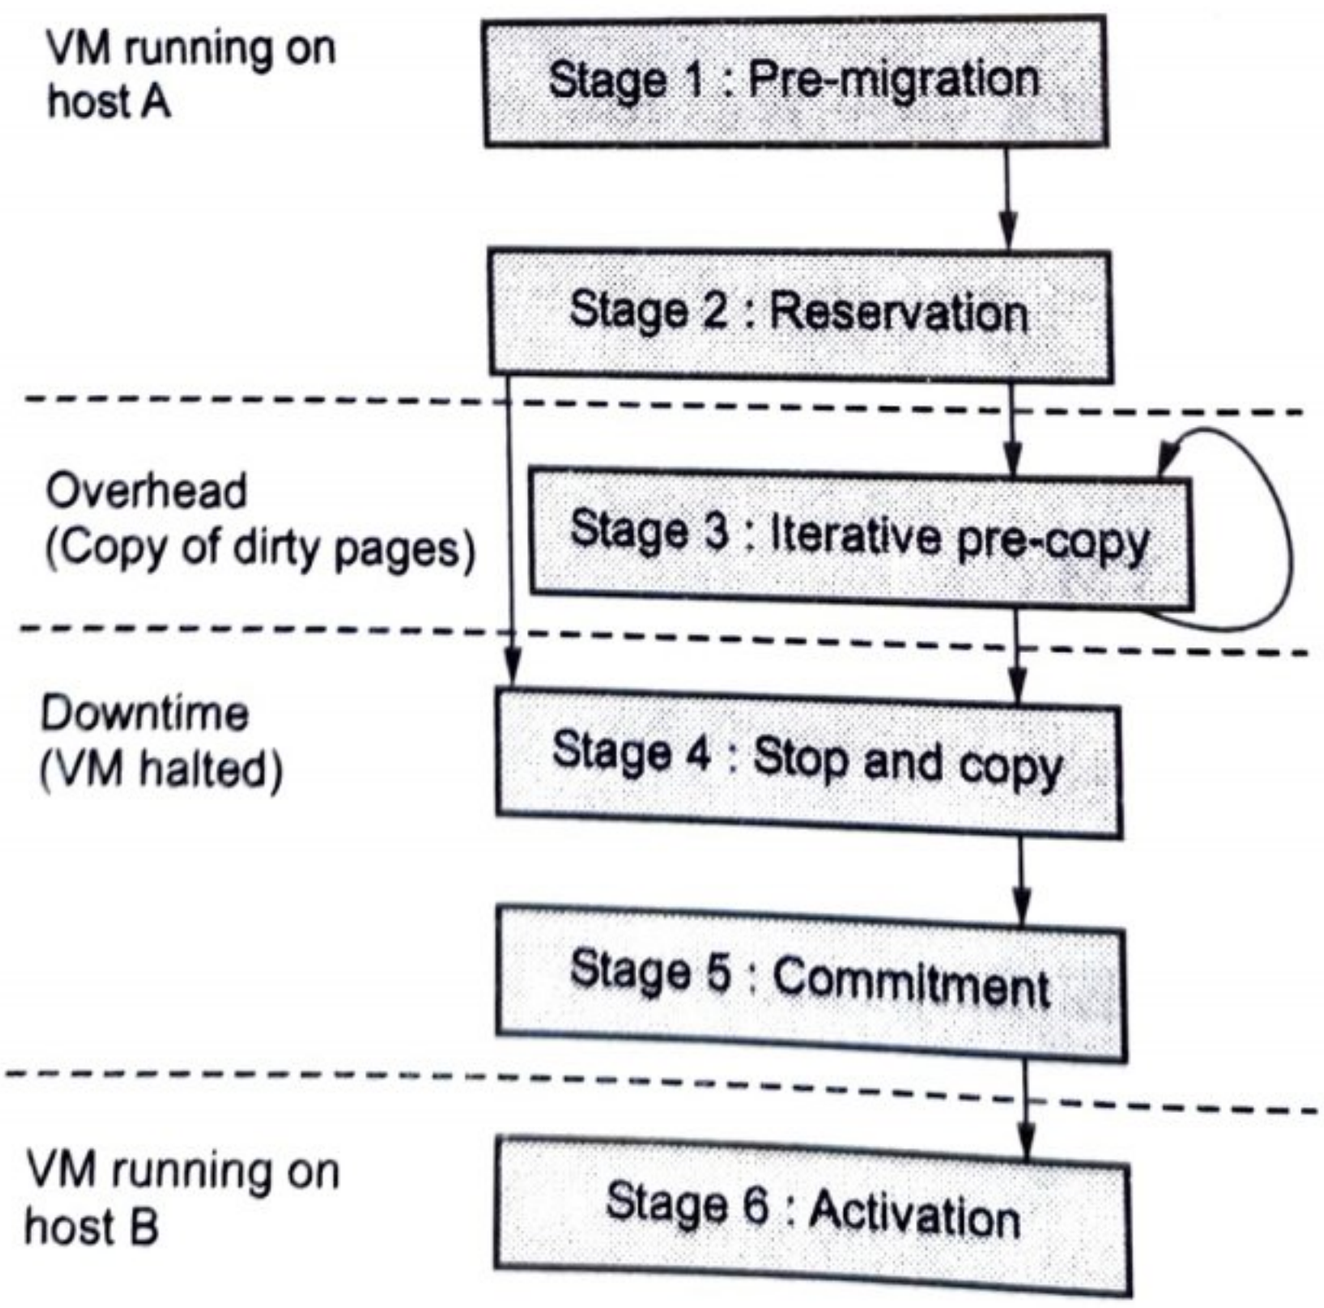

Fig. 2.4.2 Live migration timelines

Stage 1 : Pre-Migration stage : A target host will be preselected where the resources required to receive migration will be guaranteed.

Stage 2 : Reservation : A request is submitted to migrate a VM from Host-A to Host-B. If the request is not fulfilled, then VM will continue to run on Host-A.

Stage 3 : Repetitive Pre-Copy : During the first iteration, all memory pages are transferred from Host-A to Host-B. Subsequent iterations copy are only those pages dirtied during the previous transfer.

Stage 4 : Stop and Copy : In this phase, VM will be suspended on Host-A and redirect its network traffic to Host-B. CPU state and any remaining inconsistent memory pages are then transferred like a final sync. This process will reach a consistent suspended copy of the VM at both Host-A and Host-B. Host-A will remain primary and it will be resumed in case of failure at this stage.

Stage 5 : Commitment to the hosts : Host-B sends the signal to Host-A that it has successfully received a consistent VM OS image. Host-A acknowledges the signal and destroys the VM. Host-B becomes the primary host for migrated VM.

Stage 6 : Activation of VM : The migrated VM on Host-B is now activated. Post-migration code connects to the local resources and resumes the operation

2. Cold Migration

- Cold migration is the migration of a powered-off virtual machine. In cold migration, the VM is shut down at the source, moved to the destination and restarted at the destination host.
- In hot migration, the VM is suspended at the source, moved to and resumed at the destination. In live migration, the VM remains in its execution state while moving it from source to the destination host.
- The cold migration process is simple to implement.

VM migration, SLA and On-demand computing

- System attempts to allocate a maximum number of resources in a manner that ensures that all the service level agreements (SLAs) are maintained. Virtualization is considered as a core technology of cloud computing.
- Virtual machine instances allow cloud providers to utilize data center resources more efficiently. Moreover, by using dynamic VM consolidation using live migration, VMs can be placed according to their current resource requirements on the minimal number of physical nodes and consequently maintaining SLAs.
- A particular VM is consuming more than its fair share of resources at the expense of other VMs on the same host.

2.5 Platform As a Service

- Platform as a Service (PaaS) is an abstracted and integrated cloud-based computing environment that supports the development, running, and management of applications. Application components may exist in a cloud environment or may integrate with applications managed in private clouds or in data centers.
- The goal of the PaaS provider is to create an abstracted environment that supports an efficient, cost-effective, and repeatable process for the creation and deployment of high-quality applications. These applications are designed to be implemented in public or private cloud environments.
- Private clouds are virtual distributed systems that rely on a private infrastructure and provide internal users with dynamic provisioning of computing resources.

2.5.1 Integration of Private and Public Cloud

- Aneka is a software platform for developing cloud computing applications.
- Aneka is a platform and a framework for developing distributed applications on the cloud. It harnesses the spare CPU cycles of a heterogeneous network of desktop PCs and servers or data centers on demand.
- Aneka provides developers with a rich set of APIs for transparently exploiting such resources and expressing the business logic of applications by using the preferred programming abstractions.
- One of the key features of Aneka is the ability of providing different ways for expressing distributed applications by offering different programming models; execution services are mostly concerned with providing the middleware with an implementation for these models.
- Additional services such as persistence and security are transversal to the entire stack of services that are hosted by the Container.
- At the application level, a set of different components and tools are provided to :
  1) simplify the development of applications (SDK);
  2) porting existing applications to the Cloud; and
  3) monitoring and managing the Aneka Cloud.
- Aneka provides APIs and tools that enable applications to be virtualized over a heterogeneous network.
- The container is the building block of the middleware and represents the runtime environment for executing applications; it contains the core functionalities of the system and is built up from an extensible collection of services that allow administrators to customize the Aneka cloud.

- There are three classes of services that characterize the container :
  1. **Execution services** : They are responsible for scheduling and executing applications. Each of the programming models supported by Aneka defines specialized implementations of these services for managing the execution of a unit of work defined in the model.
  2. **Foundation services** : These are the core management services of the Aneka container. They are in charge of metering applications, allocating resources for execution, managing the collection of available nodes, and keeping the services registry updated.
  3. **Fabric services** : They constitute the lowest level of the services stack of Aneka and provide access to the resources managed by the cloud

2.5.2 Technologies and Tools for Cloud Computing

- Cloud computing covers the entire computing stack from hardware infrastructure to end-user software applications.
- Amazon is probably the major player for what concerns the infrastructure as-a-Service solutions in the case of public clouds. Amazon Web Services deliver a set of services that, when composed together, form a reliable, scalable, and economically accessible cloud.
- By using the GoGrid Web interface users can create their custom virtual images, deploy database and application servers, and mount new storage volumes for their applications.
- Both GoGrid and Amazon EC2 charge their customers on a pay-as-you-go basis, and resources are priced per hours of usage.
- Tera AppLogic lays at the foundation of many public clouds, it provides a grid operating system that includes workload distribution.

2.6 Multiple Choice Questions

Q.1 Azure is cloud operating system from _____.

a Google
b Amazon
c Flipkart
d Microsoft

Q.2 SQS is a classic example for understanding what happens when a familiar _____ service is recast as a cloud service.

a on-premise
b off-premise
c gateway
d None

Q.3 Force.com is a _____, enabling developers to create and deliver any kind of on-demand business application.

a software as a service
b infrastructure as a service
c platform as a service
d IT as a Service

Q.4 The .NET service bus is the _____ of the service.

a back end
b front end
c middleware
d All of these

Q.5 Informatica cloud is a _____ solution and platform that works like software as a service.

a data integration
b service integration
c storage integration
d service integration

Q.6 Online MQ is an Internet-based _____ system.

a messaging
b networking
c queuing
d All of these

Q.7 _____ are a set of agreements that are signed between the user and service providers.

a Service level agreement
b Service oriented architecture
c Service layer agreement
d Software level agreement

Q.8 Which of the following is NOT Cloud application features ?

a Multitenancy
b Elasticity
c Homogeneous cloud platform
d On-demand service

Q.9 Usually, when accessing the public or private cloud, the users require minimum _____,which is sometimes defined by the cloud providers ?

a Frequency
b bandwidth
c Internet
d all of these

Q.10 Which of the following network resources can be load balanced ?

a Connection through intelligent switches

b DNS

c Storage resources

d All of the mentioned

Q.11 How many types of virtual private server instances are partitioned in an IaaS stack ?

a One

b Two

c Three

d All of the mentioned

Answer Keys for Multiple Choice Questions :

Q.1	d	Q.2	a
Q.3	c	Q.4	b
Q.5	a	Q.6	c
Q.7	a	Q.8	c
Q.9	b	Q.10	d
Q.11	c		

❑❑❑

3 Abstraction and Virtualization

Syllabus

Introduction to Virtualization Technologies, Load Balancing and Virtualization, Understanding Hyper visors, Understanding Machine Imaging, Porting Applications, Virtual Machines Provisioning and Manageability Virtual Machine Migration Services, Virtual Machine Provisioning and Migration in Action, Provisioning in the Cloud Context, Virtualization of CPU, Memory, I/O Devices, Virtual Clusters and Resource management, Virtualization for Data Center Automation.

Contents

3.1 Introduction to Virtual Machine

GTU : Summer-18, Winter-17,18

- In a pure virtual machine architecture the operating system gives each process the illusion that it is the only process on the machine. The user writes an application as if only its code were running on the system.
- Each user interacts with the computer by typing commands to the virtual machine on a virtual system console and receiving results back from the machine as soon as they are computed.
- Each user directs the virtual machine to perform different commands. These commands are then executed on the physical machine in a multiprogramming environments.
- Virtualization is an abstraction layer that decouples the physical hardware from the operating system to deliver greater IT resource utilization and flexibility.
- It allows multiple virtual machines, with heterogeneous operating systems to run in isolation, side-by-side on the same physical machine.
- Fig. 3.1.1 shows virtual machine.
- Each virtual machine has its own set of virtual hardware (e.g., RAM, CPU, NIC, etc.) upon which an operating system and applications are loaded.
- The operating system creates the illusion of multiple processes, each executing on its own processor with its own (virtual) memory.
- The main components of virtual machine are the control program, conversational monitor system, remote spooling communication and interactive problem control system.
- The control program creates the environments in which virtual machines can executes. It also manages the real machines underlying the virtual machine environment.

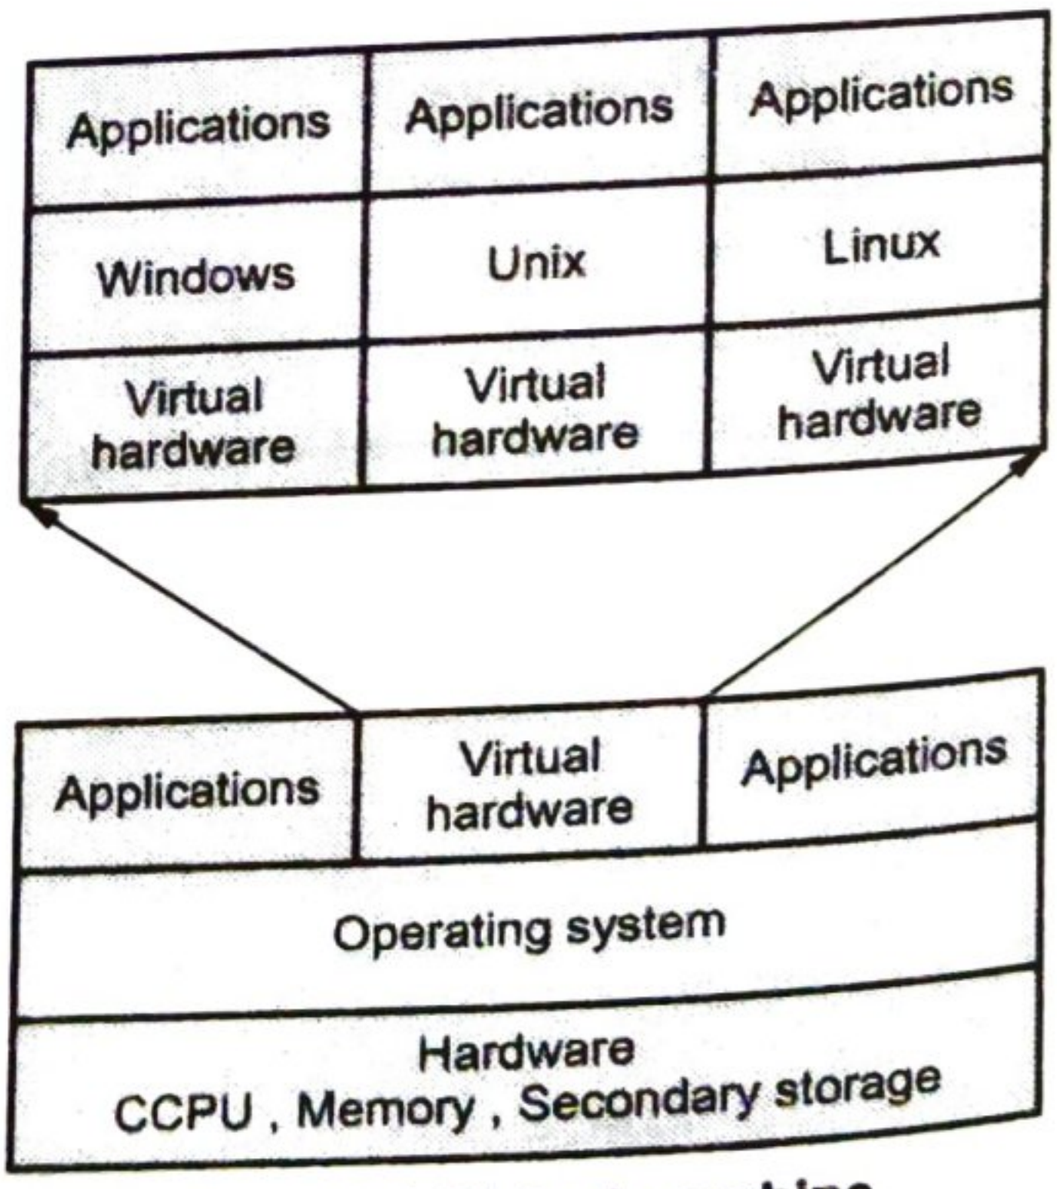

Fig. 3.1.1 Virtual machine

3.1.1 Virtualization Technologies

- Different types of virtualization that are characteristic of cloud computing:
  1. Access : A client can request access to a cloud service from any location.
  2. Application : A cloud has multiple application instances and directs requests to an instance based on conditions.
  3. CPU : Computers can be partitioned into a set of virtual machines with each machine being assigned a workload. Alternatively, systems can be virtualized through load-balancing technologies.
  4. Storage : Data is stored across storage devices and often replicated for redundancy.
- Features can be defined in software and hardware which enable flexibility as conforming to one or more of the following mobility patterns:
  a) P2V : Physical to Virtual
  b) V2V : Virtual to Virtual
  c) V2P : Virtual to Physical
  d) P2P : Physical to Physical
  e) D2C : Datacenter to Cloud
  f) C2C : Cloud to Cloud
  g) C2D : Cloud to Datacenter
  h) D2D : Datacenter to Datacenter

1. P2V : Physical to Virtual

- Physical to Virtual (P2V) is a term that refers to the migration of an operating system, application programs and data from a computer's main hard disk to a virtual machine. It is also called hardware virtualization.
- Data migrated in P2V includes an OS, applications, programs and data from a computer's main hard disk to a VM or a disk partition.
- The end result of a P2V migration is a VM with the same data, applications and system configurations as the physical server being virtualized.
- Modern operating systems and processors allow multiple processes to run simultaneously.
- A protection mechanism should exist in the processor so that all instructions from different processes will not access the hardware directly, this will lead to a system crash.
- All processors should have at least two modes : user mode and supervisor mode.

- Modes are used to control the access to the hardware directly. Instructions running in the supervisor mode are called privileged instructions and the others are unprivileged.
- Ex : VMware Workstation
- In a virtualized environment, it is more difficult to make OSes and applications run correctly because there are more layers in the machine stack.
- P2V is commonly used to accomplish server virtualization. P2V is also popular as a way for Mac users to run Windows applications.

2. V2V : Virtual to Virtual

- Virtual to Virtual (V2V) is a term that refers to the migration of an Operating System (OS), application programs and data from a virtual machine or disk partition to another virtual machine or disk partition.
- The target can be a single system or multiple systems. To streamline the operation, part or all of the migration can be carried out automatically by means of specialized programs known as migration tools.
- Virtual to Virtual (V2V) is a process of copying, migrating or replicating a Virtual Machine (VM) image, data or disk partition to another VM. It facilitates the migration of data or a machine instance between VMs and/or virtualization environments.

3. V2P : Virtual to Physical

- In Virtual to Physical (V2P) , the migration of an Operating System (OS), application programs and data from a virtual machine or disk partition to a computer's main hard disk. The target can be a single computer or multiple computers.
- V2P can be done manually by defining the target physical environment such as hard disk and then installing the OS, applications and data on it from the virtual environment.
- This can be complex and uncertain process, especially if the new environment contains substantially different hardware than the old. To streamline the operation, part or all of the migration can be carried out automatically by means of specialized programs known as migration tools.
- V2P can be used to restore the hard disk contents of a failed computer or network server from a backup storage medium such as a tape drive.

4. C2D : Cloud to Datacenter

- Data centers require extensive network hardware in order to enable multiple levels of connectivity.

- For networking infrastructure, the data center is broken down into five network subsystems: Carrier and External Networks Interconnection, Web-Tier Load Balancing and Acceleration.
- Carrier and External Networks Interconnection : It consists of backbone routers, firewall and VPN gateways. Backbone routers provide routing between external WAN connections and data center LAN.
- Web-Tier Load Balancing and Acceleration : It contains web acceleration device such as XML pre-processors, encryption/decryption appliances and layer 7 switching devices that perform content-aware routing.

3.1.2 Load Balancing

- Load balancing can be defined as the process of task distribution among multiple computers, processes, disk or other resources in order to get optimal resource utilization and to reduce the computation time.
- Fig. 3.1.2 shows load balancing in cloud computing.

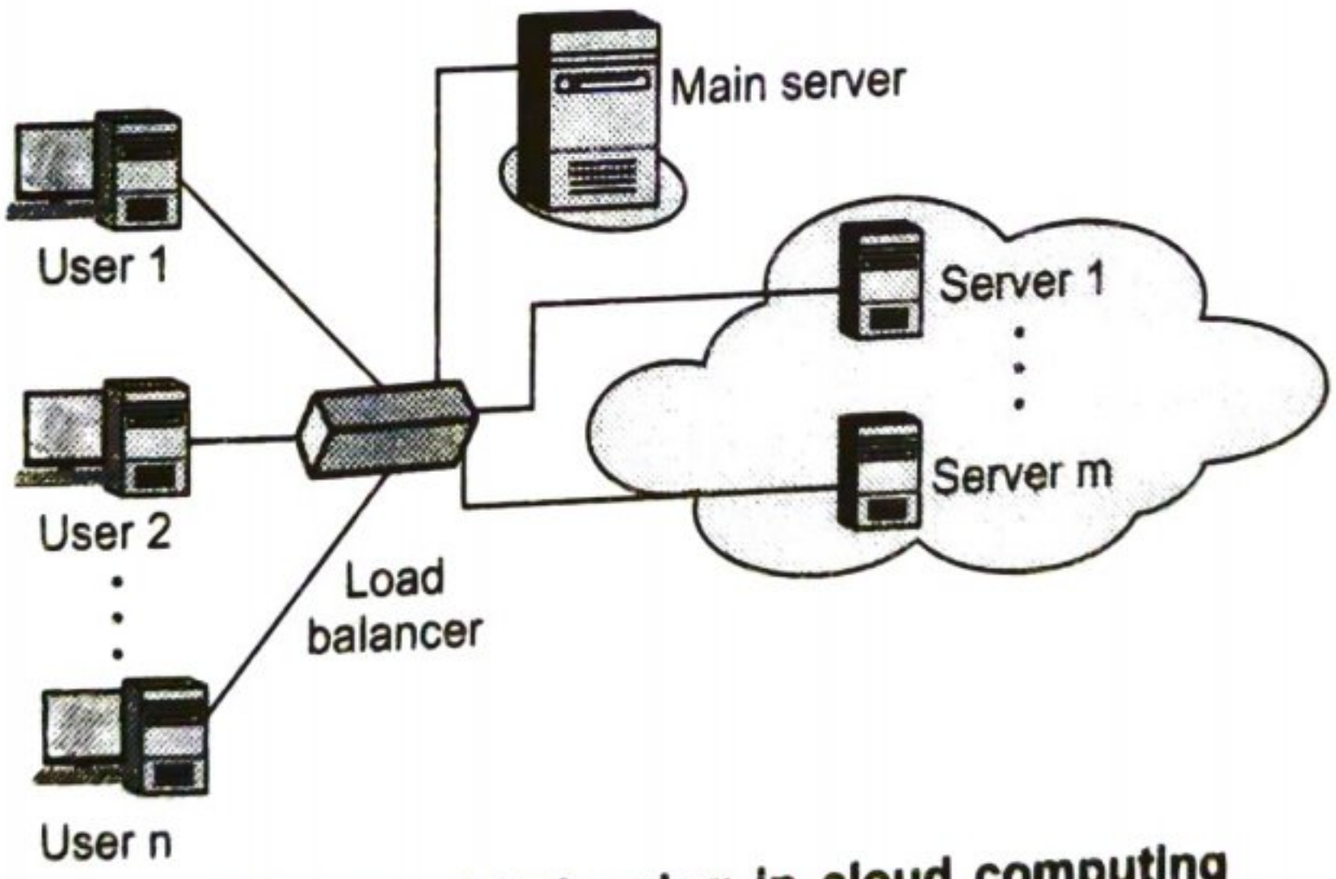

Fig. 3.1.2 Load balancing in cloud computing

- Load balancing is an important means to achieve effective resource sharing and utilization.
- Cloud load balancing is the process of distributing workloads and computing resources in a cloud computing environment. Load balancing allows enterprises to manage application or workload demands by allocating resources among multiple computers, networks or servers.
- The technology used to distribute service requests to resources is referred to as load balancing. Load balancing can be implemented in hardware.
- Google, Yahoo!, Amazon, and Microsoft experience millions of user hits per day. Across the web, sites experience a wide range of network traffic requirements.

- To handle such web requests, the sites use a technique known as load balancing, to share the requests across multiple servers.
- Load balancing uses a server to route traffic to multiple servers which, in turn, share the workload
- In general, load balancing algorithms can be divided into following three types :
  1. Centralized approach : In this approach, a single node is responsible for managing the distribution within the whole system.
  2. Distributed approach : In this approach, each node independently builds its own load vector by collecting the load information of other nodes. Decisions are made locally using local load vectors. This approach is more suitable for widely distributed systems such as cloud computing.
  3. Mixed approach : A combination between the two approaches to take advantage of each approach.
- Load balancing is an optimization technique; it can be used to increase utilization and throughput, lower latency, reduce response time and avoid system overload.
- The following network resources can be load balanced :
  a. Network interfaces and services such as DNS, FTP and HTTP
  b. Connections through intelligent switches
  c. Processing through computer system assignment
  d. Storage resources
  e. Access to application instances
- Without load balancing, cloud computing would very difficult to manage. Load balancing provides the necessary redundancy to make an intrinsically unreliable system reliable through managed redirection. It also provides fault tolerance when coupled with a failover mechanism.

Metrics For Load Balancing In Clouds :

- Various metrics considered in existing load balancing techniques in cloud computing are discussed below :
  1. Throughput is used to calculate the number of tasks whose execution has been completed. It should be high to improve the performance of the system.
  2. Overhead associated determines the amount of overhead involved while implementing a load-balancing algorithm. It is composed of overhead due to movement of tasks, interprocessor and inter-process communication. This should be minimized so that a load balancing technique can work efficiently.
  3. Fault tolerance is the ability of an algorithm to perform uniform load balancing in spite of arbitrary node or link failure. The load balancing should be a good fault-tolerant technique.

4. Migration time is the time to migrate the jobs or resources from one node to other. It should be minimized in order to enhance the performance of the system.
5. Response time is the amount of time taken to respond by a particular load balancing algorithm in a distributed system. This parameter should be minimized.
6. Resource utilization is used to check the utilization of resources. It should be optimized for an efficient load balancing.
7. Scalability is the ability of an algorithm to perform load balancing for a system with any finite number of nodes. This metric should be improved.
8. Performance is used to check the efficiency of the system. This has to be improved at a reasonable cost, e.g., reduce task response time while keeping acceptable delays

- In cloud computing, load balancing is required to distribute the dynamic local workload evenly across all the nodes. It helps to achieve a high user satisfaction and resource utilization ratio by ensuring an efficient and fair allocation of every computing resource.
- The various benefits of load balancing are as follows :
  1. Increase resource utilization
  2. Maximize throughput
  3. Lower latency
  4. Reduce response time
  5. Avoid system overload
  6. Increased reliability

University Questions

1. *What is virtual machine ? Explain virtual machine types.* **GTU : Winter-17, Marks 4**
2. *Explain in brief P2V, V2V, V2P, P2P, D2C, C2C and C2D virtual machine conversions in VMM.* **GTU : Winter-17, Marks 7**
3. *Define load balancing. What is need of load balancing in cloud computing ?* **GTU : Summer-18, Winter-18, Marks 4**

3.2 Understanding Hypervisors

GTU : Summer-18

- In computing, a hypervisor is a virtualization platform that allows multiple operating systems to run on a host computer at the same time. The term usually refers to an implementation using full virtualization.

- Hypervisors are currently classified in two types :

1. Type 1 hypervisor is software that runs directly on a given hardware platform. A "guest" operating system thus runs at the second level above the hardware.
   - Type 1 VMs have no host operating system because they are installed on a bare system. An operating system running on a Type 1 VM is a full virtualization because it is a complete simulation of the hardware that it is running on.
2. Type 2 hypervisor is software that runs within an operating system environment. A "guest" operating system thus runs at the third level above the hardware.
   - Bochs and QEMU are PC emulators that allow operating systems such as Windows or Linux to be run in the user-space of a Linux operating system. VMware is a popular commercial full-virtualization solution that can virtualizes unmodified operating systems.
   - Xen is an open source para-virtualization solution that requires modifications to the guest operating systems but achieves near native performance by collaborating with the hypervisor.
   - Microsoft Virtual PC is a para-virtualization virtual machine approach. User-Mode Linux (UML) is another para-virtualization solution that is open source.
   - Each guest operating system executes as a process of the host operating system. Cooperative Linux, is a virtualization solution that allows two operating systems to cooperatively share the underlying hardware.
   - Linux-V server is an operating system-level virtualization solution for GNU/Linux systems with secure isolation of independent guest servers.
   - The Linux KVM is virtualization technology that has been integrated into the mainline Linux kernel. Runs as a single kernel loadable module, a Linux kernel running on virtualization-capable hardware is able to act as a hypervisor and support unmodified Linux and Windows guest operating systems.

University Question

1. What are hypervisors? List it's importance. **GTU : Summer-18, Marks 3**

3.2.1 Xen Architecture

- Xen is a type 1 hypervisor that creates logical pools of system resources so that many virtual machines can share the same physical resources.

- Xen is a hypervisor that runs directly on the system hardware. It inserts a virtualization layer between the system hardware and the virtual machines, turning the system hardware into a pool of logical computing resources that Xen can dynamically allocate to any guest operating system.
- Fig. 3.2.1 shows Xen architecture.

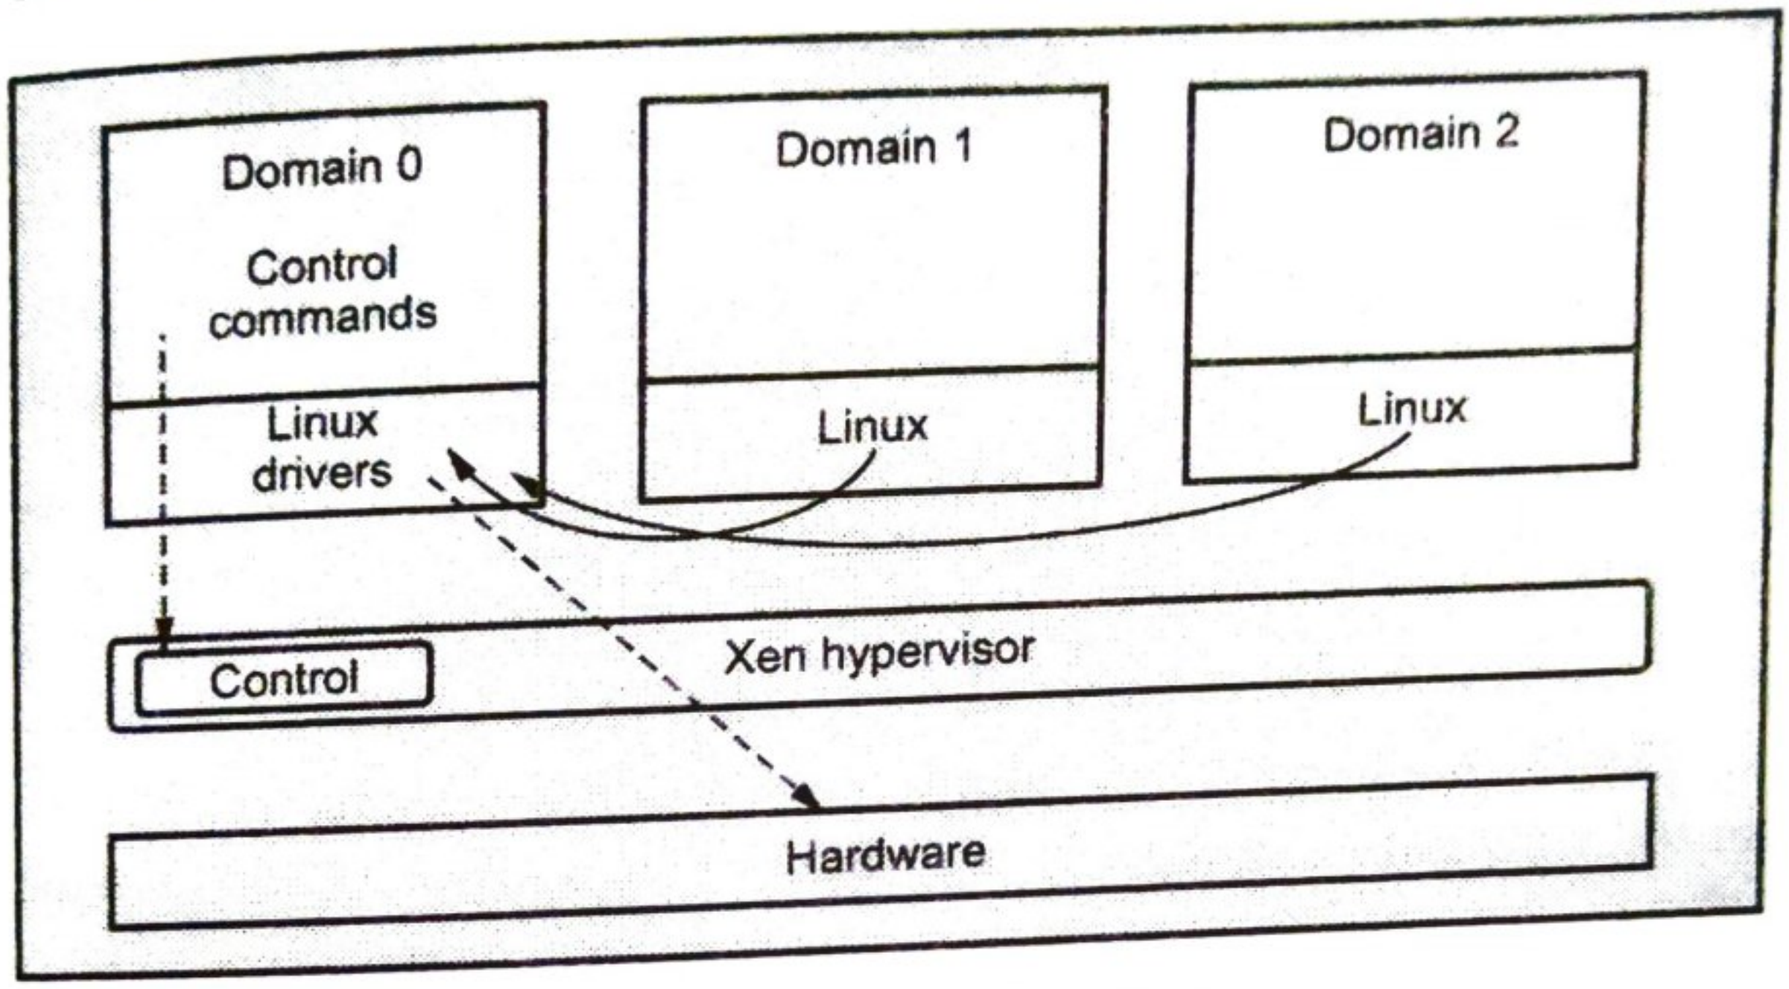

Fig. 3.2.1 Xen architecture

- The operating systems running in virtual machines interact with the virtual resources as if they were physical resources. Xen provides a virtual environment located between the hardware and the OS.
- Xen doesn't include any device drivers; it provides a mechanism by which a guest-OS can have direct access to the physical devices.
- The core components of Xen are the hypervisor, kernel and applications. Many guest operating systems can run on the top of the hypervisor; but it should be noted that one of these guest OS controls the others.
- This guest OS with the control ability is called Domain 0, the others are called Domain U. Domain 0 is first loaded when the system boots and can access the hardware directly and manage devices by allocating the hardware resources for the guest domains (Domain U).
- The Control Domain (or Domain 0) is a specialized Virtual Machine that has special privileges like the capability to access the hardware directly, handles all access to the system's I/O functions and interacts with the other Virtual Machines.
- It also exposes a control interface to the outside world, through which the system is controlled. The Xen Project Hypervisor is not usable without Domain 0, which is the first VM started by the system.

3.3 Machine Migration Services

GTU : Summer-17

- Machine imaging is a process that is used to provide system portability, and provision and deploy systems in the cloud through capturing the state of systems using a system image.
- A system image makes a copy or a clone of the entire computer system inside a single file. The image is made by using a program called system imaging program and can be used later to restore a system image.
- For example : Amazon Machine Image (AMI) is a system image that is used in the cloud computing. The Amazon Web Services uses AMI to store copies of a virtual machine.
- An AMI is a file system image that contains an operating system, all device drivers and any applications and state information that the working virtual machine would have.
- The AMI files are encrypted and compressed for security purpose and stored in Amazon S3 (Simple Storage System) buckets as a set of 10 MB chunks.
- Machine imaging is mostly run on virtualization platform due to this it is also called as virtual appliances and running virtual machines are called instances.
- The AMI file system is not a standard bit-for-bit image of a system that is common to many disk imaging programs. AMI omits the kernel image and stores a pointer to a particular kernel that is part of the AWS kernel library.
- Among the choices are Red Hat Linux, Ubuntu, Microsoft Windows, Solaris and others. Files in AMI are compressed and encrypted and an XML file is written that describes the AMI archive.
- Machine images are sometimes referred to as "virtual appliances", systems that are meant to run on virtualization platforms.

University Question

1. *How machine imaging help to achieve the goal of cloud computing ?*

GTU : Summer-17, Marks 7

3.4 Exploring Virtualization

GTU : Summer-17, Winer-18

- Virtualization is a broad term that refers to the abstraction of resources across many aspects of computing. For our purposes : One physical machine to support multiple virtual machines that run in parallel.
- Virtualization is a frame work or methodology of dividing the resources of computer into multiple execution environments.

- Virtualization is an abstraction layer that decouples the physical hardware from the operating system to deliver greater IT resource utilization and flexibility.
- It allows multiple virtual machines, with heterogeneous operating systems to run in isolation, side-by-side on the same physical machine.
- Fig. 3.4.1 shows before and after virtualization.
- Virtualization means running multiple machines on a single hardware. The "Real" hardware invisible to operating system. OS only sees an abstracted out picture. Only Virtual Machine Monitor (VMM) talks to hardware.
- It is "a technique for hiding the physical characteristics of computing resources from the way in which other systems, applications, or end users interact with those resources.
- This includes making a single physical resource appear to function as multiple logical resources; or it can include making multiple physical resources appear as a single logical resource."

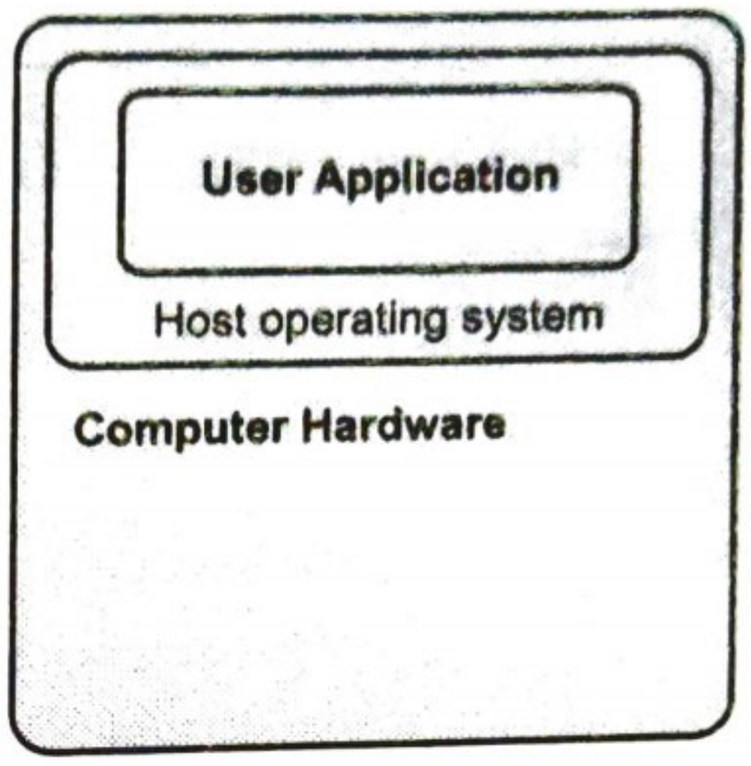

(a) : Before virtualization

(b) After virtualization

Fig. 3.4.1

- It is divided into two main categories :
  1. Platform virtualization involves the simulation of virtual machines.
  2. Resource virtualization involves the simulation of combined, fragmented or simplified resources.
- Fig. 3.4.2 shows taxonomy of virtualization. (See Fig. 3.4.2 on next page)
- Virtualization is mainly used to emulate execution environment, storage and network. Execution environment classified into two types : process level and system level.

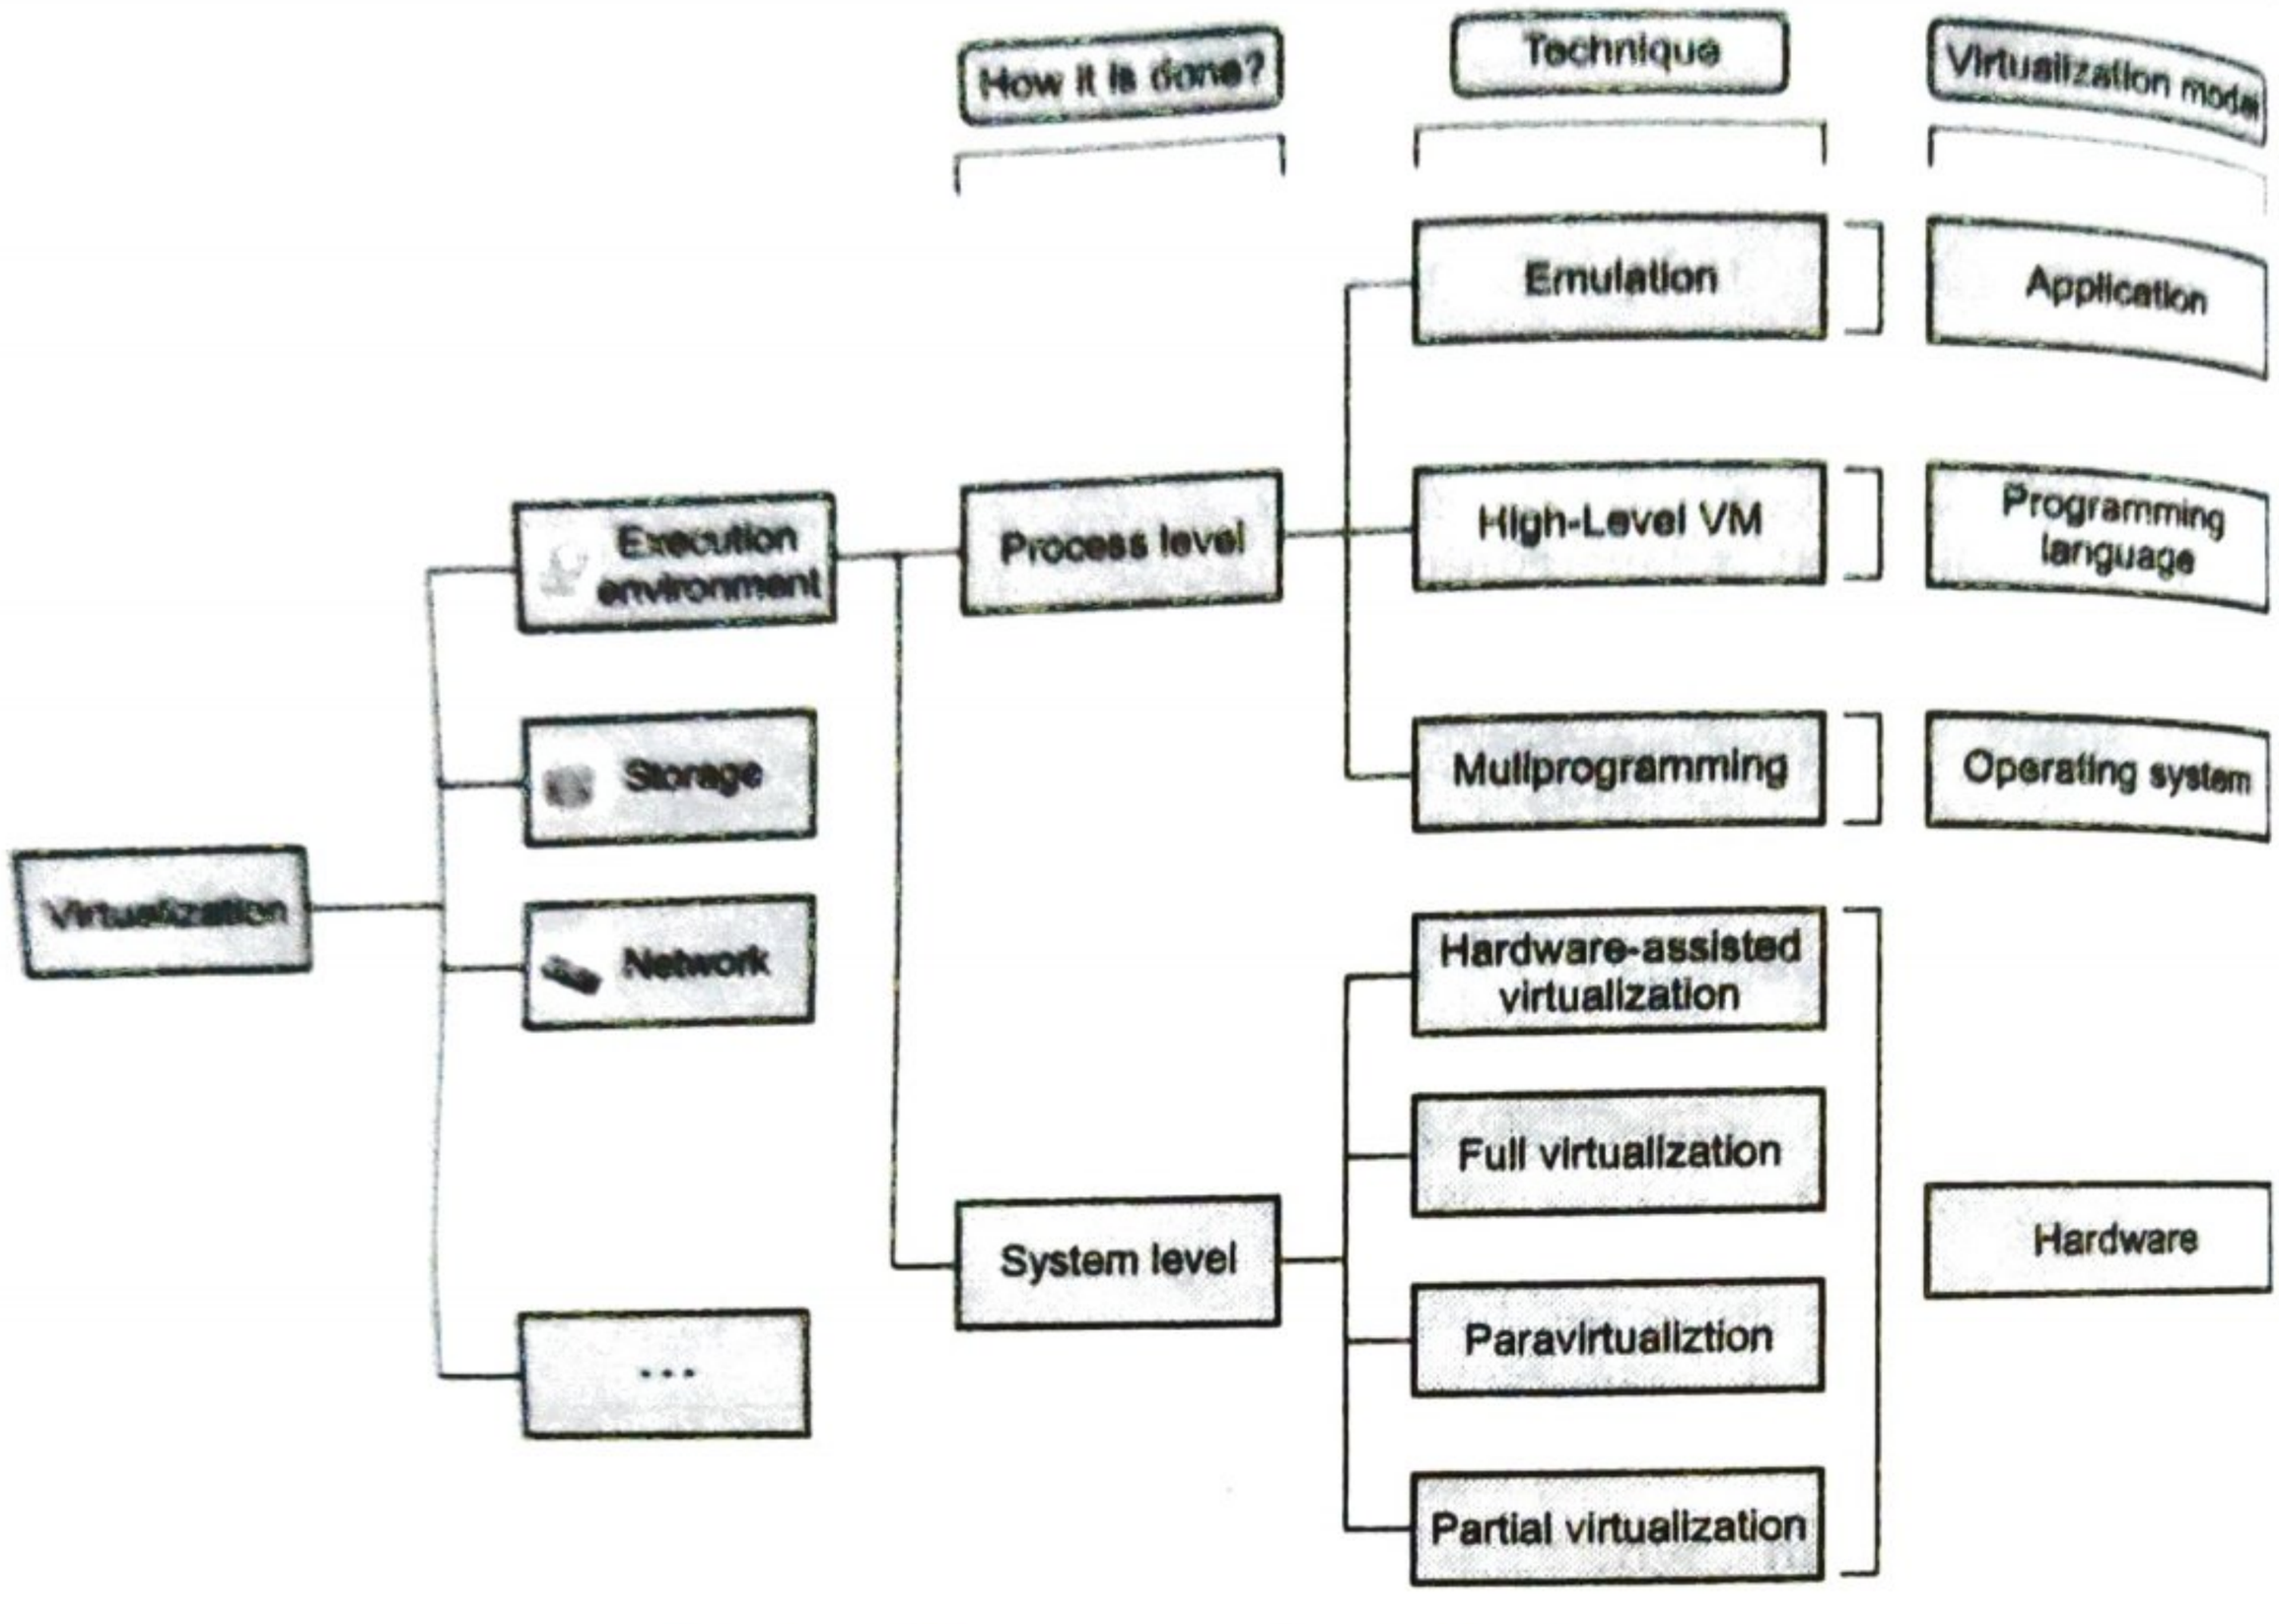

Fig. 3.4.2 Taxonomy of virtualization

- Process level is implemented on top of an existing operating system.
- System level is implemented directly on hardware and do not or minimum requirement of existing operating system.

3.4.1 Platform Virtualization

- The creation of a virtual machine using a combination of hardware and software is referred to as platform virtualization.
- Platform virtualization is performed on a given hardware platform by "host" software, which creates a simulated computer environment for its "guest" software.
- The "guest" software, which is often itself a complete operating system, runs just as if it were installed on a stand-alone hardware platform. Typically, many such virtual machines are simulated on a given physical machine.
- For the "guest" system to function, the simulation must be robust enough to support all the guest system's external interfaces, which may include hardware drivers.

3.4.2 Resource Virtualization

- The basic concept of platform virtualization was later extended to the virtualization of specific system resources, such as storage volumes, name spaces and network resources.
- Resource aggregation, spanning or concatenation combines individual components into larger resources or resource pools. For example : RAID and volume managers combine many disks into one large logical disk.
- Virtual Private Network (VPN), Network Address Translation (NAT) and similar networking technologies create a virtualized network namespace within or across network subnets. Multiprocessor and multi-core computer systems often present what appears as a single, fast processor
- Virtualization means running multiple machines on a single hardware. The "Real" hardware invisible to operating system. OS only sees an abstracted out picture. Only Virtual Machine Monitor (VMM) talks to hardware.
- It is "a technique for hiding the physical characteristics of computing resources from the way in which other systems, applications or end users interact with those resources.
- This includes making a single physical resource appear to function as multiple logical resources; or it can include making multiple physical resources appear as a single logical resource."

3.4.3 Pros and Cons of Virtualization

a) Pros

1. Data center and energy-efficiency savings : As companies reduce the size of their hardware and server footprint, they lower their energy consumption
2. Operational expenditure savings : Once servers are virtualized, your IT staff can greatly reduce the ongoing administration and management of manual work.
3. Reduced costs : It reduced cost of IT infrastructure.
4. Data does not leak across virtual machine.
5. Virtual machine is completed isolated from host machine and other virtual machine.
6. Simplifies resource management by pooling and sharing resources.
7. Significantly reduce downtime.
8. Improved performance of IT resources.

b) Cons

1. Not all hardware or software can be virtualized.
2. Not all servers are applications are specifically designed to be virtualization-friendly.

3.4.4 Difference between Virtualization and Cloud Computing

Sr. No.	Virtualization	Cloud Computing
1.	Virtualization is the process of creating a virtual environment on an existing server to run your desired program, without interfering with any of the other services provided by the server or host platform to other users.	Cloud computing means storing and accessing data and programs over the Internet instead of your computer's hard drive.
2.	Location of virtual machine is on a specific host.	Location of virtual machine is on any host.
3.	Instance storage is persistent.	Instance storage is shortly lived.
4.	Virtualization uses customizable VM resource like CPU and RAM.	Cloud computing uses standard VM resource like CPU and RAM
5.	Recovery from failures: attempt to recover failed VM.	Recovery from failures : Discard instance spin up new one.

3.4.5 Implementation Levels of Virtualization

- Virtualization is implemented at various levels :
  1. Instruction Set Architecture Level
  2. Hardware Abstraction Level
  3. Operating System Level
  4. Library Support Level
  5. User Application level

1. Instruction Set Architecture Level

- The definition of the storage resources and the instructions that manipulate data are documented in what is referred to as Instruction Set Architecture (ISA)
- ISA view of a machine corresponds to the machine and assembly language levels. For example, MIPS binary code can run on an x86-based host machine with the help of ISA emulation.

- Instruction set emulation leads to virtual ISAs created on any hardware machine. The basic emulation method is through code interpretation. An interpreter program interprets the source instructions to target instructions one by one.
- The key to virtualize a CPU lies in the execution of the guest instructions, including both system-level and user-level instructions Virtualizing a CPU can be achieved in one of two ways :

1. Emulation : the only processor virtualization mechanism available when the ISA of the guest is different from the ISA of the host.
2. Direct native execution : Possible only if the ISA of the host is identical to the ISA of the guest .
   - **Emulation** is the process of implementing the interface and functionality of one system (or subsystem) on a system (or subsystem) having different interface and functionality.
   - In other words, emulation allows a machine implementing one ISA (the target), to reproduce the behavior of a software compiled for another ISA (the source). Emulation can be carried out using :
     1. Interpretation 2. Binary translation

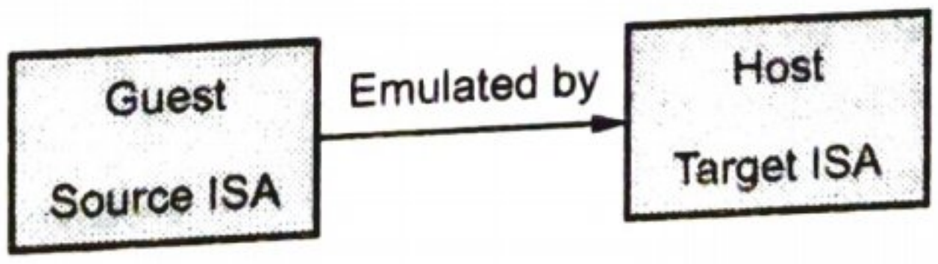

Fig. 3.4.3

2. Hardware Abstraction Level

- This type of virtualization is performed right on top of the bare hardware. On the one hand, this approach generates a virtual hardware environment for a VM. On the other hand, the process manages the underlying hardware through virtualization.
- The idea is to virtualizes a computer's resources, such as its processors, memory, and I/O devices. The intention is to upgrade the hardware utilization rate by multiple users concurrently.
- The Xen hypervisor has been applied to virtualize x86-based machines to run Linux or other guest OS applications.

3.4.6 Operating System Level Virtualization

- Operating-system-level virtualization is a server-virtualization method where the kernel of an operating system allows for multiple isolated user-space instances, instead of just one. Such instances, which are sometimes called containers and software containers.

- This refers to an abstraction layer between traditional OS and user applications.
- This type of virtualization creates isolated containers on a single physical server and the OS instances to utilize the hard-ware and software in data centers.
- Containers behave like real servers. With containers you can create a portable, consistent operating environment for development, testing, and deployment.
- This virtualization creates virtual hosting environments to allocates hardware resources among a large number of mutually distrusting users.
- Operating-system-level virtualization usually imposes little to no overhead, because programs in virtual partitions use the operating system's normal system call interface and do not need to be subjected to emulation or be run in an intermediate virtual machine.
- Operating system-level virtualization is not as flexible as other virtualization approaches since it cannot host a guest operating system different from the host one, or a different guest kernel.
- Instead of trying to run an entire guest OS, container virtualization isolates the guests, but doesn't try to virtualize the hardware. Instead, you have containers for each virtual environment.
- With container-based technologies, you'll need a patched kernel and user tools to run the virtual environments. The kernel provides process isolation and performs resource management.

Why operating system level virtualization is required ?

- Operating system level virtualization provides feasible solution for hardware level virtualization issue. It inserts a virtualization layer inside an operating system to partition a machine's physical resources.
- It enables multiple isolated VMs within a single operating system kernel. This kind of VM is often called a virtual execution environment (VE), Virtual Private System (VPS), or simply container.
- From the user's point of view, virtual execution environment look like real servers.
- This means a virtual execution environment has its own set of processes, file system, user accounts, network interfaces with IP addresses, routing tables, firewall rules etc.
- Although VEs can be customized for different people, they share the same operating system kernel. Therefore, OS-level virtualization is also called single-OS image virtualization.

Challenges to cloud computing in OS level virtualization ?

- Cloud computing is transforming the computing landscape by shifting the hardware and staffing costs of managing a computational center to third parties.
- Cloud computing has at least two challenges :
  1. The ability to use a variable number of physical machines and virtual machine instances depending on the needs of a problem. For example, a task may need only a single CPU during some phases of execution but may need hundreds of CPUs at other times.
  2. It is related to slow operation of instantiating new virtual machine. Currently, new virtual machines originate either as fresh boots or as replicates of a template VM, unaware of the current application state. Therefore, to better support cloud computing, a large amount of research and development should be done.

Advantages of OS virtualization :

1. OS virtualization provide least overhead among all types of virtualization solution.
2. They offer highest performance and highest density of virtual environment.
3. Low resource requirements.
4. High Scalability.

Disadvantage of OS virtualization :

1. They support only one operating system as base and guest OS in a single server.
2. It supports library level virtualization.

3.4.7 Library-Level Virtualization

- Library-level virtualization is also known as user-level Application Binary Interface (ABI).
- This type of virtualization can create execution environments for running alien programs on a platform rather than creating a VM to run the entire operating system.
- It is done by API call interception and remapping.
- Virtualization with library interfaces is possible by controlling the communication link between applications and the rest of a system through API hooks.
- Example : Wine, WAB, LxRun , Visual MainWin
- Advantage : It has very low implementation effort
- Shortcoming and limitation: poor application flexibility and isolation

3.4.8 VMM Design Requirements and Providers

- Hardware-level virtualization inserts a layer between real hardware and traditional OS. This layer is commonly called the Virtual Machine Monitor (VMM) and it manages the hardware resources of a computing system.
- Requirements for a VMM is as follows :
  1. It should provide an environment for programs which is same as the original machine.
  2. Programs run in this environment should show, at worst, only minor decreases in speed.
  3. VMM should be in complete control of the system resources.
- Any program run under a VMM should exhibit a function identical to that which it runs on the original machine directly. VMM is tightly related to the architectures of processors.

3.4.9 Middleware Support for Virtualization

- Library-level virtualization is also known as user-level Application Binary Interface (ABI) or API emulation.
- Windows Application Binary Interface (WABI) is a software package from Sun Microsystems to allow certain Microsoft Windows applications under the X Window System.
- Wabi 2.2 runs under Solaris on SPARC, Intel, and PowerPC. Wabi works by providing translated versions of the three core Windows libraries, user.dll, kernel.dll, and gdi.dll which redirect Windows calls to Solaris equivalents.
- For code other than core library calls Wabi either executes the instructions directly on the hardware, if it is Intel, or emulates them, either one instruction at a time or by translating a block of instructions and caching the result (e.g. for a loop).
- WabiServer allows the Windows application and X display to be on different computers

CUDA architecture :

- The programming model provided by NVIDIA for GPU is called Compute Unified Device architecture (CUDA). It provides abstraction of thread group hierarchy and shared memory hierarchy to the programmer.
- Fig. 3.4.4 shows the basic concept of the vCUDA architecture.
- The vCUDA employs a client-server model to implement CUDA virtualization. It consists of three user space components: the vCUDA library, a virtual GPU in the guest OS , and the vCUDA stub in the host OS.

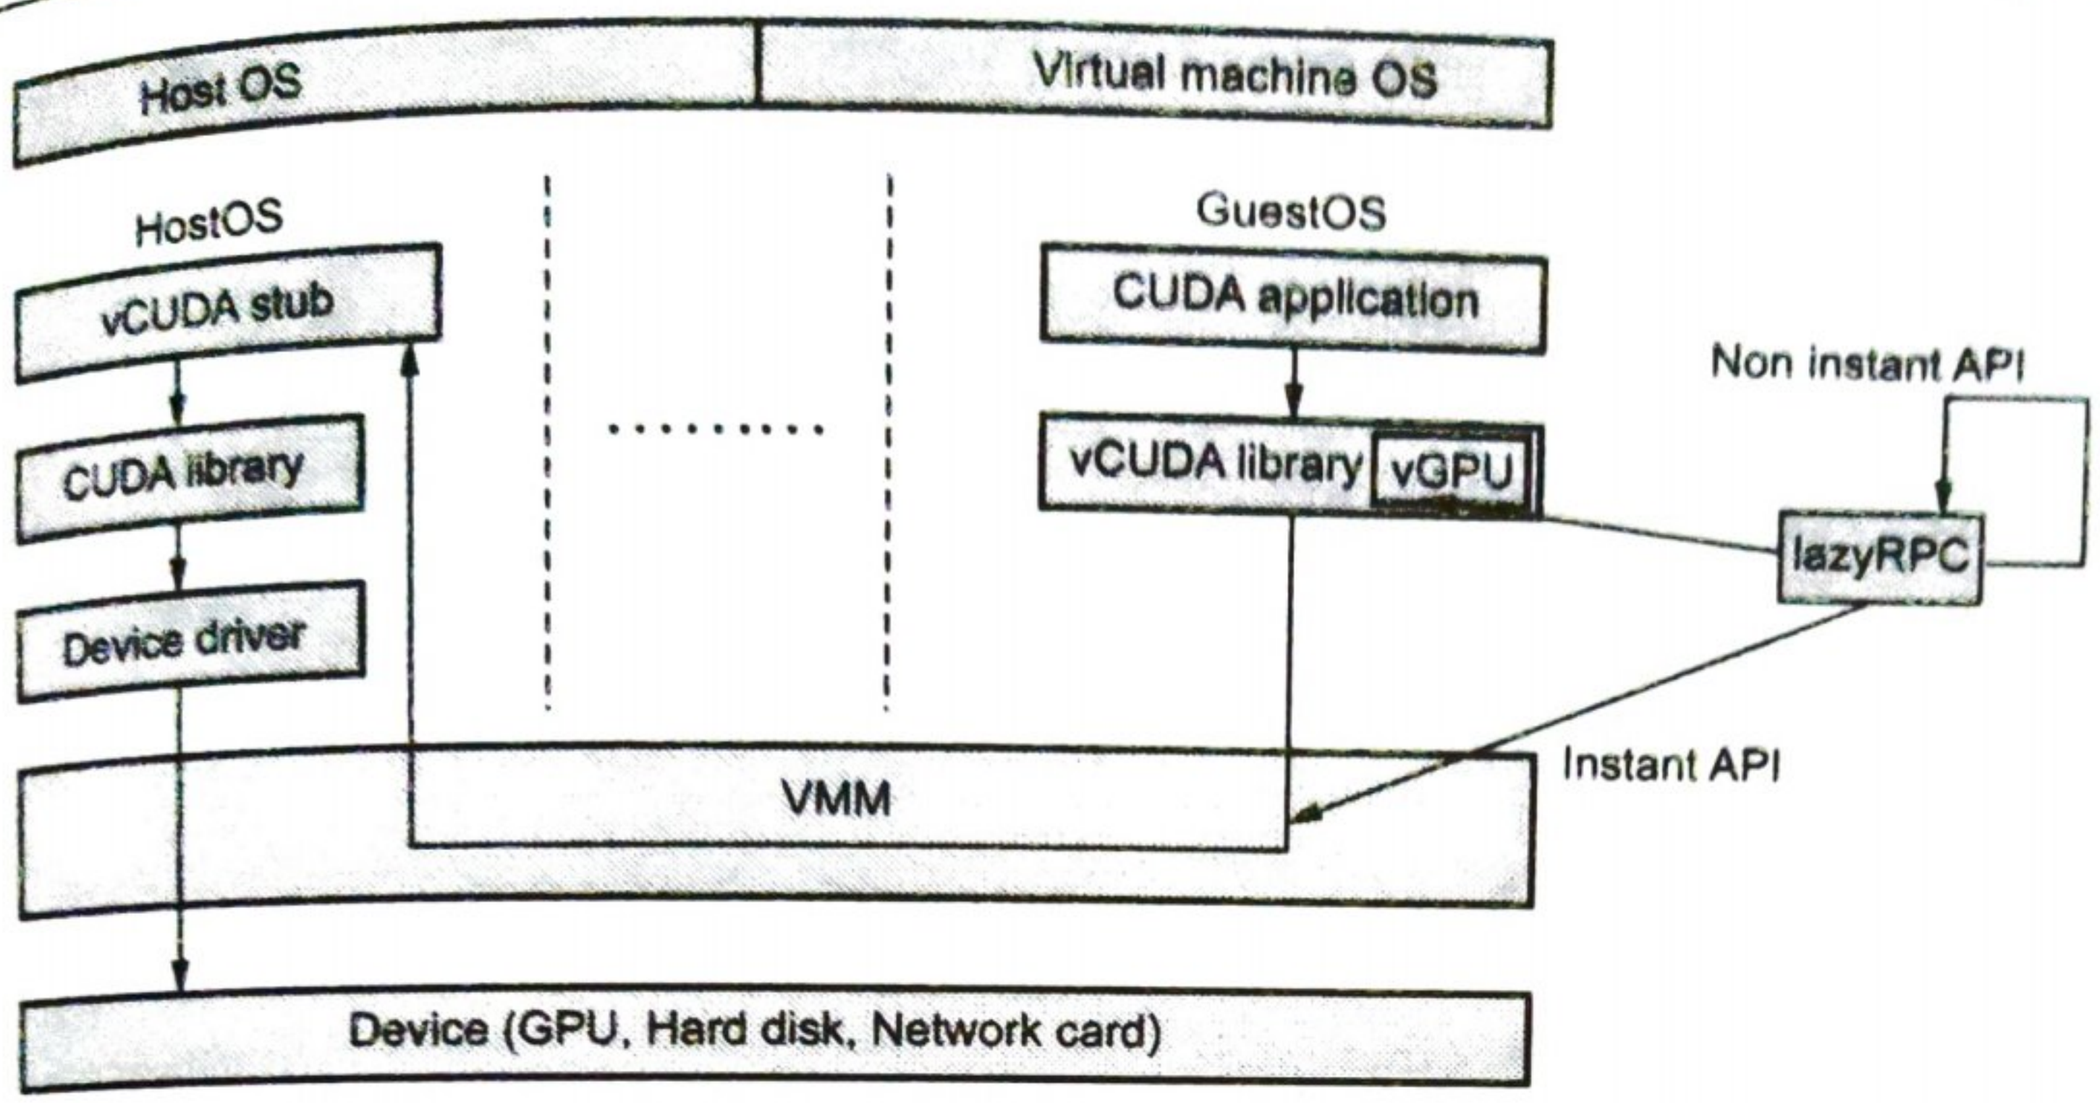

Fig. 3.4.4 Basic concept of the vCUDA architecture

- The vCUDA library resides in the guest OS as a substitute for the standard CUDA library. It is responsible for intercepting and redirecting API calls from the client to the stub. Besides these tasks, vCUDA also creates vGPUs and manages them.
- The vGPU is abstracts the GPU structure and gives applications a uniform view of the underlying hardware; when a CUDA application in the guest OS allocates a device's memory the vGPU can return a local virtual address to the application and notify the remote stub to allocate the real device memory, and the vGPU is responsible for storing the CUDA API flow.
- The vCUDA stub receives and interprets remote requests and creates a corresponding execution context for the API calls from the guest OS, then returns the results to the guest OS.
- The vCUDA stub also manages actual physical resource allocation.

3.4.10 Network Virtualization

- Network virtualization refers to the technology that enables partitioning or aggregating a collection of network resources and presenting them to various users in a way that each user experiences an isolated and unique view of the physical network.
- Network virtualization creates virtual networks whereby each application sees its own logical network independent of the physical network.
- A virtual LAN (VLAN) is an example of network virtualization that provides an easy, flexible, and less expensive way to manage networks.

- VLANs make large networks more manageable by enabling a centralized configuration of devices located in physically diverse locations.
- Fig. 3.4.5 shows network virtualization.
- Consider a company in which the users of a department are separated over a metropolitan area with their resources centrally located at one office.
- In a typical network, each location has its own network connected to the others through routers. When network packets cross routers, latency influences network performance.
- With VLANs, users with similar access requirements can be grouped together into the same virtual network. This setup eliminates the need for network routing.

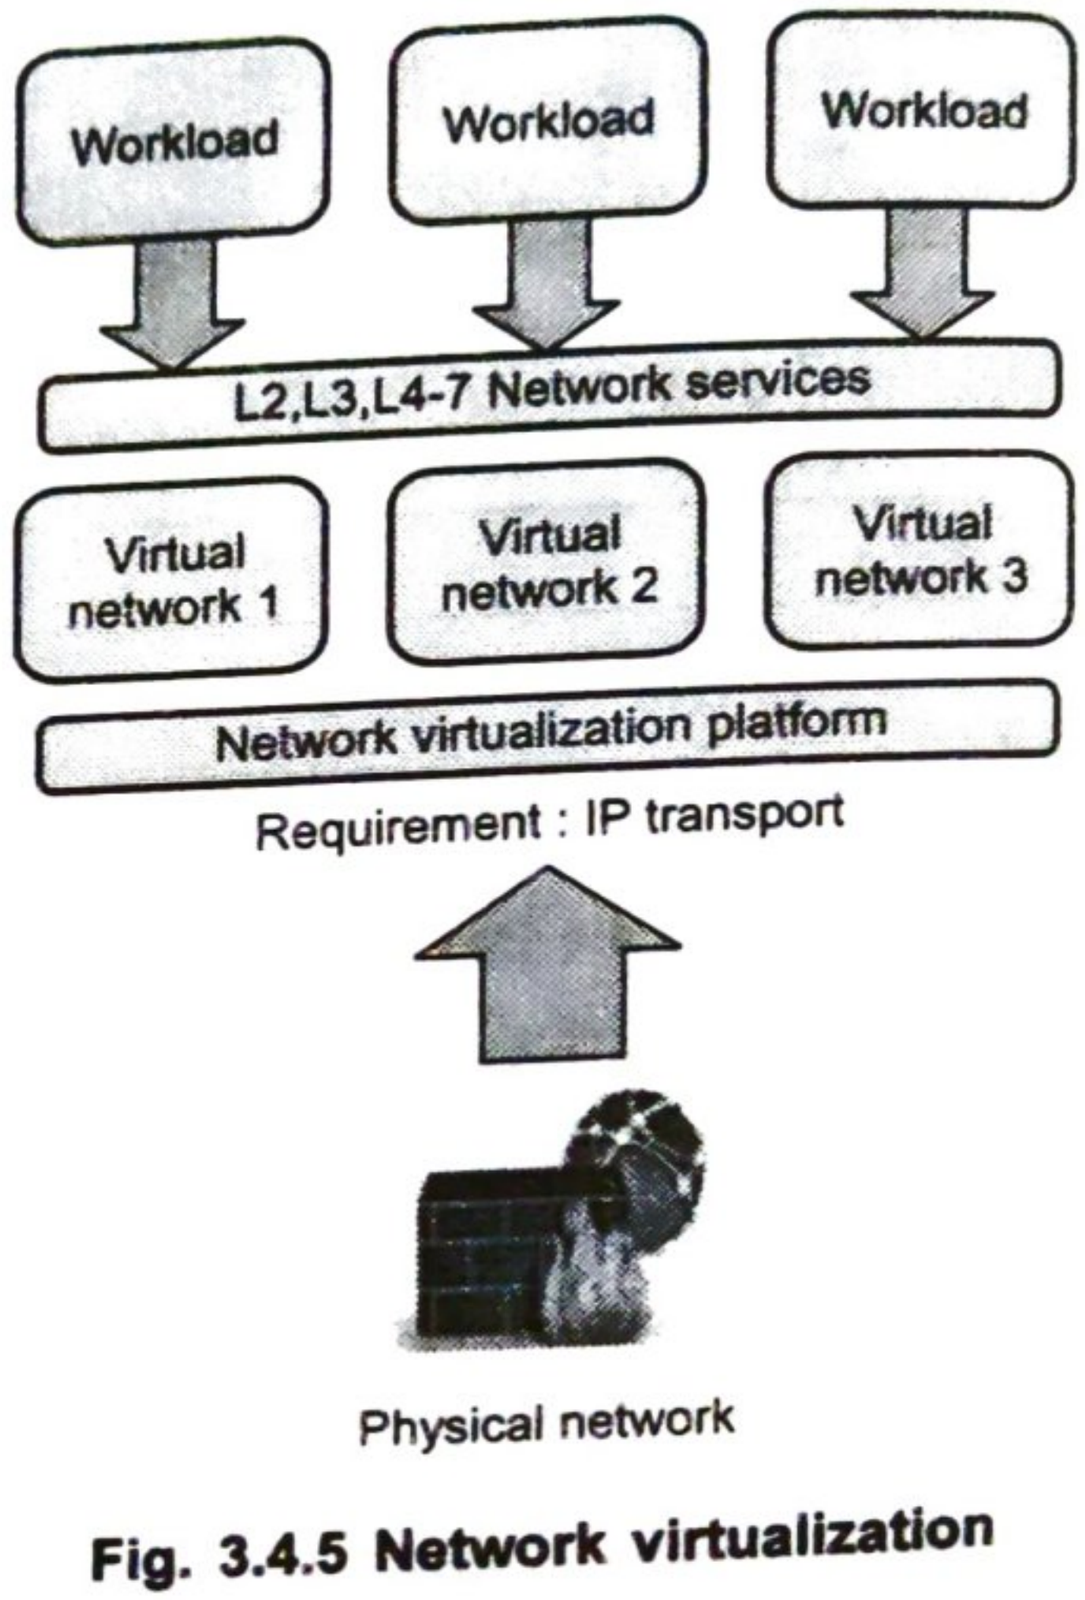

Fig. 3.4.5 Network virtualization

- As a result, although users are physically located at disparate locations, they appear to be at the same location accessing resources locally.
- In addition to improving network performance, VLANs also provide enhanced security by isolating sensitive data from the other networks and by restricting access to the resources located within the networks.
- Network virtualization decouples the roles of the traditional Internet service providers (ISPs) into infrastructure providers (InPs) and service providers (SPs)
- Benefits :
  1. Reduces the number of physical devices needed
  2. Easily segment networks
  3. Permits rapid change / scalability and agile deployment
  4. Security from destruction of physical devices

3.4.11 Application Level Virtualization

- Virtualization at the application level virtualizes an application as a VM. On a traditional OS, an application often runs as a process. Therefore, application-level virtualization is also known as process-level virtualization.
- A fully virtualized application is not installed in the traditional sense, although it is still executed as if it were. The application behaves at runtime like it is directly interfacing with the original operating system and all the resources managed by it, but can be isolated to varying degrees.
- Full application virtualization requires a virtualization layer. Application virtualization layers replace part of the runtime environment normally provided by the operating system.
- The layer intercepts all disk operations of virtualized applications and transparently redirects them to a virtualized location, often a single file.
- The application remains unaware that it accesses a virtual resource instead of a physical one. Since the application is now working with one file instead of many files spread throughout the system, it becomes easy to run the application on a different computer and previously incompatible applications can be run side-by-side.
- The most popular approach is to deploy High Level Language (HLL) VMs. Here the virtualization layer sits as an application program on top of the operating system, and the layer exports an abstraction of a VM that can run programs written and compiled to a particular abstract machine definition. Any program written in the HLL and compiled for this VM will be able to run on it.
- Benefits :
  1. Application virtualization uses fewer resources than a separate virtual machine.
  2. Application virtualization also enables simplified operating system migrations.
  3. Applications can be transferred to removable media or between computers without the need of installing them, becoming portable software.
- **Limitations :**
  1. Not all computer programs can be virtualized
  2. Lower performance

University Question

1. *Explain virtualization and hypervisor.* **GTU : Summer-17,Winer-18, Marks 7**

3.5 Full Virtualization

- Full Virtualization doesn't need to modify the host OS; it relies upon binary translation to trap and to virtualize certain sensitive instructions.
- Fig. 3.5.1 shows full virtualization.
- VMware Workstation applies full virtualization, which uses binary translation to automatically modify x86 software on-the-fly to replace critical instructions
- Normal instructions can run directly on the host OS. This is done to increase the performance overhead - normal instructions are carried out in the normal manner, but the difficult and precise executions are first discovered using a trap and executed in a virtual manner.
- This is done to improve the security of the system and also to increase the performance.

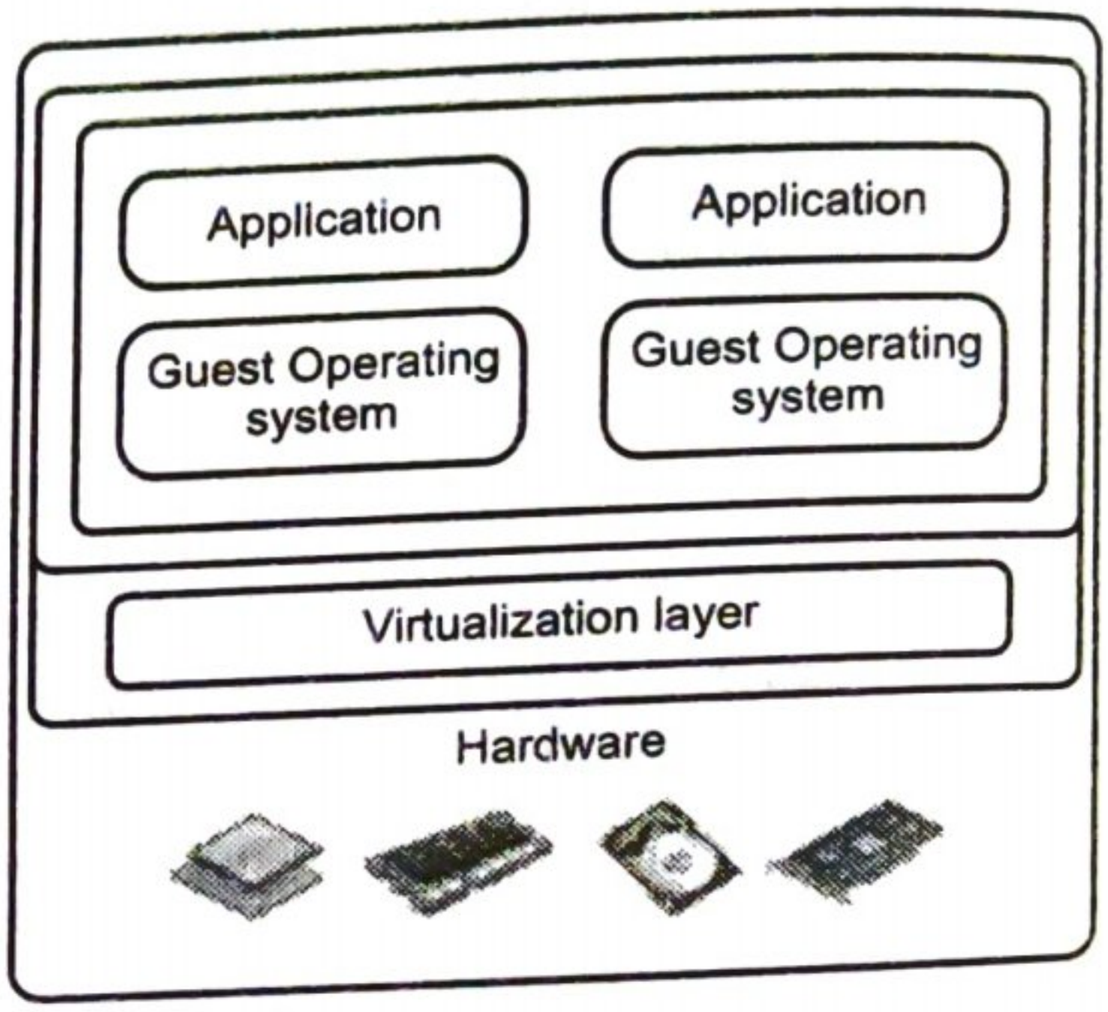

Fig. 3.5.1 Full virtualization

Host based virtualization :

- Virtualization implemented in a host computer rather than in a storage subsystem or storage appliance.
- Virtualization can be implemented either in host computers, in storage subsystems or storage appliances, or in specific virtualization appliances in the storage interconnect fabric.
- The guest OS are installed and run on top of the virtualization layer. Dedicated applications may run on the VMs. Certainly, some other applications can also run with the host OS directly.
- **Advantages** of host-based architecture :
  1. The user can install this VM architecture without modifying the host OS.
  2. The host-based approach appeals to many host machine configurations

3.5.1 Memory Virtualization

- Memory virtualization features allow abstraction isolation and monitoring of memory on a per Virtual Machine (VM) basis. These features may also make live migration of VMs possible, add to fault tolerance, and enhance security.
- Example features include Direct Memory Access (DMA) remapping and Extended Page Tables (EPT), including their extensions: accessed and dirty bits, and fast switching of EPT contexts.
- The VMkernel manages all machine memory. The VMkernel dedicates part of this managed machine memory for its own use. The rest is available for use by virtual machines.
- Virtual machines use machine memory for two purposes : each virtual machine requires its own memory and the VMM requires some memory and a dynamic overhead memory for its code and data.
- The virtual memory space is divided into blocks, typically 4KB, called pages. The physical memory is also divided into blocks, also typically 4KB.
- When physical memory is full, the data for virtual pages that are not present in physical memory are stored on disk. ESX/ESXi also provides support for large pages.
- The VMM is responsible for mapping the guest physical memory to the actual machine memory.
- Each page table of a guest OS has a page table allocated for it in the VMM. The page table in the VMM which handles all these is called a shadow page table.
- As it can be seen all this process is nested and inter-connected at different levels through the concerned address.
- If any change occurs in the virtual memory page table or TLB, the shadow page table in the VMM is updated accordingly.

3.5.2 I/O Virtualization

- I/O Virtualization involves managing of the routing of I/O requests between virtual devices and shared physical hardware.
- The there are three ways to implement this are full device emulation, para-VZ and direct I/O
- I/O virtualization features facilitate offloading of multi-core packet processing to network adapters as well as direct assignment of virtual machines to virtual functions, including disk I/O.

- Examples include Virtual Machine Device Queues (VMDQ), Single Root I/O Virtualization.
- Fig. 3.5.2 shows I/O virtualization.

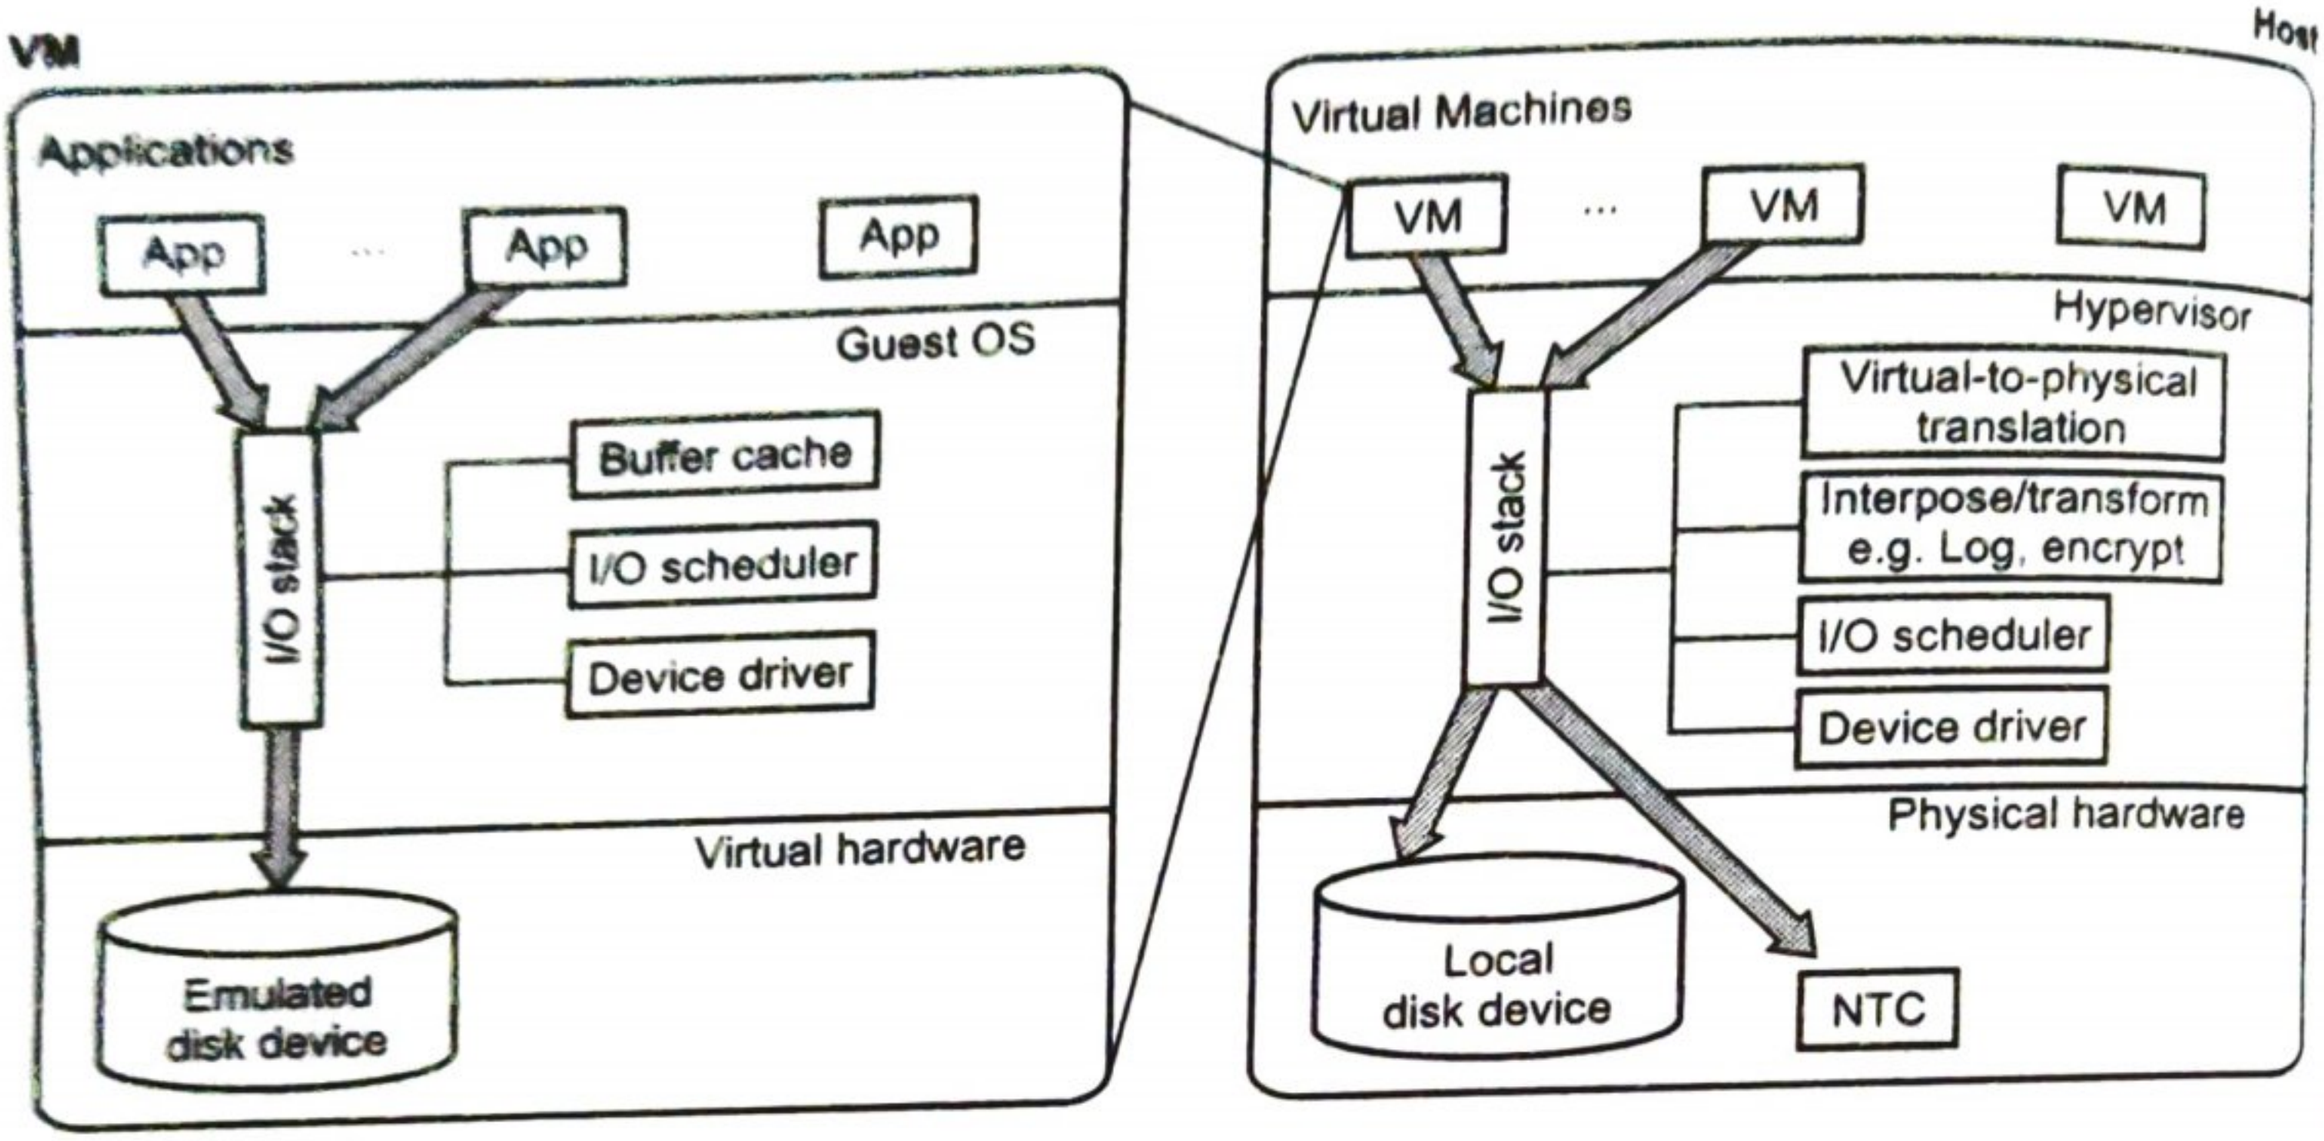

Fig. 3.5.2 I/O virtualization

1. **Full Device Emulation :** This process emulates well-known and real-world devices. All the functions of a device or bus infrastructure such as device enumeration, identification, interrupts etc. are replicated in the software, which itself is located in the VMM and acts as a virtual device. The I/O requests are trapped in the VMM accordingly.
2. **Para-virtualization :** This method of I/O VZ is taken up since software emulation runs slower than the hardware it emulates. In para-VZ, the frontend driver runs in Domain-U; it manages the requests of the guest OS. The backend driver runs in Domain-0 and is responsible for managing the real I/O devices. This methodology (para) gives more performance but has a higher CPU overhead.
3. **Direct I/O virtualization :** This lets the VM access devices directly; achieves high performance with lower costs. Currently, it is used only for the mainframes.

3.6 Virtual Clusters and Resource Management

- As with traditional physical servers, Virtual Machines (VMs) can also be clustered. A VM cluster starts with two or more physical servers.
- Most virtualization platforms, including XenServer and VMware ESX Server, support a bridging mode which allows all domains to appear on the network as individual hosts. By using this mode, VMs can communicate with one another

freely through the virtual network interface card and configure the network automatically.

- Virtual clusters enable admins to deploy, track and manage containers across various systems to ensure performance, security and governance, and low costs.
- With many VMs, an inefficient configuration always causes problems with overloading or underutilization.
- Amazon's EC2 provides elastic computing power in a cloud. EC2 permits customers to create VMs and to manage user accounts over the time of their use. Xen Server and VMware ESXi Server support a bridging mode which allows all domains to appear on the network as individual hosts. With this mode VMs can communicate with one another freely through the virtual network interface card and configure the network automatically

Physical versus Virtual Clusters :

- Virtual Clusters are built with VMs installed at one or more physical clusters. The VMs in a virtual cluster are interconnected by a virtual network across several physical networks.

Virtual cluster features :

a) Virtual machines can be restarted on other hosts if the host where the virtual machine running fails.

b) Distributed Resource Scheduler : virtual machines can be load balanced so that none of the hosts is too overloaded or too much empty in the cluster.

c) Live migration : of virtual machines from one host to other

- Fig. 3.6.1 shows cloud platform with virtual cluster.

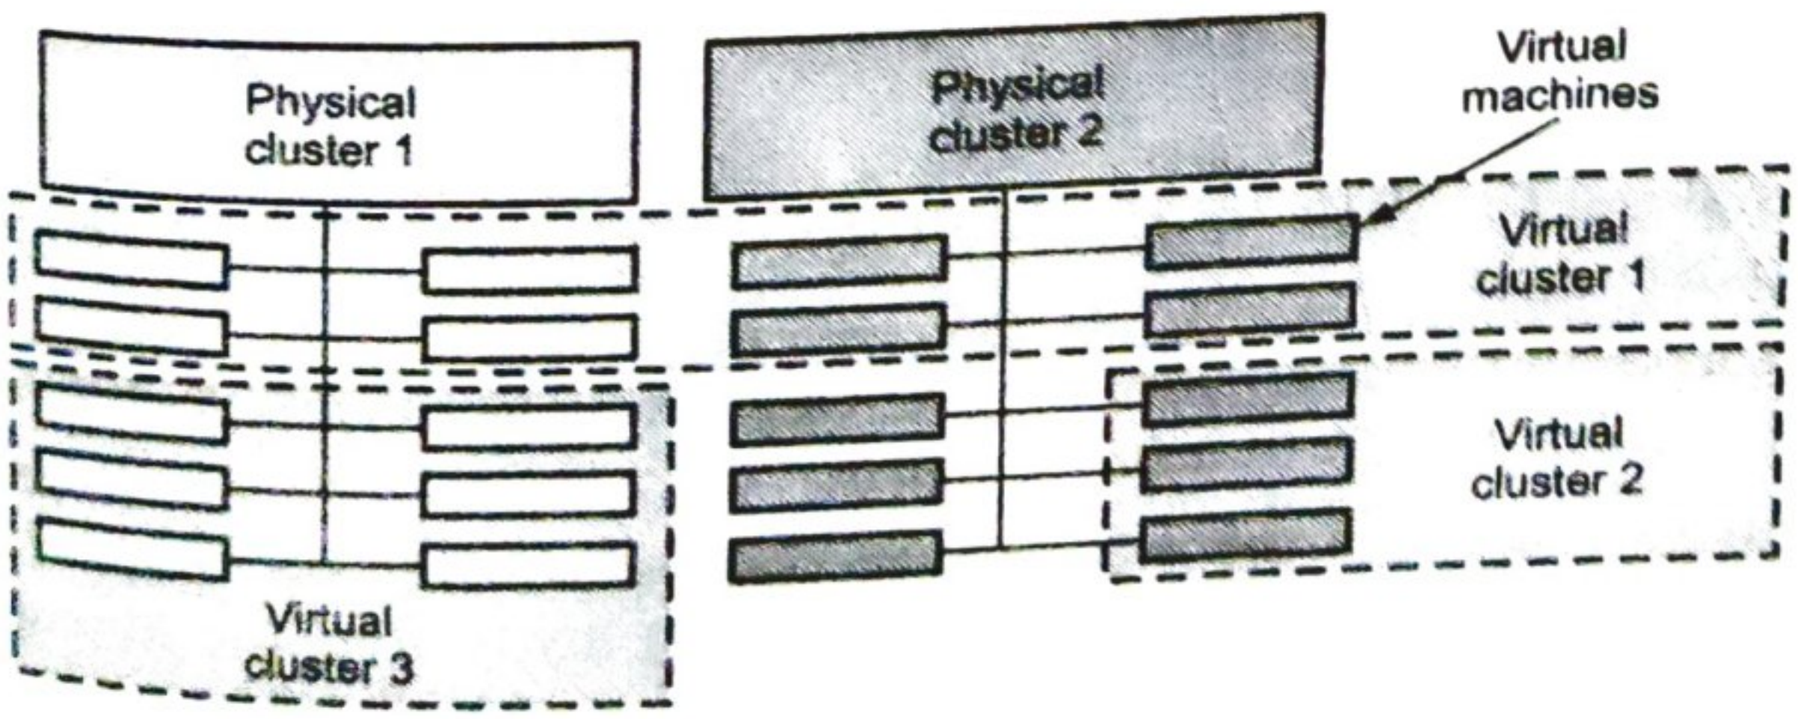

Fig. 3.6.1 cloud platform example with three virtual clusters over two physical clusters

- The provisioning of VMs to a virtual cluster is done dynamically and they have the following properties :

a) Virtual cluster nodes can be either physical or virtual with different operating systems.

b) VM runs with a guest OS that manages the resources in the physical machine.

c) The purpose of using VMs is to consolidate multiple functionalities on the same server.

d) VMs can be replicated in multiple servers to promote parallelism, fault tolerance and disaster discovery.

e) The no. of nodes in a virtual cluster can grow or shrink dynamically.

f) The failure of some physical nodes will slow the work but the failure of VMs will cause no harm.

Characteristics Virtual Cluster :

1. Virtual machine or physical machine is used as virtual cluster nodes. Multiple VM running with different types of OS can be deployed on the same physical node.
2. Virtual machine runs with guest operating system. Host OS and VM OS are different but it manages the resources in the physical machine.
3. Virtual machine can be replicated in multiple servers and it support distributed parallelism, fault tolerance and disaster recovery.
4. Number of nodes of a virtual cluster may change accordingly.
5. If Virtual machine failes, it can not affect the host machine.

- **Virtual cluster is managed by four ways :**

1. We can use a guest-based manager, by which the cluster manager resides inside a guest OS. Ex. : A Linux cluster can run different guest operating systems on top of the Xen hypervisor.
2. We can bring out a host-based manager which itself is a cluster manager on the host systems. Ex. : VMware HA (High Availability) system that can restart a guest system after failure.
3. An independent cluster manager, which can be used on both the host and the guest - making the infrastructure complex.
4. Finally, we might also use an integrated cluster (manager), on the guest and host operating systems; here the manager must clearly distinguish between physical and virtual resources.

3.7 Virtualization for Data Center Automation

- Data centers have grown rapidly in recent years, and all major IT companies are pouring their resources into building new data centers. Data Centers are specialized environments that safeguard company's most valuable equipment and intellectual property.
- The data center infrastructure is central to the IT architecture, from which all content is sourced or passes through. Proper planning of the data center infrastructure design is critical, and performance, resiliency, and scalability need to be carefully considered.
- Data center automation is the process by which routine workflows and processes of a data center, scheduling, monitoring, maintenance, application delivery are managed and executed without human administration.
- Data-center automation means that huge volumes of hardware, software, and database resources in these data centers can be allocated dynamically to millions of Internet users simultaneously, with guaranteed QoS and cost effectiveness.
- Data center automation increases agility and operational efficiency. It reduces the time IT needs to perform routine tasks and enables them to deliver services on demand in a repeatable, automated manner. These services can then be rapidly consumed by end users.
- Why data center automation is important
  a) Delivers insight into server nodes and configurations
  b) Automates routine procedures like patching, updating, and reporting
  c) Produces and programs all data center scheduling and monitoring tasks
  d) Enforces data center processes and controls in agreement with standards and policies

3.7.1 Server Consolidation in Data Centers

- It is common practice to dedicate each server to a single application. If several applications only use a small amount of processing power, the network administrator can combine several machines into one server running multiple virtual environments.
- In data centers, a large number of heterogeneous workloads can run on servers at various times. These heterogeneous workloads can be roughly divided into two categories : chatty workloads and non inter-active workloads.

- Consolidation enhances hardware utilization. Many underutilized servers are consolidated into fewer servers to enhance resource utilization. Consolidation also facilitates backup services and disaster recovery.
- The heterogeneous workloads in the data center is divided into two categories : chatty workloads and noninteractive workloads.
- Chatty workloads may burst at some point and return to a silent state at some other point. For example, video services can be used by a lot of people at night and few people use it during the day.
- Noninteractive workloads do not require people's efforts to make progress after they are submitted. Server consolidation is an approach to improve the low utility ratio of hardware resources by reducing the number of physical servers.
- The use of VMs increases resource management complexity.
- It enhances hardware utilization. Many underutilized servers are consolidated into fewer servers to enhance resource utilization. Consolidation also facilitates backup services and disaster recovery.
- In a virtual environment, the images of the guest OSes and their applications are readily cloned and reused.
- Total cost of ownership is reduced
- Improves availability and business continuity
- Automation of data-center operations includes resource scheduling, architectural support, power management, automatic or autonomic resource management, performance of analytical models, and so on.
- In virtualized data centers, an efficient, on-demand, fine-grained scheduler is one of the key factors to improve resource utilization.
- Dynamic CPU allocation is based on VM utilization and application-level QoS metrics.
- One method considers both CPU and memory flowing as well as automatically adjusting resource overhead based on varying workloads in hosted services.
- Another scheme uses a two-level resource management system to handle the complexity involved. A local controller at the VM level and a global controller at the server level are designed.
- Three resource managers are as follows :

1. Instance Manager controls the execution, inspection, and terminating of VM instances on the host where it runs.
2. Group Manager collects an information about schedules VM execution on specific instance managers and it manages virtual instance network.

3. Cloud Manager is the entry-point into the cloud for users and administrators. It queries node managers for information about resources, makes scheduling decisions, and implements them by making requests to group managers.

3.7.2 Trust Management in Virtualized Data Center

- Virtual machine in the host machine entirely encapsulates the state of the guest operating system running inside it.
- Encapsulated machine state can be copied and shared over the network and removed like a normal file, which proposes a challenge to VM security.
- In general, a VMM can provide secure isolation and a VM accesses hardware resources through the control of the VMM, so the VMM is the base of the security of a virtual system.
- Normally, one VM is taken as a management VM to have some privileges such as creating, suspending, resuming, or deleting a VM.

1. VM-Based Intrusion Detection

- Intrusion detection is a set of techniques and methods that are used to detect suspicious activity both at the network and host level. Intrusion Detection System is software, hardware or combination of both used to detect intruder activity.
- Fig. 3.7.1 shows IDS.

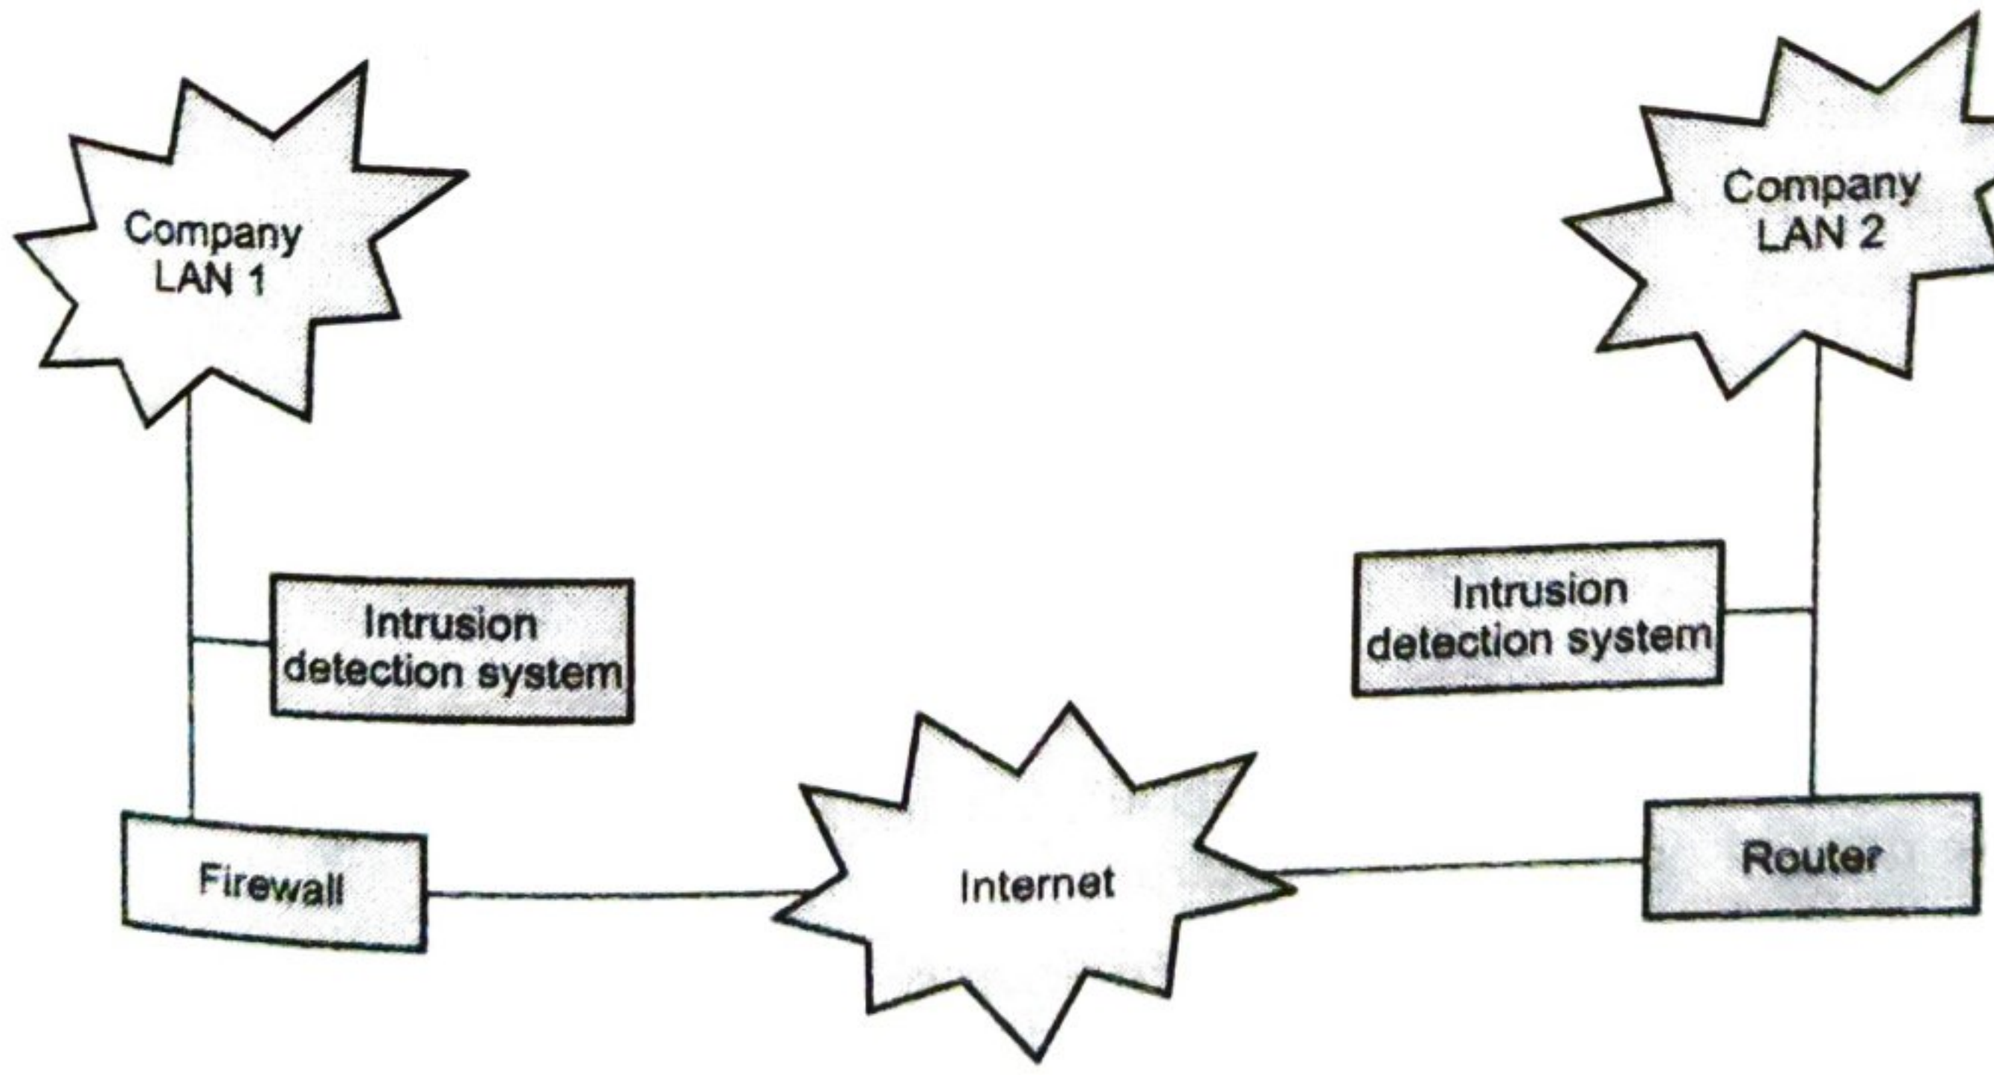

Fig. 3.7.1 IDS

- A lightweight intrusion detection system can easily be deployed on most any node of a network, with minimal disruption to operations. Snort is a libpcap based packet sniffer and logger that can be used as a lightweight network intrusion detection system.
- IDSs serve three essential security functions; monitor, detect and respond to unauthorized activity
- Functions of intrusion detection systems

1. IDS monitor and do analysis of user and system activity.
2. Auditing of system configurations and vulnerabilities.
3. Assessing the integrity of critical system and data files.
4. Recognition of activity patterns reflecting known attacks.
5. Statistical analysis for abnormal activity patterns.

Benefits of intrusion detection

1. Improving integrity of other parts of the information security infrastructure.
2. Improved system monitoring.
3. Tracing user activity from the point of entry to point of exit or impact.
4. Recognizing and reporting alterations to data files.
5. Spotting errors of system configuration and sometimes correcting them.
6. Recognizing specific types of attack and alerting appropriate staff for defensive responses.
7. Keeping system management personnel up to date on recent corrections to programs.
8. Allowing non-expert staff to contribute to system security.
9. Providing guidelines in establishing information security policies.

Limitations of IDS

1. Detect attack only after they have entered the network.
2. Cannot expect to detect all malicious activity at all-time handling alert to trigger false positive or false negative alarm.
3. Cannot integrated with filtering rules security to stop traffic from attacking

3.8 Multiple Choice Questions

Q.1 Hardware-level virtualization is performed right on top of the _____ hardware.

a bare
b layered
c closed
d none

Q.2 Operating system level virtualization refers to an _____ layer between traditional OS and user applications.

a middleware
b abstraction
c bottom
d all of these

Q.3 Application-level virtualization is also known as _____ level virtualization.

a system
b program
c thread
d process

Q.4 The hypervisor supports _____ virtualization on bare metal devices like CPU, memory, disk and network interfaces.

a I/O level
b OS level
c hardware level
d software level

Q.5 Xen is a _____ hypervisor, which separates the policy from the mechanism.

a kernel
b microkernel
c monolithic kernel
d hybrid kernel

Q.6 I/O virtualization involves managing the routing of I/O requests between virtual devices and the shared physical hardware.

a Virtual cluster
b Virtual OS
c Virtual memory
d Virtual devices

Q.7 A physical cluster is a collection of servers interconnected by a physical network such as _____.

a MAN
b WAN
c LAN
d All of these

Answer Keys for Multiple Choice Questions :

Q.1	a	Q.2	b
Q.3	d	Q.4	c
Q.5	b	Q.6	d
Q.7	c		

❑❑❑

4 Cloud Infrastructure and Cloud Resource Management

Syllabus

Architectural Design of Compute and Storage Clouds, Layered Cloud Architecture Development, Design Challenges, Inter Cloud Resource Management, Resource Provisioning and Platform Deployment, Global Exchange of Cloud Resources. Administrating the Clouds, Cloud Management Products, Emerging Cloud Management Standards.

Contents

4.1 Architectural Design of Compute and Storage Clouds

- Major design goals of a cloud computing platform is scalability, virtualization, efficiency, and reliability. Clouds support Web 2.0 applications.
- The cloud management receives the user request and then finds the correct resources, and then calls the provisioning services which invoke resources in the cloud. The cloud management software need to support both physical and virtual machines.
- The platform needs to establish a very large-scale HPC infrastructure. The hardware and software systems are combined together to make it easy and efficient to operate. The system scalability can benefit from cluster architecture.
- A cloud platform should be built to serve many users simultaneously. Therefore, multitasking is a necessity to assess distributed system performance.
- Five basic performance metrics are shown in Fig. 4.1.1. Refined performance models could be extended form basic attributes to include program behavior, environmental demand, QoS and cost-effectiveness.

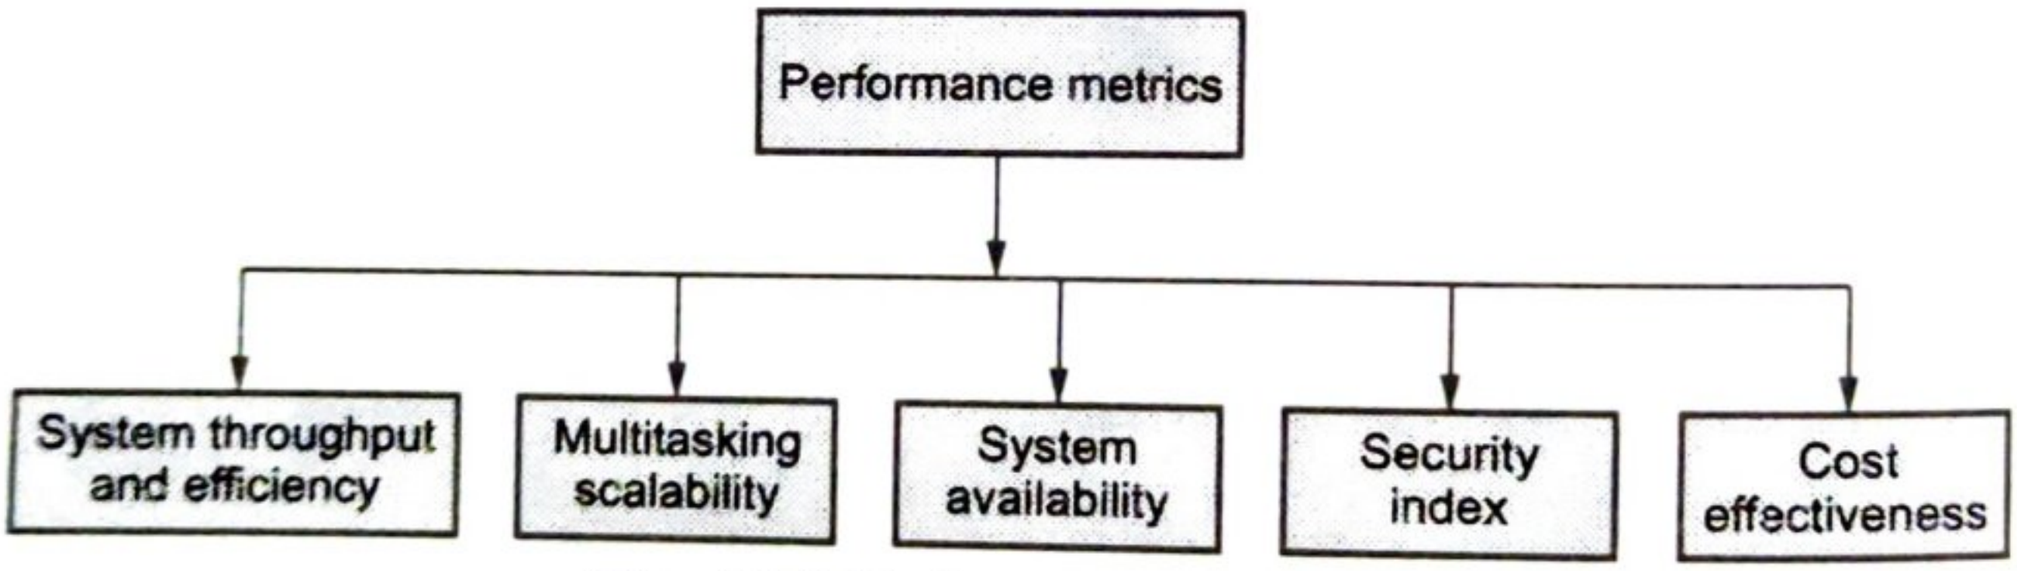

Fig. 4.1.1 Performance Metrics

- Enabling technologies for clouds : The key driving forces behind cloud computing are the ubiquity of broadband and wireless networking, falling storage costs, and progressive improvements in Internet computing software.
- Cloud users are able to demand more capacity at peak demand, reduce costs, experiment with new services, and remove unneeded capacity, whereas service providers can increase the system utilization via multiplexing, virtualization and dynamic resource provisioning.
- Resource virtualization enables rapid cloud deployment faster and fast disaster recovery. Service-oriented architecture (SOA) also plays a vital role. The progress in providing Software as a Service, Wed.2.0 standards and Internet performance have all contributed to the emergence of cloud services.
- The cloud computing resources are built in data centers, which are typically owned and operated by a third-party provider. Consumers do not need to know the underlying technologies.

- Web service providers offer special APIs that enable developers to exploit Internet clouds. Monitoring and metering units are used to track the usage and performance of resources provisioned. The software infrastructure of a cloud platform must handle all resource management and do most of the maintenance, automatically

4.1.1 Layered Cloud Architecture Development

- Fig. 4.1.2 shows layered architectural development of the cloud platform for IaaS, PaaS, and SaaS applications over the Internet and intranet.

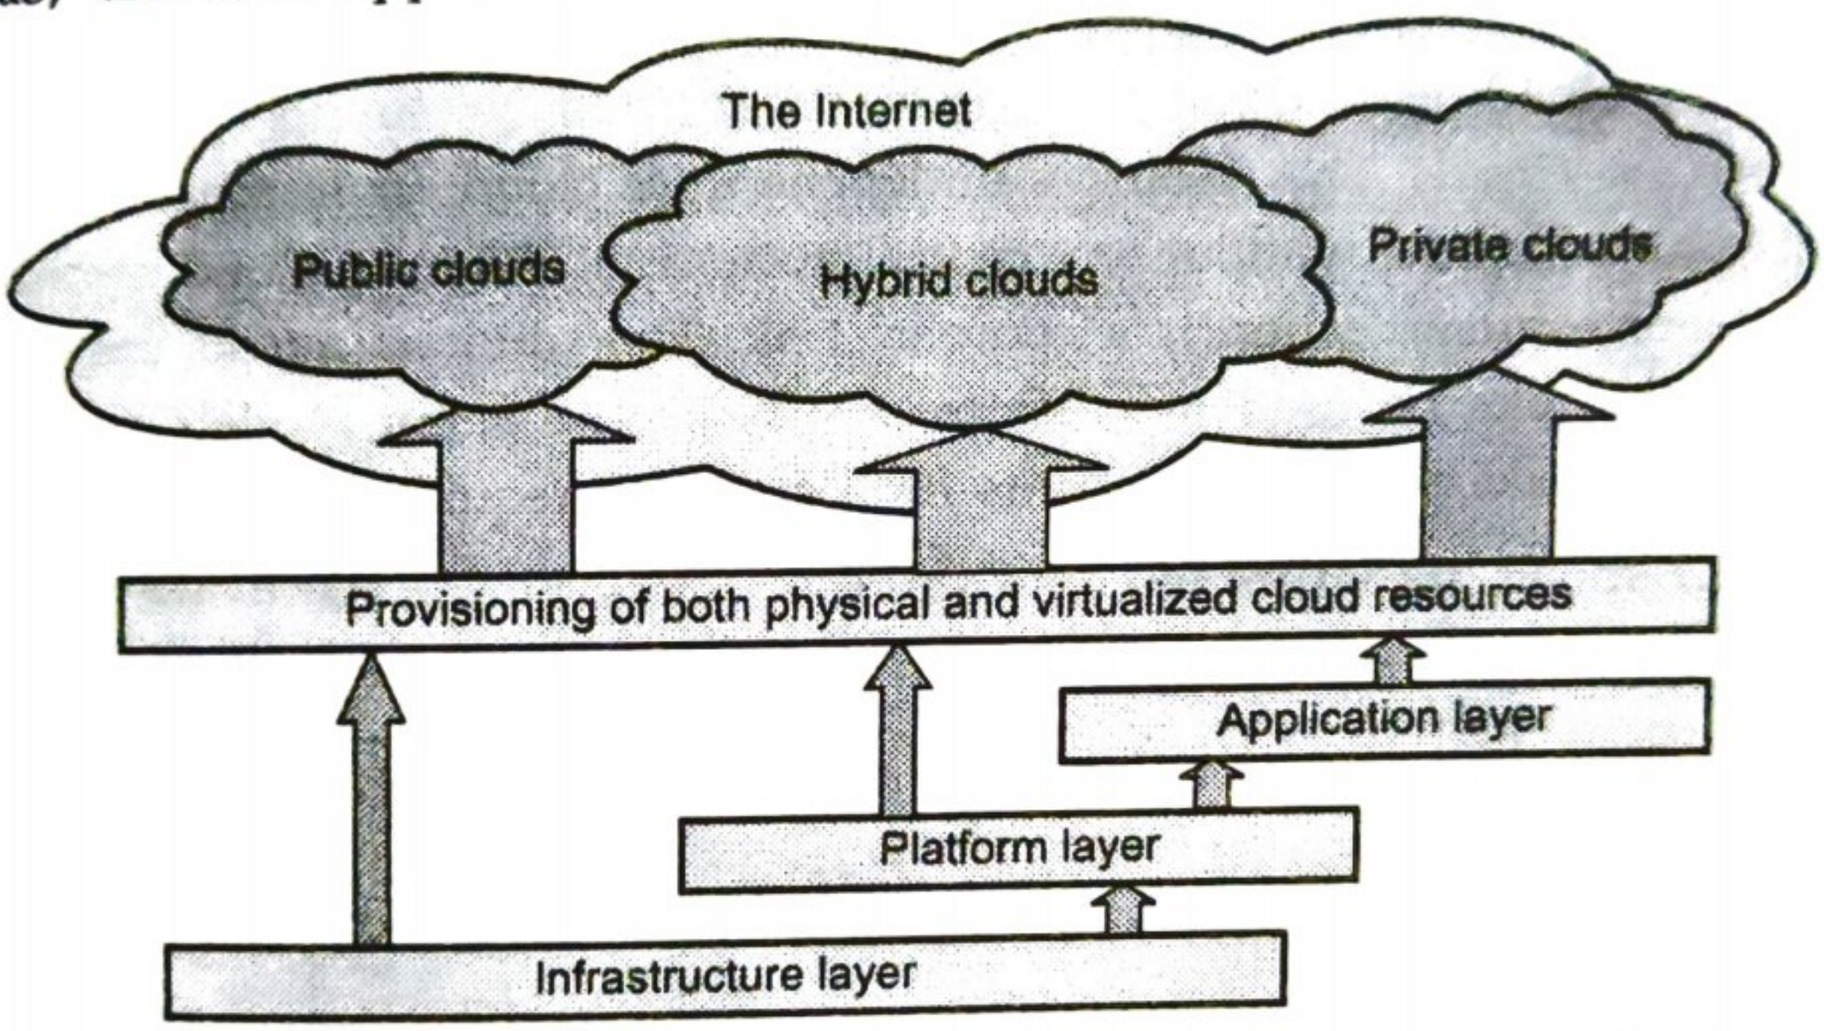

Fig. 4.1.2 Layered architectural development of the cloud platform

- The architecture of a cloud is developed at three layers : Infrastructure, platform, and application. These three development layers are implemented with virtualization and standardization of hardware and software resources provisioned in the cloud.
- The services to public, private, and hybrid clouds are conveyed to users through the networking support over the Internet and intranets involved. It is clear that the infrastructure layer is deployed first to support IaaS type of services.
- This infrastructure layer serves as the foundation to build the platform layer of the cloud for supporting PaaS services. The infrastructure layer is built with virtualized compute, storage and network resource.
- The platform layer is for general-purpose and repeated usage of the collection of software resources. The application layer is formed with a collection of all needed software modules for SaaS application.

4.1.2 Design Challenges

1. Service availability and data lock-in problem
2. Data privacy and security concerns
3. Unpredictable performance and bottlenecks
4. Distributed storage and wide-spread software bug
5. Cloud scalability, interoperability and standardization
6. Software licensing and reputation sharing

4.2 Inter Cloud Resource Management

- The inter cloud is a cloud of clouds constructed to support resource sharing between the clouds. The resources under the inter cloud environment are managed in distributed model without any central authority. The inter cloud communication and resource identification is a complex task. The software agents are small piece of code that can be used to perform any task. The agent models are applied to execute the tasks as small fragments for a specified requirement.
- Fig. 4.2.1 shows six layers of cloud services, ranging from hardware, network and collocation to infrastructure, platform and software applications

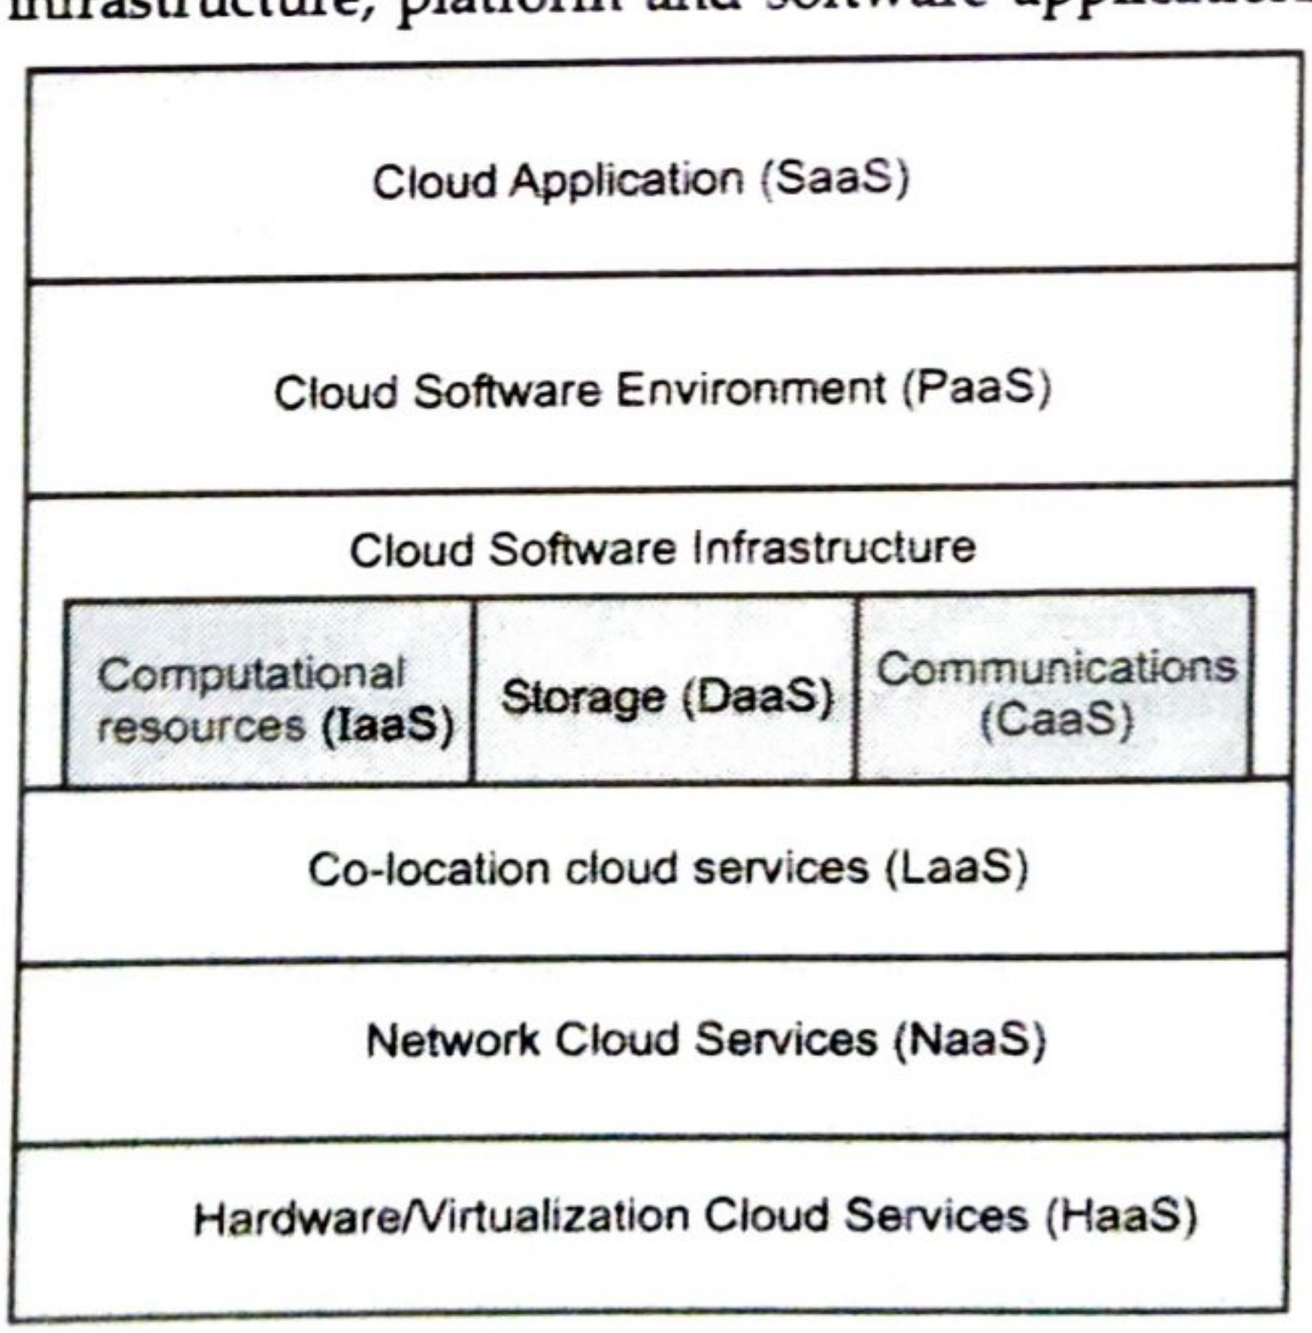

Fig. 4.2.1 Six layer stack

- The bottom most layer provides Hardware as a Service (HaaS). The next layer is for interconnecting all the hardware components, and is simply called Network as a Service (NaaS). Virtual LANs fall within the scope of NaaS.

- The next layer up offers Location as a Service (LaaS), which provides a collocation service to house, power, and secure all the physical hardware and network resources. The cloud infrastructure layer can be further subdivided as data as a service and communication as a service in addition to compute and storage in IaaS.
- The top layer is for SaaS applications. For example, CRM is heavily practiced in business promotion, direct sales, and marketing services. CRM offered the first SaaS on the cloud successfully.
- PaaS is provided by Google, Salesforce.com and Facebook, among others. IaaS is provided by Amazon, Windows Azure, and RackRack, among others.
- Runtime support services : As in a cluster environment, there are also some runtime supporting services in the cloud computing environment. Cluster monitoring is used to collect the runtime status of the entire cluster. Runtime support is software needed in browser-initiated applications applied by thousands of cloud customers.

4.2.1 Resource Provisioning and Platform Deployment

- Cloud architecture puts more emphasis on the number of processor cores.
- Provisioning of compute resources : Providers supply cloud services by signing SLAs with end users. The SLAs must commit sufficient resources such as CPU, memory and bandwidth that the user can use for a pre-set period.
- Under-provisioning of resources will lead to broken SLAs and penalties. Overprovisioning of resources will lead to resource underutilization and consequently, a decrease in revenue for the provider. Efficient VM provisioning depends on the cloud architecture and management of cloud infrastructures.
- Resource provisioning methods :
  a) The demand-driven method provides static resources and has been used in grid computing for many years. When a resource has surpassed a threshold for a certain amount of time, the scheme increases that resource based on demand.
  b) The event driven method is based on predicted workload by time. This scheme adds or removes machine instances based on a specific time event.
  c) The popularity-driven method is based on Internet traffic monitored. the Internet searches for popularity of certain applications and creates the instances by popularity demand.

4.2.2 Global Exchange of Cloud Resources

- In cloud computing, large numbers of customers use cloud services from all over the world. To ensure reliability in the cloud server, the service provider established various data centers in different locations worldwide.
- For example, the famous e-commerce website AMAZON has data centers in different geographical areas across the world. Even though the site has different data centers, it has specific limitations; for example, they don't have an automatic mechanism by which data centers at different locations can cooperate better and scale their different hosting services.
- Fig. 4.2.2 shows Inter-cloud exchange of cloud resources through brokering.

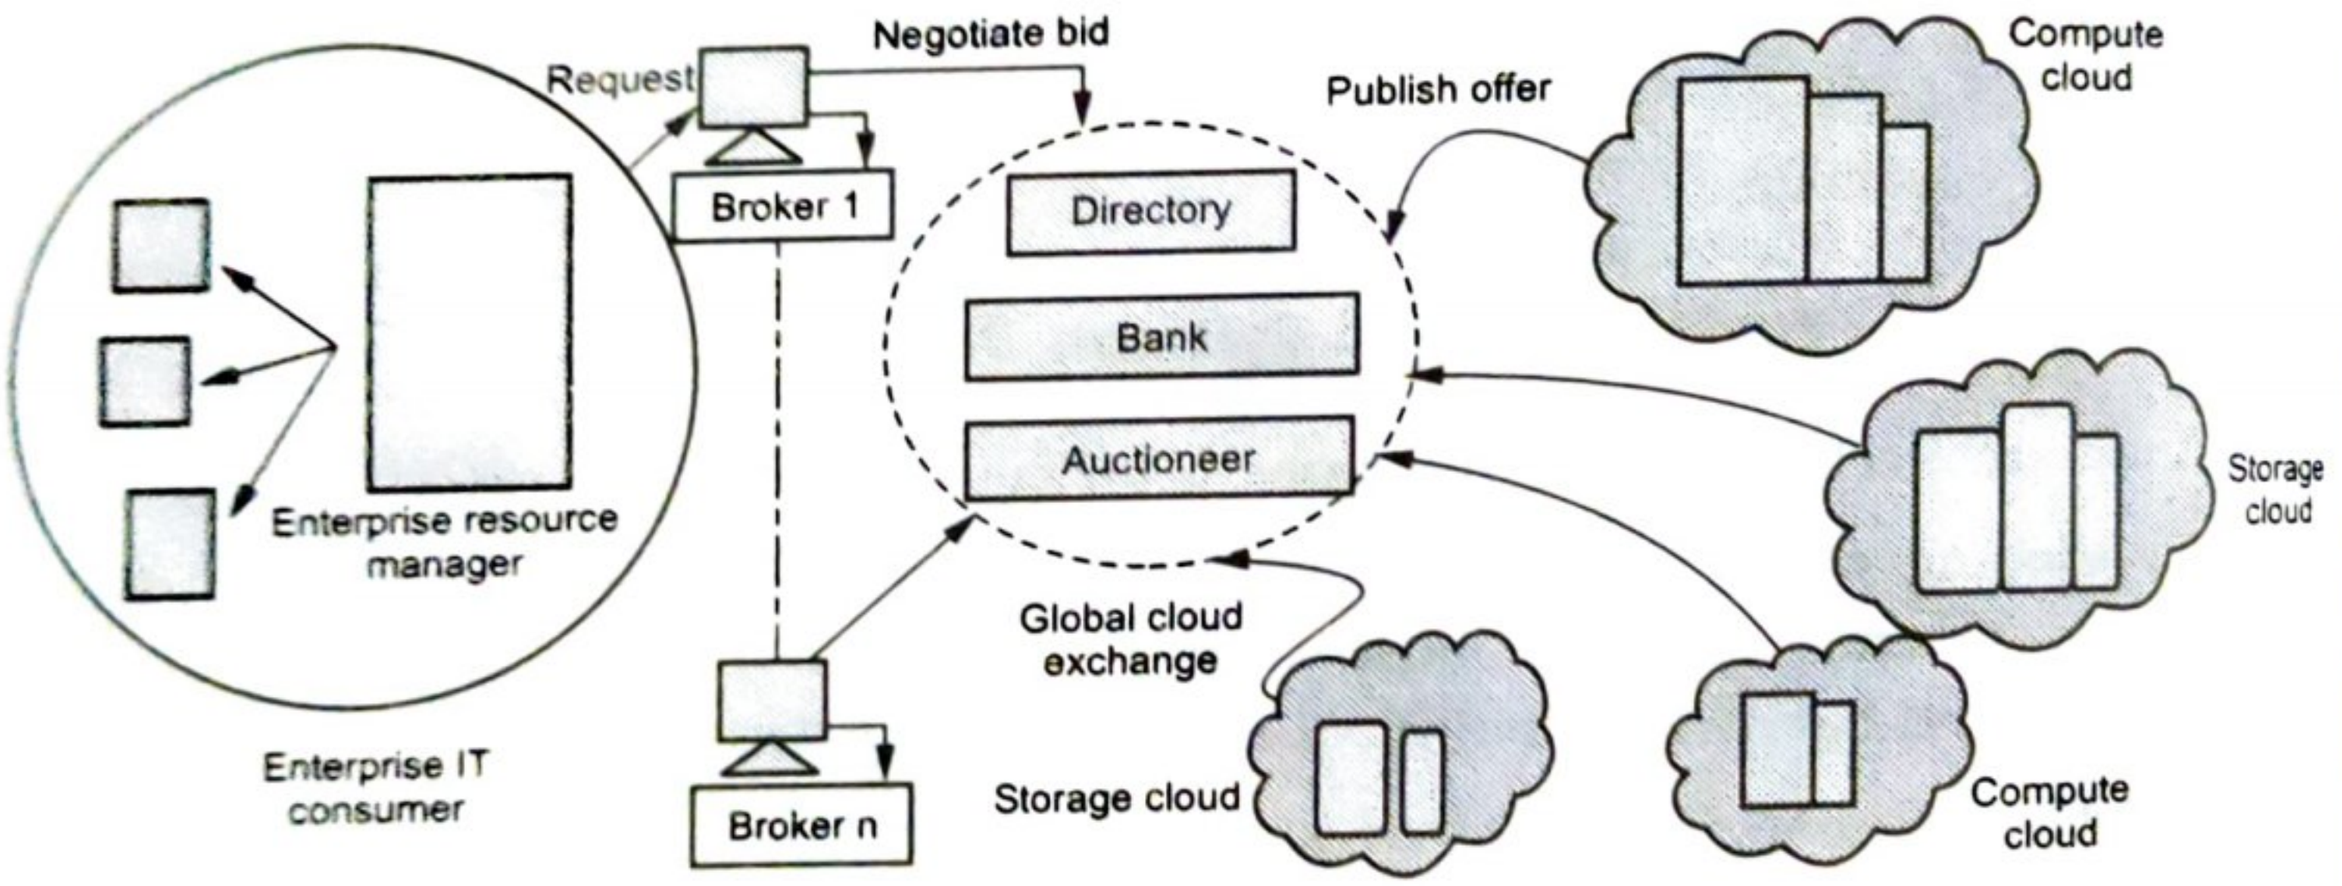

Fig. 4.2.2 Inter-cloud exchange of cloud resources through brokering

- The cloud exchange acts as a market maker for bringing together service producers and consumers. It aggregates the infrastructure demands from application brokers and evaluates them against the available supply currently published by the cloud coordinators.

4.3 Administrating the Clouds

- Administering cloud computing services is an important process when you have hosted your business data on the cloud. The business owners need to know whether the performance is at the right level and whether the deleted data is permanently gone.
- Cloud service provider can definitely build and provide a stable service that are cost effective and efficient. However, there can be a serious gap between the actual service and the promised services.

- You would need evaluate the solution providers when you are choosing a cloud application. Some of the questions that you can ask your vendors are :
  a) Are the vendors available to solve any software issues ?
  b) How will they manage if there is an outage ?
  c) How much experience do they hold in managing customer issues ?
  d) Will they provide training to the customers ?

4.3.1 Cloud Management Products

- Cloud management is the organized oversight, control, administration and maintenance of public cloud, private cloud or more commonly, hybrid multi-cloud computing infrastructure, services and resources.
- Cloud management services combine different technologies and products to deliver a cohesive, consistent strategy and process. Administrators can orchestrate delivery and management of cloud infrastructure, applications, data, services and access control. They can access resources, automate processes, make changes as needed and monitor utilization and cost.
- Cloud management platforms help IT teams secure and optimize cloud infrastructure, including the applications and data residing on it. Administrators can manage compliance, set up real-time monitoring, and pre-empt cyberattacks and data breaches.
- Typically, a cloud management system will be installed on the target cloud. It captures information on activity and performance then sends analysis to a web-based dashboard where administrators can see and act accordingly. Where there is an issue, administrators can issue commands back to the cloud through the cloud management platform, that servers as a consolidated point of control.

4.3.1.1 Dynamo

- Dynamo is propriety key value structured storage system. It can act as database and also distributed hash table.
- Dynamo dynamically partitions a set of keys over a set of storage nodes
- It is most powerful relational database available in WWW. Relational databases have been used a lot in retail sites, to make visitors browse and search for product easily.
- Dynamo does not support replication.
- Dynamo is used to manage the state of services that have very high reliability requirements and need tight control over the tradeoffs between availability, consistency, cost-effectiveness and performance.

- There are many services on Amazon's platform that only need primary-key access to a data store. For many services, such as those that provide best seller lists, shopping carts, customer preferences, session management, sales rank, and product catalog, the common pattern of using a relational database would lead to inefficiencies and limit scale and availability. Dynamo provides a simple primary-key only interface to meet the requirements of these applications.
- Dynamo is a completely decentralized system with minimal need for manual administration. Storage nodes can be added and removed from Dynamo without requiring any manual partitioning or redistribution.
- Compared to Bigtable, Dynamo targets applications that require only key/value access with primary focus on high availability where updates are not rejected even in the wake of network partitions or server failures.
- Dynamo stores objects associated with a key through a simple interface; it exposes two operations : get() and put().
- Dynamo treats both the key and the object supplied by the caller as an opaque array of bytes. It applies a MD5 hash on the key to generate a 128-bit identifier, which is used to determine the storage nodes that are responsible for serving the key.
- Dynamo's partitioning scheme relies on consistent hashing to distribute the load across multiple storage hosts. In consistent hashing , the output range of a hash function is treated as a fixed circular space or "ring".
- Dynamo provides eventual consistency, which allows for updates to be propagated to all replicas asynchronously.
- Dynamo uses vector clocks in order to capture causality between different versions of the same object. A vector clock is effectively a list of (node, counter) pairs. One vector clock is associated with every version of every object.
- In Dynamo, when a client wishes to update an object, it must specify which version it is updating. This is done by passing the context it obtained from an earlier read operation, which contains the vector clock information.
- In Dynamo, each storage node has three main software components: request coordination, membership and failure detection, and a local persistence engine. All these components are implemented in Java.

4.3.2 Emerging Cloud Management Standards

- The following working groups produce the standards and technologies promoted by the cloud management initiative :
  1. **Cloud Management Working Group (CMWG) :** Models the management of cloud services and the operations and attributes of the cloud service lifecycle through its work on the Cloud Infrastructure Management Interface (CIMI).

2. **Cloud Auditing Data Federation Working Group (CADF)** : Defines the CADF standard, a full event model anyone can use to fill in the essential data needed to certify, self-manage and self-audit application security in cloud environments.
3. **Software Entitlement Working Group (SEWG)** : Focuses on the interoperability with which software inventory and product usage are expressed, allowing the industry to better manage licensed software products and product usage.
4. **Open Virtualization Working Group (OVF)** : Produces the OVF standard, which provides the industry with a standard packaging format for software solutions based on virtual systems.

4.3.2.1 Open Cloud Consortium

- The Open cloud Consortium is a 501(c)(3) non-profit venture which provides cloud computing and data commons resources to support "scientific, environmental, medical and health care research."
- The Open Cloud Consortium has four working groups, one of which is the Open Science Data Cloud (OSDC).
- The infrastructure of the OSDC has been designed to address the challenges inherent in transporting large datasets, to balance the needs of data management and data analysis, and to archive data.
- The OSDC is based on a shared community infrastructure where hardware and software are shared among researchers and projects at the scale where it is most efficient to centrally locate and process data.
- It supports the development of standards for cloud computing and frameworks for interoperating between clouds; develops benchmarks for cloud computing; and supports reference implementations for cloud computing, preferably open source reference implementations.
- The OCC has a particular focus in large data clouds. It has developed the "MalStone Benchmark" for large data clouds and is working on a reference model for large data clouds

4.3.2.2 Open Virtualization Format

- OVF is an open standard, specified by the Distributed Management Task Force (DMTF), for packaging and distributing a virtual appliance consisting of one or more virtual machines (VMs).

- An OVF Package is composed of metadata and file elements that describe virtual machines, plus additional information that is important to the deployment and operation of the applications in the OVF package. Its file extension is .ovf.
- An OVF Package always includes a descriptor file (*.ovf) and may also include a number of other files

File type	Description
Descriptor	The descriptor specifies the virtual hardware requirements of the service and can also include other information such as descriptions of virtual disks, the service itself, and guest operating systems, a license agreement (EULA), instructions to start and stop VMs in the appliance, and instructions to install the service. The descriptor file extension is .ovf.
Manifest	The manifest is an SHA-1 digest of every file in the package, allowing the package contents to be verified by detecting any corruption. The manifest file extension is .mf.
Signature	The signature is the digest of the manifest signed with the public key from the X.509 certificate included in the package, and allows the package author to be verified. The signature file extension is .cert.
Virtual disks	OVF does not specify a disk image format. An OVF package includes files comprising virtual disks in the format defined by the virtualization product that exported the virtual disks. XenServer produces OVF packages with disk images in Dynamic VHD format; VMware products and Virtual Box produce OVF packages with virtual disks in Stream-Optimized VMDK format

4.4 Multiple Choice Questions

Q.1 Which of the following is cloud performance metric?

a Multitasking scalability

b System availability

c Security index

d All of these

Q.2 The _____ driven method is used on internet traffic monitored.

a demand

b event

c popularity

d None

Q.3 Dynamo is propriety key value structured storage system, it can act as database and also distributed _____ .

a hash table

b system

c memory system

d All of these

Q.4 Dynamo uses _____ clocks in order to capture causality between different versions of the same object.

a physical

b vector

c logical

d None

Q.5 Cloud Auditing Data Federation Working Group defines

a Models the management of cloud services and the operations and attributes of the cloud service lifecycle

b OVF standard, which provides the industry with a standard packaging format for software solutions based on virtual systems

c CADF standard, a full event model anyone can use to fill in the essential data needed to certify

d All of these

Answer Keys for Multiple Choice Questions :

Q.1	d	Q.2	c
Q.3	a	Q.4	b
Q.5	c		

❑❑❑

5 Security

Syllabus

Security Overview, Cloud Security Challenges and Risks, Software-as-a Service Security, Cloud computing security architecture : Architectural Considerations, General Issues Securing the Cloud, Securing Data, Data Security, Application Security, Virtual Machine Security, Identity and Presence, Identity Management and Access Control, Autonomic Security Establishing Trusted Cloud computing, Secure Execution Environments and Communications, Identity Management and Access control Identity management, Access control, Autonomic Security Storage Area Networks, Disaster Recovery in Clouds.

Contents

5.1 Security Overview

- Cloud security is the protection of data stored online via cloud computing platforms from theft, leakage, and deletion. Methods of providing cloud security include firewalls, penetration testing, tokenization, Virtual Private Networks (VPN), and avoiding public internet connections.
- Cloud security refers to an array of policies, technological procedures, services, and solutions designed to support safe functionality when building, deploying, and managing cloud-based applications and associated data.
- Cloud security is designed to protect the following, regardless of your responsibilities :
  a) **Physical networks** - Routers, electrical power, cabling, climate controls, etc.
  b) **Data storage** - Hard drives, etc.
  c) **Data servers** - Core network computing hardware and software
  d) **Computer virtualization frameworks** - Virtual machine software, host machines, and guest machines
  e) **Operating systems (OS)** - Software that houses
  f) **Middleware** - Application programming interface (API) management,
  g) **Runtime environments** - Execution and upkeep of a running program
  h) **Data** - All the information stored, modified, and accessed
  i) **Applications** - Traditional software services (email, tax software, productivity suites, etc.)
  j) **End-user hardware** - Computers, mobile devices, Internet of Things (IoT) devices, etc.
- Cloud computing security addresses both physical and logical security issues across all the different service models of software, platform and infrastructure. It also addresses how these services are delivered in the Public, Private, Hybrid and Community delivery models.

5.1.1 Cloud Security Challenges and Risks

- Cloud computing security challenges fall into three broad categories :
  1. Data protection : Securing your data both at rest and in transit.
  2. User authentication : Limiting access to data and monitoring who accesses the data.
  3. Disaster and data breach : Contingency planning.

- Data protection : Data needs to be encrypted at all times, with clearly defined roles when it comes to who will be managing the encryption keys.
- User authentication : Data resting in the cloud needs to be accessible only by those authorized to do so, making it critical to both restrict and monitor who will be accessing the company's data through the cloud. In order to ensure the integrity of user authentication, companies need to be able to view data access logs and audit trails to verify that only authorized users are accessing the data.
- Contingency planning : With the cloud serving as a single centralized repository for a company's mission-critical data, the risks of having that data compromised due to a data breach or temporarily made unavailable due to a natural disaster are real concerns.
- If information is encrypted while passing through the cloud, who controls the encryption/decryption keys ? Is it the customer or the cloud vendor ? Most customers probably want their data encrypted both ways across the Internet using Secure Sockets Layer protocol.
- They also most likely want their data encrypted while it is at rest in the cloud vendor's storage pool. Be sure that you, the customer, control the encryption/decryption keys, just as if the data were still resident on your own servers.
- Data integrity means ensuring that data is identically maintained during any operation.
- Cloud-based services will result in many mobile IT users accessing business data and services without traversing the corporate network. This will increase the need for enterprises to place security controls between mobile users and cloud-based services.
- Placing large amounts of sensitive data in a globally accessible cloud leaves organizations open to large distributed threats, attackers no longer have to come onto the premises to steal data, and they can find it all in the one "virtual" location.
- Virtualization efficiencies in the cloud require virtual machines from multiple organizations to be co-located on the same physical resources. Although traditional data center security still applies in the cloud environment, physical segregation and hardware-based security cannot protect against attacks between virtual machines on the same server.
- Operating system and application files are on a shared physical infrastructure in a virtualized cloud environment and require system, file, and activity monitoring to provide confidence and auditable proof to enterprise customers that their resources have not been compromised or tampered with.

- In the cloud computing environment, the enterprise subscribes to cloud computing resources, and the responsibility for patching is the subscriber's rather than the cloud computing vendor's.
- The need for patch maintenance vigilance is imperative. Lack of due diligence in this regard could rapidly make the task unmanageable or impossible, leaving you with "virtual patching" as the only alternative.
- Confidentiality : Confidentiality refers to limiting information access. Sensitive information should be kept secret from individuals who are not authorized to see the information.
- In cloud environments, confidentiality primarily pertains to restricting access to data in transit and storage.
- Integrity can extend to how data is stored, processed, and retrieved by cloud services and cloud-based IT resources.
- Some common cloud security threats include :
  a) Risks of cloud-based infrastructure including incompatible legacy IT frameworks, and third-party data storage service disruptions.
  b) Internal threats due to human error such as misconfiguration of user access controls.
  c) External threats caused almost exclusively by malicious actors, such as malware, phishing, and DDoS attacks.

5.2 Software-as-a Service Security

- Software-as-a-Service is a model of software deployment in which an application is licensed for use as a service provided to customers on demand. On-demand licensing and use relieves the customer of the burden of equipping a device with every application to be used.
- SaaS Security refers to securing user privacy and corporate data in subscription-based cloud applications. SaaS applications carry a large amount of sensitive data and can be accessed from almost any device by a mass of users, thus posing a risk to privacy and sensitive information.
- Pillars to SaaS-specific security :
  1. Access Management : The vendor must provide a unified framework to manage user authentication through business rules.
  2. Network Control : Security groups control who can access specific instances across the network.

3. Perimeter Network Control : Perimeter defence has traditionally been about controlling traffic flowing into and out of a data center network. The primary technology that underpins perimeter protection is a firewall.
4. VM Management : Ensuring your infrastructure is secure requires frequent updates directly to your virtual machine.
5. Data Protection : The most important practice of all is the SaaS provider's methodology for preventing a data breach, primarily by using various methods for data encryption both at rest and in transit.
6. Governance and Incident Management : Certain types of incidents must be captured, reported, and tracked to closure, and there must be procedures in place for investigating any potential security breaches.
7. Scalability and Reliability : Vertical scaling is limited by only being able to get as large as the size of the server. Horizontal scaling means the ability to connect multiple hardware or software entities, such as servers, so that they work as a single logical unit.

5.2.1 SaaS Security Challenges

a. **Data Security** : In SaaS scenario, data resides in the database which is outside the boundary of the enterprise and depends on the provider for proper security measures.

b. **Application Security** : SaaS applications are mostly used and managed over the web. They are presented to users in a browser. This makes it inevitable to confront the security challenges such as SQL injection, Cross-site scripting and Cross-site Request Forgery.

c. **Software-as-a-Service Deployment Security** : Virtualization refers to the act of creating different instances on hardware and on each instance a guest OS is installed.

- **Risk Management** : Effective risk management entails identification of technology assets; identification of data and its links to business processes, applications, and data stores; and assignment of ownership.
- **Risk Assessment** : Security risk assessment is critical to helping the information security organization make informed decisions when balancing the duelling priorities of business utility and protection of assets.
- **Security Portfolio Management** : Lack of portfolio and project management discipline can lead to projects never being completed; unsustainable and unrealistic workloads and expectations because projects are not prioritized according to strategy, goals, and resource capacity.

- **Security Awareness :** Not providing proper awareness and training to the people who may need them can expose the company to a variety of security risks for which people.
- **Policies, Standards, and Guidelines :** Many resources and templates are available to aid in the development of information security policies, standards, and guidelines. Policies should be developed, documented, and implemented, along with documentation for supporting standards and guidelines.
- **Third-Party Risk Management :** Lack of a third-party risk management program may result in damage to the provider's reputation, revenue losses, and legal actions should the provider be found not to have performed due diligence on its third-party vendors.
- **Forensics :** Computer forensics is used to retrieve and analyse data. The computer forensics means responding to an event by gathering and preserving data, analysing data to reconstruct events

Business Continuity Plan (BCP)

- Business continuity refers to the activities required to keep the organization running during a period of displacement or interruption of normal operations.
- BCP helps in continuing the business even after a disaster occurs.
- Business has to stay active during the crisis; if it closes its operations even for a day or a week, they are many chances that the organization will experience losses and will have to shut down.
- Moreover, legal issues can arise if the critical services are not provided to clients. This can lead to bad reputation and many more legal problems for an organization in addition to having the pain of being in the state of disaster. Hence an efficient BCP plan can be used to actively run and maintain the business activities.

5.2.2 Secure Software Development Life Cycle

- The SDLC consists of six phases, and there are steps unique to the Secure Software Development Life Cycle (SecSLDC) in each of phases :

Phase 1. Investigation : Define project processes and goals, and document them in the program security policy.

Phase 2. Analysis : Analyse existing security policies and programs, analyse current threats and controls, examine legal issues, and perform risk analysis.

Phase 3. Logical design : Develop a security blueprint, plan incident response actions, plan business responses to disaster, and determine the feasibility of continuing and/or outsourcing the project.

Phase 4. Physical design : Select technologies to support the security blueprint, develop a definition of a successful solution, design physical security measures to support technological solutions, and review and approve plans.

Phase 5. Implementation : Buy or develop security solutions. At the end of this phase, present a tested package to management for approval.

Phase 6. Maintenance : Constantly monitor, test, modify, update, and repair to respond to changing threats.

5.3 Cloud Security Architecture

GTU : Summer-17

- Cloud security architecture describes all the hardware and technologies designed to protect data, workloads, and systems within cloud platforms.
- Fig. 5.3.1 shows NIST cloud computing security reference architecture approach. The reference architecture identifies the five major cloud actors; consumer, provider, broker, carrier, and auditor

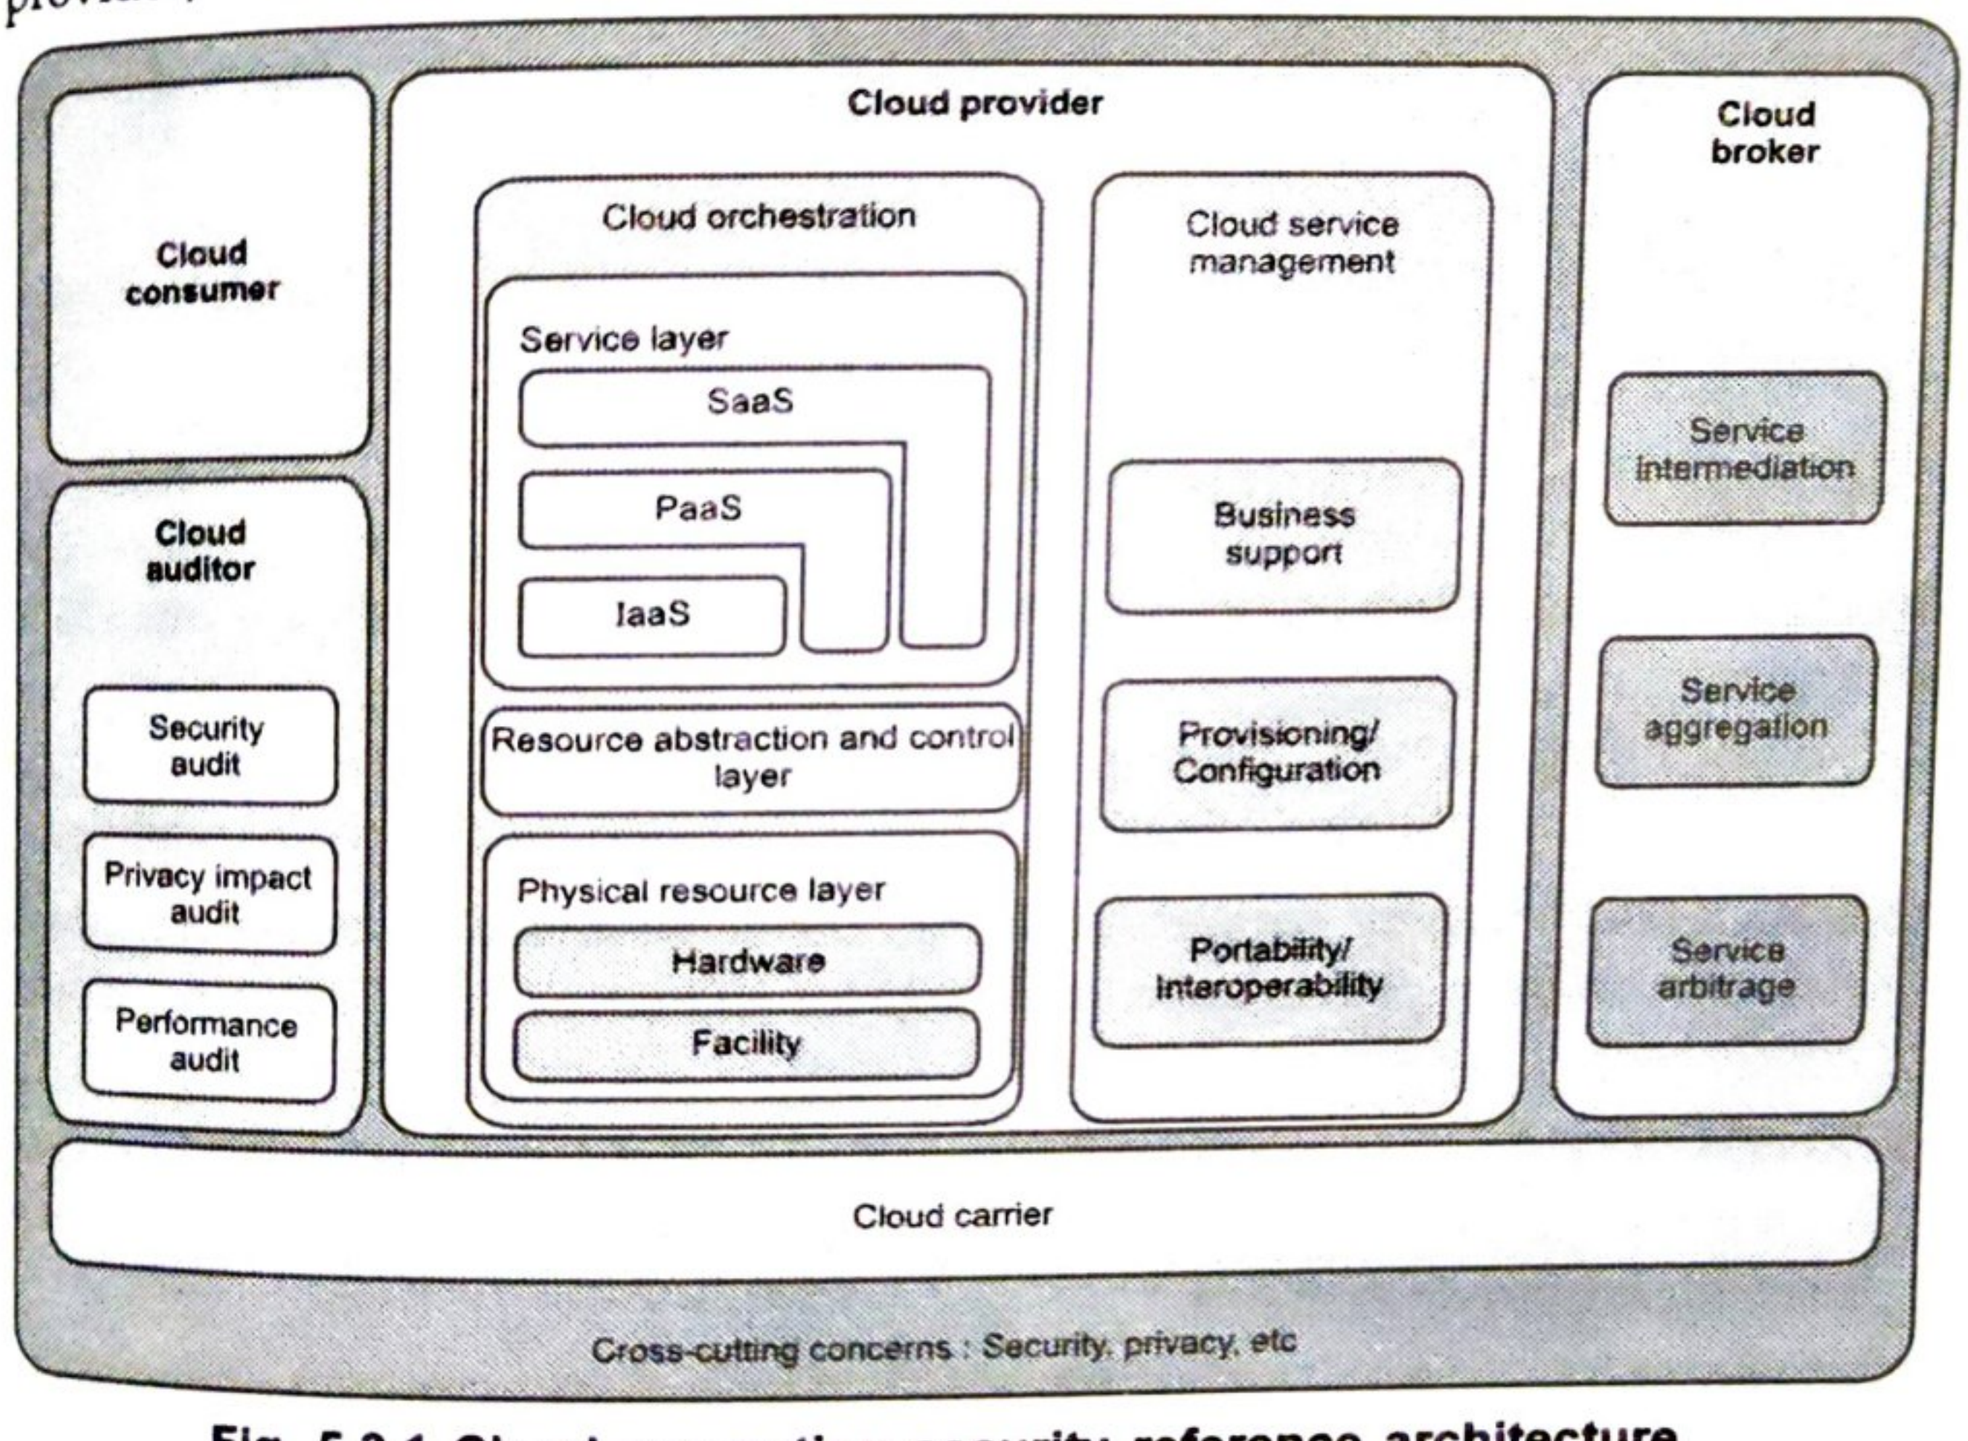

Fig. 5.3.1 Cloud computing security reference architecture

- Secure cloud computing architecture encompasses three core capabilities: confidentiality, integrity, and availability.
  1. Confidentiality is the ability to keep information secret and unreadable to the people who shouldn't have access to that data.
  2. Integrity is the idea that the systems and applications are exactly what you expect them to be and function exactly as you expect them to function.

3. Availability speaks to Denial-of-Service (DoS) attacks. Perhaps an attacker can't see or change your data. But if an attacker can make systems unavailable to you or your customers, then you can't carry out tasks that are essential to maintain your business.

Actor	Definition
Cloud Consumer	A person or organization that maintains a business relationship with and uses service form, *Cloud Providers*.
Cloud Provider	A person, organization or entity responsible for making a service available to interested parties.
Cloud Auditor	A party that can conduct an independent assessment of cloud services, information system operations, performance and security of the cloud implementation.
Cloud Broker	An entity that manages the use, performance and delivery of cloud services and negotiates relationships between *Cloud Providers* and *Cloud Consumers*.
Cloud Carrier	An intermediary that provides connectivity and transport of cloud services from *Cloud Providers* to *Cloud Consumers*.

5.3.1 General Issues Securing the Cloud

- The common security issues around cloud computing divided into four main categories :
  a) **Cloud infrastructure, platform and hosted code :** This comprises concerns related to possible virtualization, storage and networking vulnerabilities.
  b) **Data :** This category comprises the concerns around data integrity, data lock in, data remanence, provenance, and data confidentiality and user privacy specific concerns.
  c) **Access :** This comprises the concern around cloud access (authentication, authorization and access control or AAA), encrypted data communication, and user identity management.
  d) **Compliance :** Because of its size and disruptive influence, the cloud is attracting attention from regulatory agencies, especially around security audit, data location; operation trace- ability and compliance concerns.

5.3.2 Challenges to Data Security in Cloud

1. **Data residency :** Many companies face legislation by their country of origin or the local country that the business entity is operating in, requiring certain types of data to be kept within defined geographic borders. There are specific regulations that must be followed, centered around data access, management and control.

2. **Data privacy :** Business data often needs to be guarded and protected more stringently than non-sensitive data. The enterprise is responsible for any breaches to data and must be able ensure strict cloud security in order to protect sensitive information.

3. **Industry and regulation compliance :** Organizations often have access to and are responsible for data that is highly regulated and restricted.

5.3.3 Virtual Machine Security

- The virtual machines which were created have their own virtual server in which different data was processing which comes from the host machines. The service provider will ensure that the server can be protected with the help of the techniques like a firewall or other security mechanisms, and the server's database also contains some security mechanisms; and in this way, the cloud infrastructure provider provides security to the server.
- Virtualized security, or security virtualization, refers to security solutions that are software-based and designed to work within a virtualized IT environment.
- Step one in securing virtual machine security in cloud computing is to isolate the new hosted elements. In cloud-virtual security is to certify virtual features and functions for security compliance before you allow them to be deployed. Outside attacks are a real risk in virtual networking, but an insider attack is a disaster.
- To separate infrastructure management and orchestration from the service. Management APIs will always represent a major risk because they're designed to control features, functions and service behaviour. It's important to protect all such APIs, but it's critical to protect the APIs that oversee infrastructure elements that should never be accessed by service users.
- Cloud-virtual network security is to ensure that virtual network connections don't cross over between tenants or services.
- The key to virtualization security is the hypervisor, which controls access between virtual guests and host hardware. A Type 1 hypervisor is part of an operating system that runs directly on host hardware. A Type 2 hypervisor runs as an application on a normal operating system, such as Windows 10. For example : VMware ESX is a Type 1 hypervisor and VMware Workstation is Type 2.
- To secure the guest OS running in virtual machines, best practices for the protection of physical machines must be followed that include updating the OS regularly for patches and updates, using anti-virus software, securing internet and email and monitoring of guest OS regularly.

- Whenever VM is migrated from one physical machine to different, images on previous disks should be completely removed.
- Each component of virtualization layer can act as an attack vector to launch multiple attacks on the system. Attacks that target different components of virtualization environment may result in security issues such as compromise of complete Cloud infrastructure, stealing of customer data and system hacking.

Machine Imaging :

- Machine imaging is a process that is used to provide system portability, and provision and deploy systems in the cloud through capturing the state of systems using a system image.
- A system image makes a copy or a clone of the entire computer system inside a single file. The image is made by using a program called system imaging program and can be used later to restore a system image.
- For example : Amazon Machine Image (AMI) is a system image that is used in the cloud computing. The Amazon Web Services uses AMI to store copies of a virtual machine.
- An AMI is a file system image that contains an operating system, all device drivers and any applications and state information that the working virtual machine would have.
- The AMI files are encrypted and compressed for security purpose and stored in Amazon S3 (Simple Storage System) buckets as a set of 10 MB chunks.
- Machine imaging is mostly run on virtualization perform due to this it is also called as virtual appliances and running virtual machines are called **instances.**
- The AMI file system is not a standard bit-for-bit image of a system that is common to many disk imaging programs. AMI omits the kernel image and stores a pointer to a particular kernel that is part of the AWS kernel library.
- Among the choices are Red Hat Linux, Ubuntu, Microsoft Windows, Solaris and others. Files in AMI are compressed and encrypted and an XML file is written that describes the AMI archive.
- Machine images are sometimes referred to as "virtual appliances", systems that are meant to run on virtualization platforms.

University Question

1. *How machine imaging help to achieve the goal of cloud computing ?*

GTU : Summer-17, Marks 7

5.4 Identity Management and Access Control

GTU : Summer-17, 18, Winter-17, 18

- AWS Identity and Access Management (IAM) is a web service that helps you securely control access to AWS resources.
- When you first create an AWS account, you begin with a single sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS account root user and is accessed by signing in with the email address and password that you used to create the account.
- For each AWS account, you can create multiple users with different credentials. For each user, you can give different rights.
- **IAM Users** are account objects that allow an individual user to access your AWS environment with a set of credentials. You can issue user accounts to anyone you want to view or administer objects and resources within your AWS environment. Permissions can be applied individually to a user, but the best practice for permission assignments is to assign them via the use of groups.
- **IAM groups** are objects that have permissions assigned to them via policies allowing the members of the group access to specific resources. Having users assigned to these groups allows for a uniform approach to access management and control.
- **IAM roles** are again objects created within IAM which have policy permissions associated to them. However, instead of being associated with users as groups are, roles are assigned to instances at the time of launch. This allows the instance to adopt the permissions given by the role without the need to have access keys stored locally on the instance.
- Security groups are used to control access to EC2 instances. Because AWS uses flat Layer 3 networking, any instance within a user account can communicate with any other instance.
- AWS Identity Access Management allows to establish access rules and permissions to specific users and applications.
  1. Set up permissions for users and applications.
  2. Create user groups for common rules assignment.
  3. Cloud Trail allows to monitor the access.
  4. Identity federation : allow users to log in with their company credentials.
  5. Temporary security credentials, obtained by calling AWS STS APIs like AssumeRole or GetFederationToken.
- IAM policy - A document that defines the effect, actions, resources, and optional conditions.

- IAM role - An identity with permission policies, to which users can be assigned.
- IAM group - A group of users to which common policies can be attached.
- Best practices regarding security groups are as follows :
  1. Avoid using the default security group.
  2. Use meaningful names.
  3. Open only the ports you need to open.
  4. Partition applications.
  5. Restrict system administrator access.

5.4.1 Identify and Access Management

- Identity and Access Management (IAM) can help a user to manage to compute, store, manage, and application services in the AWS cloud. It uses access control techniques through which a user is familiar with which includes users, groups and permission.
- With the help of a single AWS IAM, the user can manage the customer and their needs. It provides Amazon AWS building blocks which help the user to build the applications for the security purpose.
- AWS identity and access management help the user to focus on the features and functionality which includes the security on the other side of the things. AWS IAM can also rotate access keys on the virtual machine instances.
- Functions :
  1. To manage AWS IAM users and their access.
  2. To manage Amazon IAM roles and their permissions.
  3. To manage to federate users and their permissions.

5.4.2 Security Policies

- User can manage access in AWS by creating policies and attaching them to IAM identities or AWS resources.
- A policy is an object in AWS that, when associated with an identity or resource, defines their permissions. AWS evaluates these policies when a principal entity (user or role) makes a request.
- IAM policies define permissions for an action regardless of the method that you use to perform the operation.

Types of Policy :

1. **Identity-based policies :** Attach managed and inline policies to IAM identities (users, groups to which users belong, or roles). Identity-based policies grant permissions to an identity.

2. **Resource-based policies** : Attach inline policies to resources. For example : resource-based policies are Amazon S3 bucket policies and IAM role trust policies. Resource-based policies grant permissions to a principal entity that is specified in the policy. Principals can be in the same account as the resource or in other accounts.
3. **Permissions boundaries** : Use a managed policy as the permissions boundary for an IAM entity (user or role). That policy defines the maximum permissions that the identity-based policies can grant to an entity, but does not grant permissions.
4. **Organizations SCPs** : Use an AWS Organizations service control policy (SCP) to define the maximum permissions for account members of an organization or Organizational Unit (OU). SCPs limit permissions that identity-based policies or resource-based policies grant to entities (users or roles) within the account, but do not grant permissions.
5. **Access Control Lists (ACLs)** : Use ACLs to control which principals in other accounts can access the resource to which the ACL is attached. ACLs are similar to resource-based policies, although they are the only policy type that does not use the JSON policy document structure. ACLs are cross-account permissions policies that grant permissions to the specified principal entity. ACLs cannot grant permissions to entities within the same account.
6. **Session policies** : Pass an advanced session policy when you use the AWS CLI or AWS API to assume a role or a federated user. Session policies limit the permissions that the role or user's identity-based policies grant to the session. Session policies limit permissions for a created session, but do not grant permissions.

5.4.3 IAM Abilities and Limitations

- Path names must begin and end with a forward slash (/).
- Names of users, groups, roles, policies, instance profiles, and server certificates must be alphanumeric, including the following common characters : plus (+), equal (=), comma (,), period (.), at (@), underscore (_), and hyphen (-).
- Names of users, groups and roles must be unique within the account.
- User passwords (login profiles) can contain any Basic Latin (ASCII) characters.

University Questions

1. *What do you mean by identity management and access management ?*

GTU : Summer-17, Winter-18, Marks 7

2. *What are the functionalities of Identity and Access Management (IAM) ?*

GTU : Winter-17, 18, Summer-18, Marks 4

5.5 Autonomic Security Establishing Trusted Cloud Computing

- Clients of cloud computing services currently have no means of verifying the confidentiality and integrity of their data and computation. This problem is solved by using Trusted Cloud Computing Platform (TCCP). TCCP that enables IaaS services such as Amazon EC2 to provide a closed box execution environment. TCCP guarantees confidential execution of guest VMs and allows users to attest to the IaaS provider and determine if the service is secure before they launch their VMs.
- Trust is one of the most concerned obstacles for the adoption and growth of cloud computing. Although couple of solutions have been proposed, determination of credibility of trust feedbacks is neglected in most of the cases which lead to many security failures.
- Automation of management is one of the essential characteristics of the cloud networks today. Autonomic computing is an approach to equip computer systems with capabilities to autonomously adapt their behaviour and/or structure according to dynamic operating conditions. For effective self management, a system needs context awareness, self-configuration, self-optimization, self-protecting, self-management, self-healing, anticipatory, and openness.
- A trusted computing infrastructure guarantees control of data to provide the transparency that can be verified by a customer. Trust in cloud computing requires data to be digitally signed for integrity and hence privacy preservation must be provided through efficient cryptographic techniques. Since data is physically spread across multiple data centers, a consumer would never know exactly where his data is; hence control over the data is minimal.
- Fig. 5.5.1 shows trust establishment between service provider and consumer.

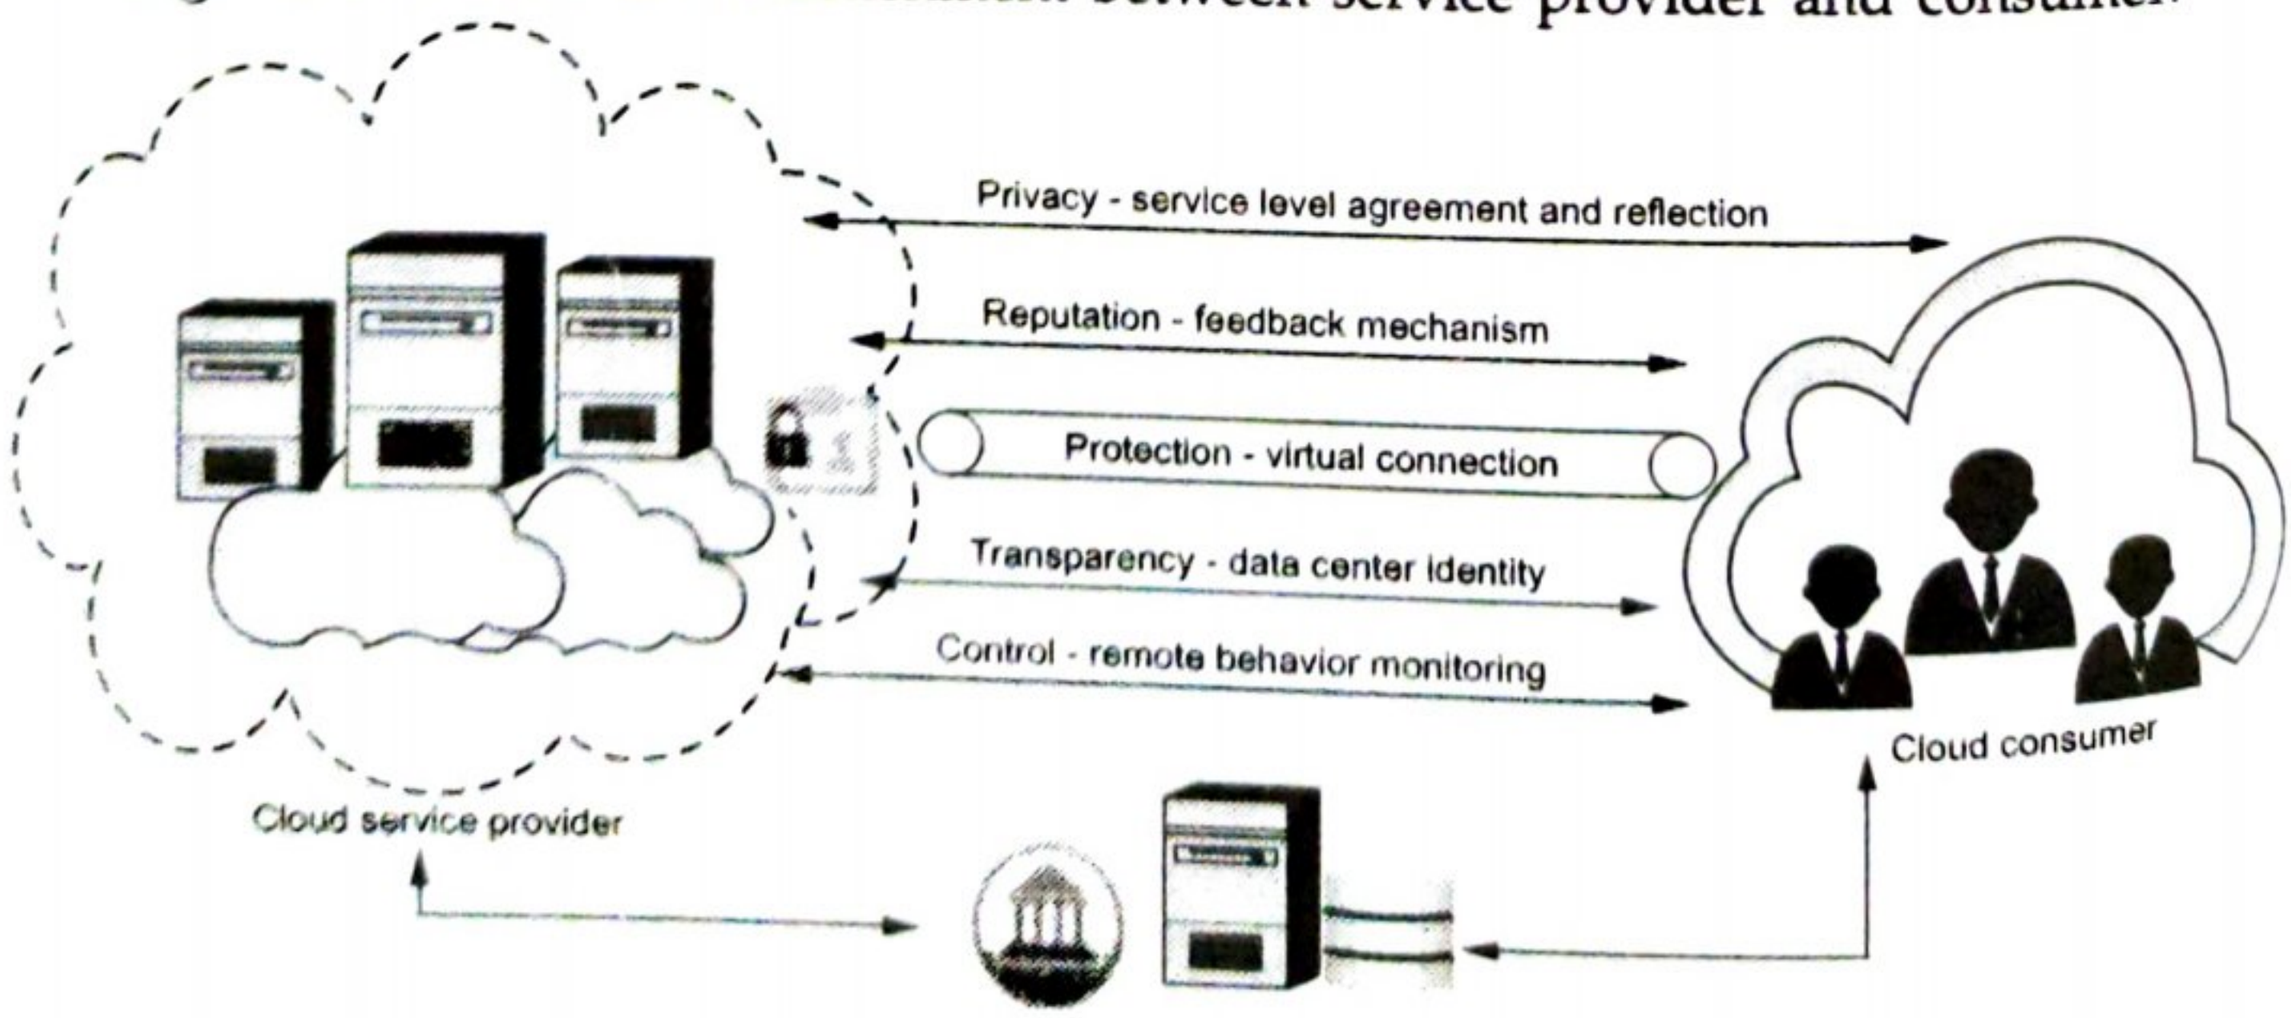

Fig. 5.5.1

- Initial trust for a service is null for the first run because reputation, feedback, and behavior are not measured. After an evaluation process to check the integrity of a system, its process and data demonstrate progress. System integrity is a binary process, indicating whether or not the system has a trustworthy execution.
- Process integrity depends on the codes executed during instantiation and should not be compromised. A modified code will yield malicious behavior on the part of the system, leading to mistrusted activity. A cryptographic process with a hash function can prevent intrusion.
- Data integrity can be ensured by checking the communication process.

5.6 Storage Area Networks

- There are hundreds of different cloud storage systems. Some have a very specific focus, such as storing Web e-mail messages or digital pictures. Others are available to store all forms of digital data.
- Some cloud storage systems are small operations, while others are so large that the physical equipment can fill up an entire warehouse. The facilities that house cloud storage systems are called data centers.
- Cloud storage systems generally -rely on hundreds of data servers. Because computers occasionally require maintenance or repair, it's important to store the same information on multiple machines. This is called redundancy.
- Without redundancy, a cloud storage system couldn't ensure clients that they could access their information at any given time.
- Most systems store the same data on servers that use different power supplies. That way, clients can access their data even if one power supply fails.
- Data centers are buildings where multiple servers and communication gear are co-located because of their common environmental requirements and physical security needs, and for ease of maintenance.
- Data centers are specialized environments that safeguard company's most valuable equipment and intellectual property.
- Data centers support the following things :
  1. Processing of users business transactions
  2. Hosting of company website
  3. Process and store intellectual property
  4. Maintain financial records
  5. Route electronic mails.

- The data center infrastructure is central to the IT architecture, from which all content is sourced or passes through. Proper planning of the data center infrastructure design is critical, and Performance, resiliency, and scalability need to be carefully considered.
- Data center equipment's environmental conditions should fall within the ranges.
- The main purpose of a data center is running the applications that handle the core business and operational data of the organization. Data Centers are the facilities that will house the equipment in order to secure, store, and exchange data.
- Networked storage devices usually fall into one of the following categories :

1. Storage Area Network (SAN) :

- Physical data storage media are connected through a dedicated network and provide block-level data storage access using industry standard protocols, such as the Small Computer System Interface (SCSI).
- The purpose of the SAN is to allow multiple servers access to a pool of storage in which any server can potentially access any storage unit.
- SAN is a network designed to transfer data from servers to targets, and it is alternative to directly attached target architecture, or to a DAS architecture, where the storage is connected to the servers on general purpose networks.
- SAN consists of a communication infrastructure, which provides physical connections; and a management layer, which organizes the connections, storage elements, and computer systems so that data transfer is secure and robust.
- SAN consists of three main components : Servers, Network infrastructure and Storage

2. Network-Attached Storage (NAS)

- Hard drive arrays are contained and managed by this dedicated device, which connects through a network and facilitates access to data using file-centric data access protocols like the Network File System (NFS) or Server Message Block (SMB).
- NAS is storage that sits on the ordinary network and is accessible by devices attached to that LAN. NAS devices provide access to file systems and as such are effectively file server appliances.
- NAS allows more hard disk storage space to be added to a network that already utilizes servers without shutting them down for maintenance and upgrades.
- NAS device does not need to be located within the server but can exist anywhere in a LAN and can be made up of multiple networked NAS devices.
- NAS systems usually contain one or more hard disks, often arranged into logical, redundant storage containers or RAID arrays.

5.6.1 Difference between NAS and SAN

NAS	SAN
Machine connected with LAN may utilize NFS, CIFS or HTTP protocol to connect to a NAS.	Server class devices that are equipped with SCSI and fibre channel adapters connect to a SAN.
File system is managed by the NAS head unit.	The SAN servers manage the file system.
Backups and mirrors are generated on files, not blocks.	Backups and mirrors require a block by block copy operation.
A NAS identifies the data by file name and byte offset, transfers file data or metadata.	SAN addresses the data by logical block numbers, and transfers the data in disk blocks.
NAS uses TCP/IP networks.	SAN uses fibre channel.

5.7 Disaster Recovery in Clouds

- Business continuity is more proactive and generally refers to the processes and procedures an organization must implement to ensure that mission-critical functions can continue during and after a disaster. BC involves more comprehensive planning geared toward long-term challenges to an organization's success.
- Disaster recovery is more reactive and comprises specific steps an organization must take to resume operations following an incident. Disaster recovery actions take place after the incident and response times can range from seconds to days.
- Disaster recovery is a piece of business continuity planning and concentrates on accessing data easily following a disaster.
- As cyberthreats increase and the tolerance for downtime decreases, business continuity and disaster recovery gain importance. These practices enable an organization to get back on its feet after problems occur and improve operations while decreasing the chance of emergencies.

Disaster recovery plan :

- Disaster recovery plan is a plan designed to recover all the vital business during a disaster with in a limited amount of time. This plan has all the procedures required to handle the emergency situations.
- A disaster recovery process should have provable recovery capability and hence it provides the most efficient method to be adopted immediately after a disaster occurs.
- Mostly the DRP has technology oriented methodologies and concentrates on getting the systems up as soon as possible, within a reasonable amount of time Recovery Time Objective (RTO) and Recovery Point Objective (RPO).

- RTO and RPO are the recovery time objective and recovery point objective, which are the targets of DRP.
- The most successful disaster recovery strategy is the one that will never be implemented; therefore, risk avoidance is a critical element in the disaster recovery process.

5.7.1 Difference between Disaster Recovery and Business Continuity Plan

Sr. No.	Disaster Recovery Plan	Business Continuity Plan
1.	Main idea : Recover from disasters.	Main idea : Continue critical business operations.
2.	Disaster recovery is data centric.	Business continuity is business centric.
3.	DR plan can be built upon a strong business continuity plan.	The business continuity process has a series of DRPs.
4.	Activities are pre-planned to react to disasters.	Planning on mitigating risk for the assets, business processes that will adversely impact company, if a disaster happens.
5.	DR plan starts with IT, not because other aspects are not important, but because IT is easiest to recover, and impact is also more.	BC plan is not an IT process; it includes the complete business as a unit.
6.	Disaster recovery is more reactive.	Business continuity is more proactive.

5.8 Multiple Choice Questions

Q.1 Pillars to SaaS-specific security are _____.

a network control
b VM management
c access management
d all of these

Q.2 The SDLC consists of _____ phases, and there are steps unique to the secure software development life cycle in each of phases.

a 4
b 6
c 8
d 9

Q.3 The key to virtualization security is the _____, which controls access between virtual guests and host hardware.

a hypervisor
b zen
c virtual machine
d none

Q.4 A type 1 hypervisor is part of an _____ that runs directly on host hardware.

a hardware
b software
c operating system
d all of these

Q.5 NAS stands for _____.

a Network Attached Storage
b New Attached Storage
c Network Area Storage
d Network Attached Standard

Answer Keys for Multiple Choice Questions :

Q.1	d	Q.2	c
Q.3	a	Q.4	c
Q.5	a		

❑❑❑

6 Cloud Middleware

Syllabus

OpenStack, Eucaluptus, Windows Azure, CloudSim, EyeOs, Aneka, Google App Engine.

Contents

6.1 OpenStack

- OpenStack is a recently open-sourced, IaaS cloud-computing platform founded by Rackspace Hosting and NASA, and is used widely in industry
- OpenStack is an open-source cloud platform. OpenStack software controls large pools of compute, storage, and networking resources throughout a data center, all managed by a dashboard that gives administrators control while empowering their users to provision resources through a web interface.
- To produce the ubiquitous Open-Source cloud computing platform that will meet the needs of public and private cloud providers regardless of size, by being simple to implement and massively scalable.
- Components of OpenStack are as follows :
  1. Horizon - Dashboard : It provides a modular web-based user interface for all the OpenStack services. With this web GUI, user can perform most operations on your cloud like launching an instance, assigning IP addresses and setting access controls.
  2. Keystone is a framework for authentication and authorization for all the OpenStack services. It handles API requests as well as providing configurable catalog, policy, token and identity services. Keystone is a framework for authentication and authorization for all the OpenStack services.
  3. Nova : It provides virtual servers upon demand. Nova is the most complicated and distributed component of OpenStack. A large number of processes cooperate to turn end user API requests into running virtual machines.
  4. Glance - Image Store : It provides discovery, registration and delivery services for disk and server images.
  5. Quantum - Network : It provides " network connectivity as a service " between interface devices managed by other OpenStack services. The service works by allowing users to create their own networks and then attach interfaces to them. Quantum has a pluggable architecture to support many popular networking vendors and technologies.
  6. Cinder allows block devices to be exposed and connected to compute instances for expanded storage & better performance.
  7. Object store allows you to store or retrieve files. It provides a fully distributed, API-accessible storage platform that can be integrated directly into applications or used for backup, archiving and data retention

6.2 Windows Azure

- Windows Azure is a cloud computing platform and infrastructure, created by Microsoft, for building, deploying and managing applications and services through a global network of Microsoft - managed data centers.
- Azure queue storage is a service for storing large numbers of messages that can be accessed from anywhere in the world via authenticated calls using HTTP or HTTPS. A single queue message can be up to 64 KB in size, and a queue can contain millions of messages, up to the total capacity limit of a storage account.
- Azure is a virtualized infrastructure to which a set of additional enterprise services has been layered on top, including, a virtualization service called Azure AppFabric that creates an application hosting environment. AppFabric is a cloud-enabled version of the .NET framework.
- Windows Azure is Microsoft's application platform for the public Cloud. Applications can be deployed on to Azure in various models
- Windows Azure is used to :
  1. Build a web application that runs and stores its data in Microsoft data centers.
  2. Store data while the applications that consume this data run on premise (outside the public Cloud).
  3. Create virtual machines to develop and test, or run SharePoint and other out-of-the-box applications.
  4. Develop massively scalable applications with many users.
  5. Offer a wide range of services
- Azure has three components : compute, storage and fabric
  1. **Compute :** Windows Azure provides a hosting environment for managed code. It provides a computation service through roles. Windows Azure supports three types of roles :
     a) Web roles used for web application programming and supported by IIS7.
     b) Worker roles are also used for background processing of web roles.
     c) Virtual Machine (VM) roles are generally used for migrating windows server applications to Windows Azure in an easy way.
  2. **Storage :** Windows Azure provides storage in the cloud. It provides four different types of storage services :
     a) Queues for messaging between web roles and worker roles.
     b) Tables for storing structural data.
     c) BLOBs (Binary Large Objects) to store text, files or large data.

d) Windows Azure Drives (VHD) to mount a page blob. They can easily be downloaded and uploaded via blobs.

3. AppFabric provides infrastructure services for developing, deploying and managing Windows Azure application. It provides five services: Service bus, Access, Caching, Integration and Composite.

- Fig. 6.2.1 shows Windows Azure platform architecture.

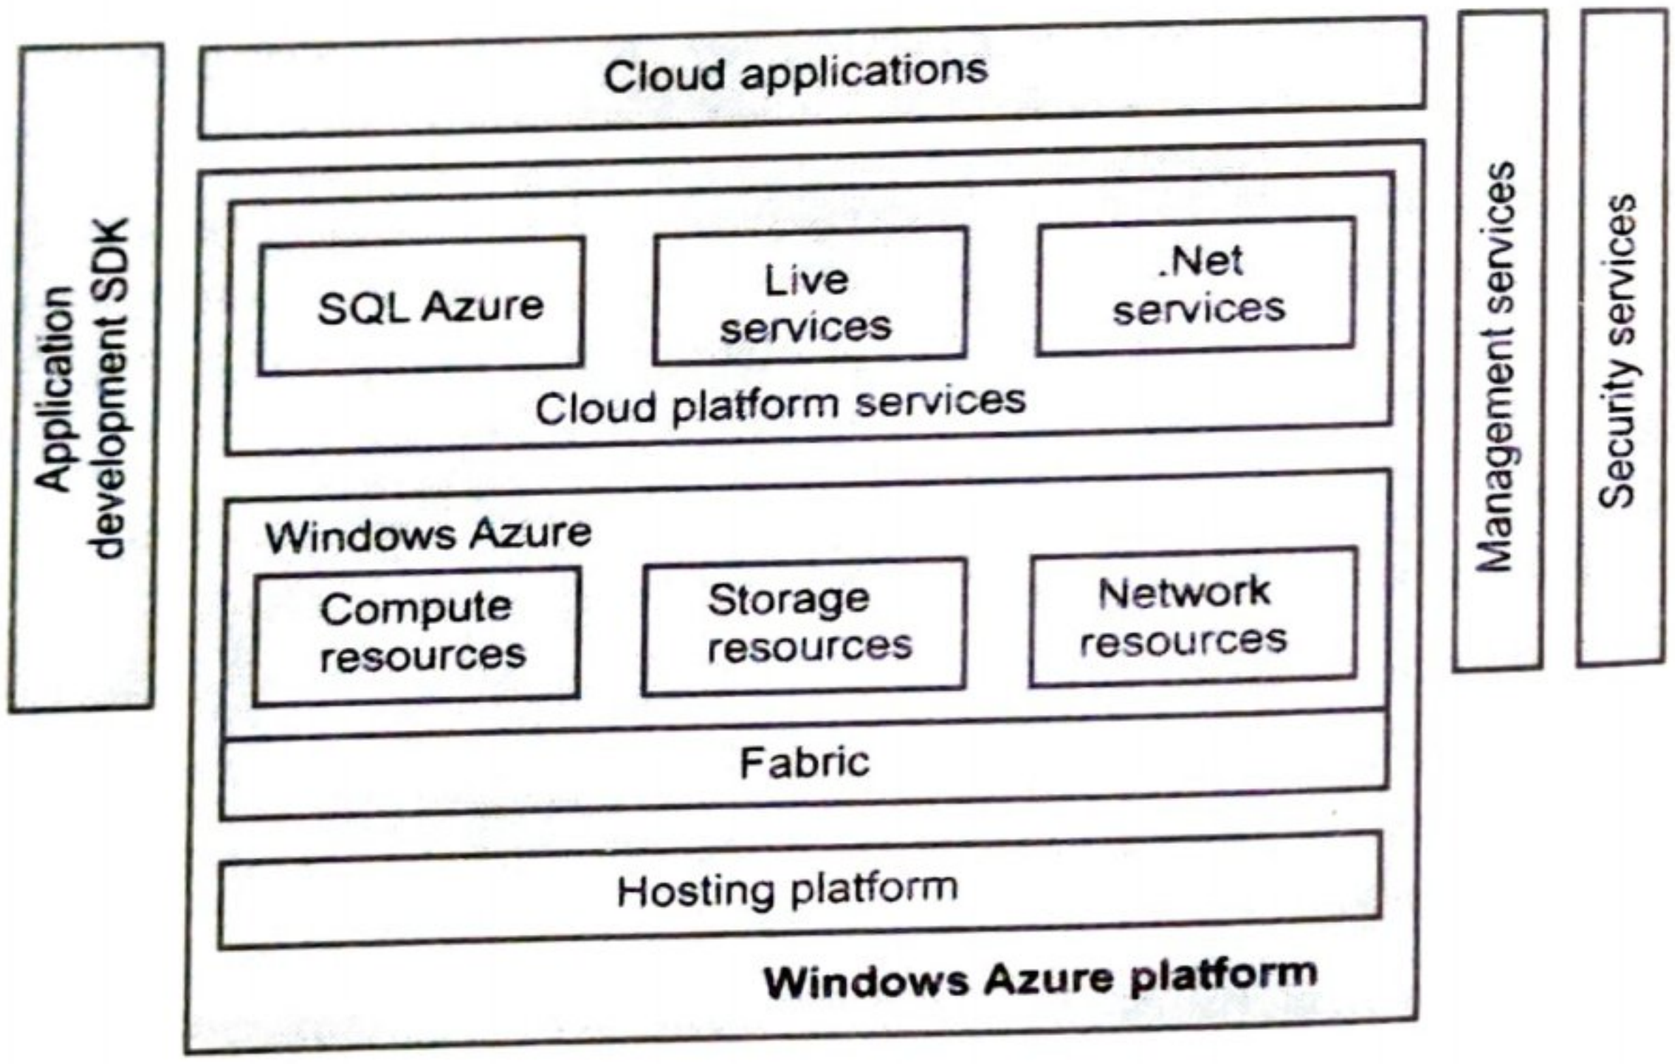

Fig. 6.2.1 Windows Azure platform architecture

- Microsoft Azure is a cloud computing service created by Microsoft for building, testing, deploying and managing applications and services through a global network of Microsoft-managed data centers.
- It provides software as a service (SaaS), platform as a service and infrastructure as a service and supports many different programming languages, tools and frameworks, including both Microsoft-specific and third-party software and systems.
- Windows Azure provides resources and services for consumers. For example, hardware is abstracted and exposed as compute resources.
- Physical storage is abstracted as storage resources and exposed through very well-defined interfaces.
- A common windows fabric abstracts the hardware and the software and exposes virtual compute and storage resources.

- Each instance of an application is automatically managed and monitored for availability and scalability.
- If an application goes down, the Fabric is notified and a new instance of the application is created. Because virtualization is a key element in cloud computing, no assumption must be made on the state of the underlying hardware hosting the application.
- Advantages of Microsoft Azure
  1. Microsoft Azure offers high availability
  2. It offers you a strong security profile
  3. It is a cost-effective solution for an IT budget.
  4. Azure allows you to use any framework, language, or tool.
  5. Azure allows businesses to build a hybrid infrastructure.

6.3 CloudSim

- CloudSim is an extensible simulation toolkit or framework that enables modeling, simulation and experimentation of Cloud computing systems and application providing environments. The CloudSim toolkit supports both system and behavior modeling of Cloud system components.
- By using CloudSim, developers can focus on specific systems design issues that they want to investigate, without getting concerned about details related to cloud-based infrastructures and services.
- CloudSim is a library for cloud computing simulation in Java language
- Features :
  1. Energy-aware computational resources.
  2. Support for modelling and simulation of large-scale cloud computing data centers.
  3. Support for data center network topologies and message-passing applications.
  4. Support for dynamic insertion of simulation elements, stop and resume of simulation.
  5. Support for user-defined policies for allocation of hosts to virtual machines and policies for allocation of host resources to virtual machines.
- Fig. 6.3.1 shows cloudSim.
- CloudSim framework consists of following components :
  1. **Regions** : It models geographical regions in which cloud service providers allocate resources to their customers.

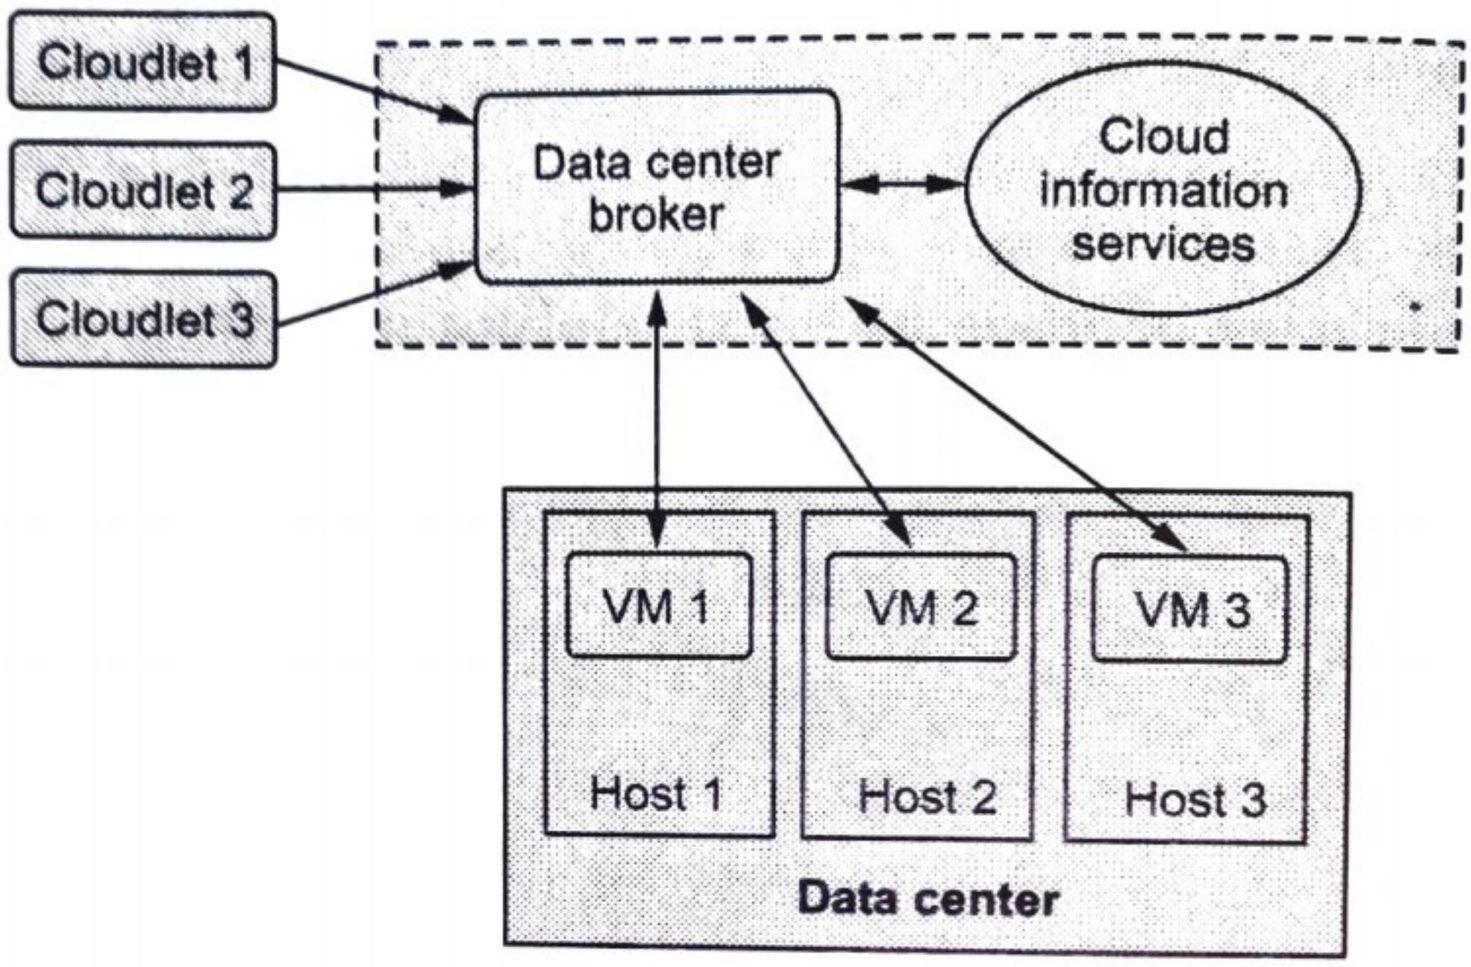

Fig. 6.3.1 : CloudSim

2. **Data centres :** Data centres is composed of a set of hosts and it is responsible for managing virtual machines
3. **Hosts :** It models physical resources. A host has a defined policy for provisioning memory, processing elements, and bandwidth to virtual machines
4. **The user base :** It models a group of users considered as a single unit in the simulation, and its main responsibility is to generate traffic for the simulation.
5. **Cloudlet :** It specifies the set of user requests. It contains the application ID, name of the user base that is the originator to which the responses have to be routed back, as well as the size of the request execution commands, and input and output files. It models the cloud-based application services.
6. **Service broker :** The service broker decides which data centre should be selected to provide the services to the requests from the user base.
7. **VMM allocation policy :** It models provisioning policies on how to allocate VMs to hosts.
8. **VM scheduler :** It models the time or space shared, scheduling a policy to allocate processor cores to VMs.

6.4 EyeOs

- EyeOS is free Cloud Computing Operating System software which let you access all your necessary files, folders, office, calendar, contacts and much more anywhere in the world.
- EyeOS is an open source web desktop following the cloud computing concept. It is mainly written in PHP, XML and JavaScript.

- The eyeOS desktop looks like any other operating system that you would come across. It can be customized on the basis of themes, though the looks of Windows Aero are obviously not a possibility. At present, the eyeOS system supports translations in 30 languages.
- Goals for eyeOS include :
  a) Being able to work from everywhere, regardless of whether or not you are using a full-featured, modern computer, a mobile gadget, or a completely obsolete PC.
  b) Sharing resources easily between different work centers at company or working from different places and countries on the same projects.
  c) Always enjoying the same applications with the same open formats and forgetting the usual compatibility problems between office suites and traditional operating systems.
  d) Being able to continue working if you have to leave your local computer or if it just crashes, without loosing data or time: Just log in to your eyeOS from another place and continue working.
- EyeOS includes the following features :
  1. **Desktop :** The desktop is similar to that of a regular operating system.
  2. **Office-related tasks :** EyeOS supports MS Office and OpenOffice.org documents, spreadsheets and presentations.
  3. **System and file management :** Uploading/downloading multiple files to the cloud, compressing in ZIP/TAR formats, and a dedicated picture viewer for slide-shows is a decent plus.
  4. **Networking :** The beauty of cloud computing is the fact that networking is at its core. There is an internal FTP client, messaging client, bulletin board and an RSS feed reader that comes bundled by default.

6.5 Aneka

- Aneka is one of platform that is used to build, accelerate and manage distributed applications with the help of .NET framework. Aneka provides developers with a rich set of APIs for transparently exploiting such resources and expressing the business logic of applications by using the preferred programming abstractions.
- Aneka based computing cloud is a collection of physical and virtualized resources connected through a network, which are either the Internet or a private intranet. Each of these resources hosts an instance of the Aneka Container representing the runtime environment where the distributed applications are executed.

- The services are broken up into fabric, foundation, and execution services. Fabric services directly interact with the node through the Platform Abstraction Layer (PAL) and perform hardware profiling and dynamic resource provisioning.
- Foundation services identify the core system of the Aneka middleware, providing a set of basic features to enable Aneka containers to perform specialized and specific sets of tasks. Execution services directly deal with the scheduling and execution of applications in the Cloud.
- ANEKA is available at PaaS in cloud environment. It means that it provides programming application programming interfaces for developing distributed applications and virtual execution environment in which the applications developed as per API can be made to run.
- Aneka is a market oriented Cloud development and management platform with rapid application development and workload distribution capabilities.
- Fig. 6.5.1 shows Aneka framework architecture.

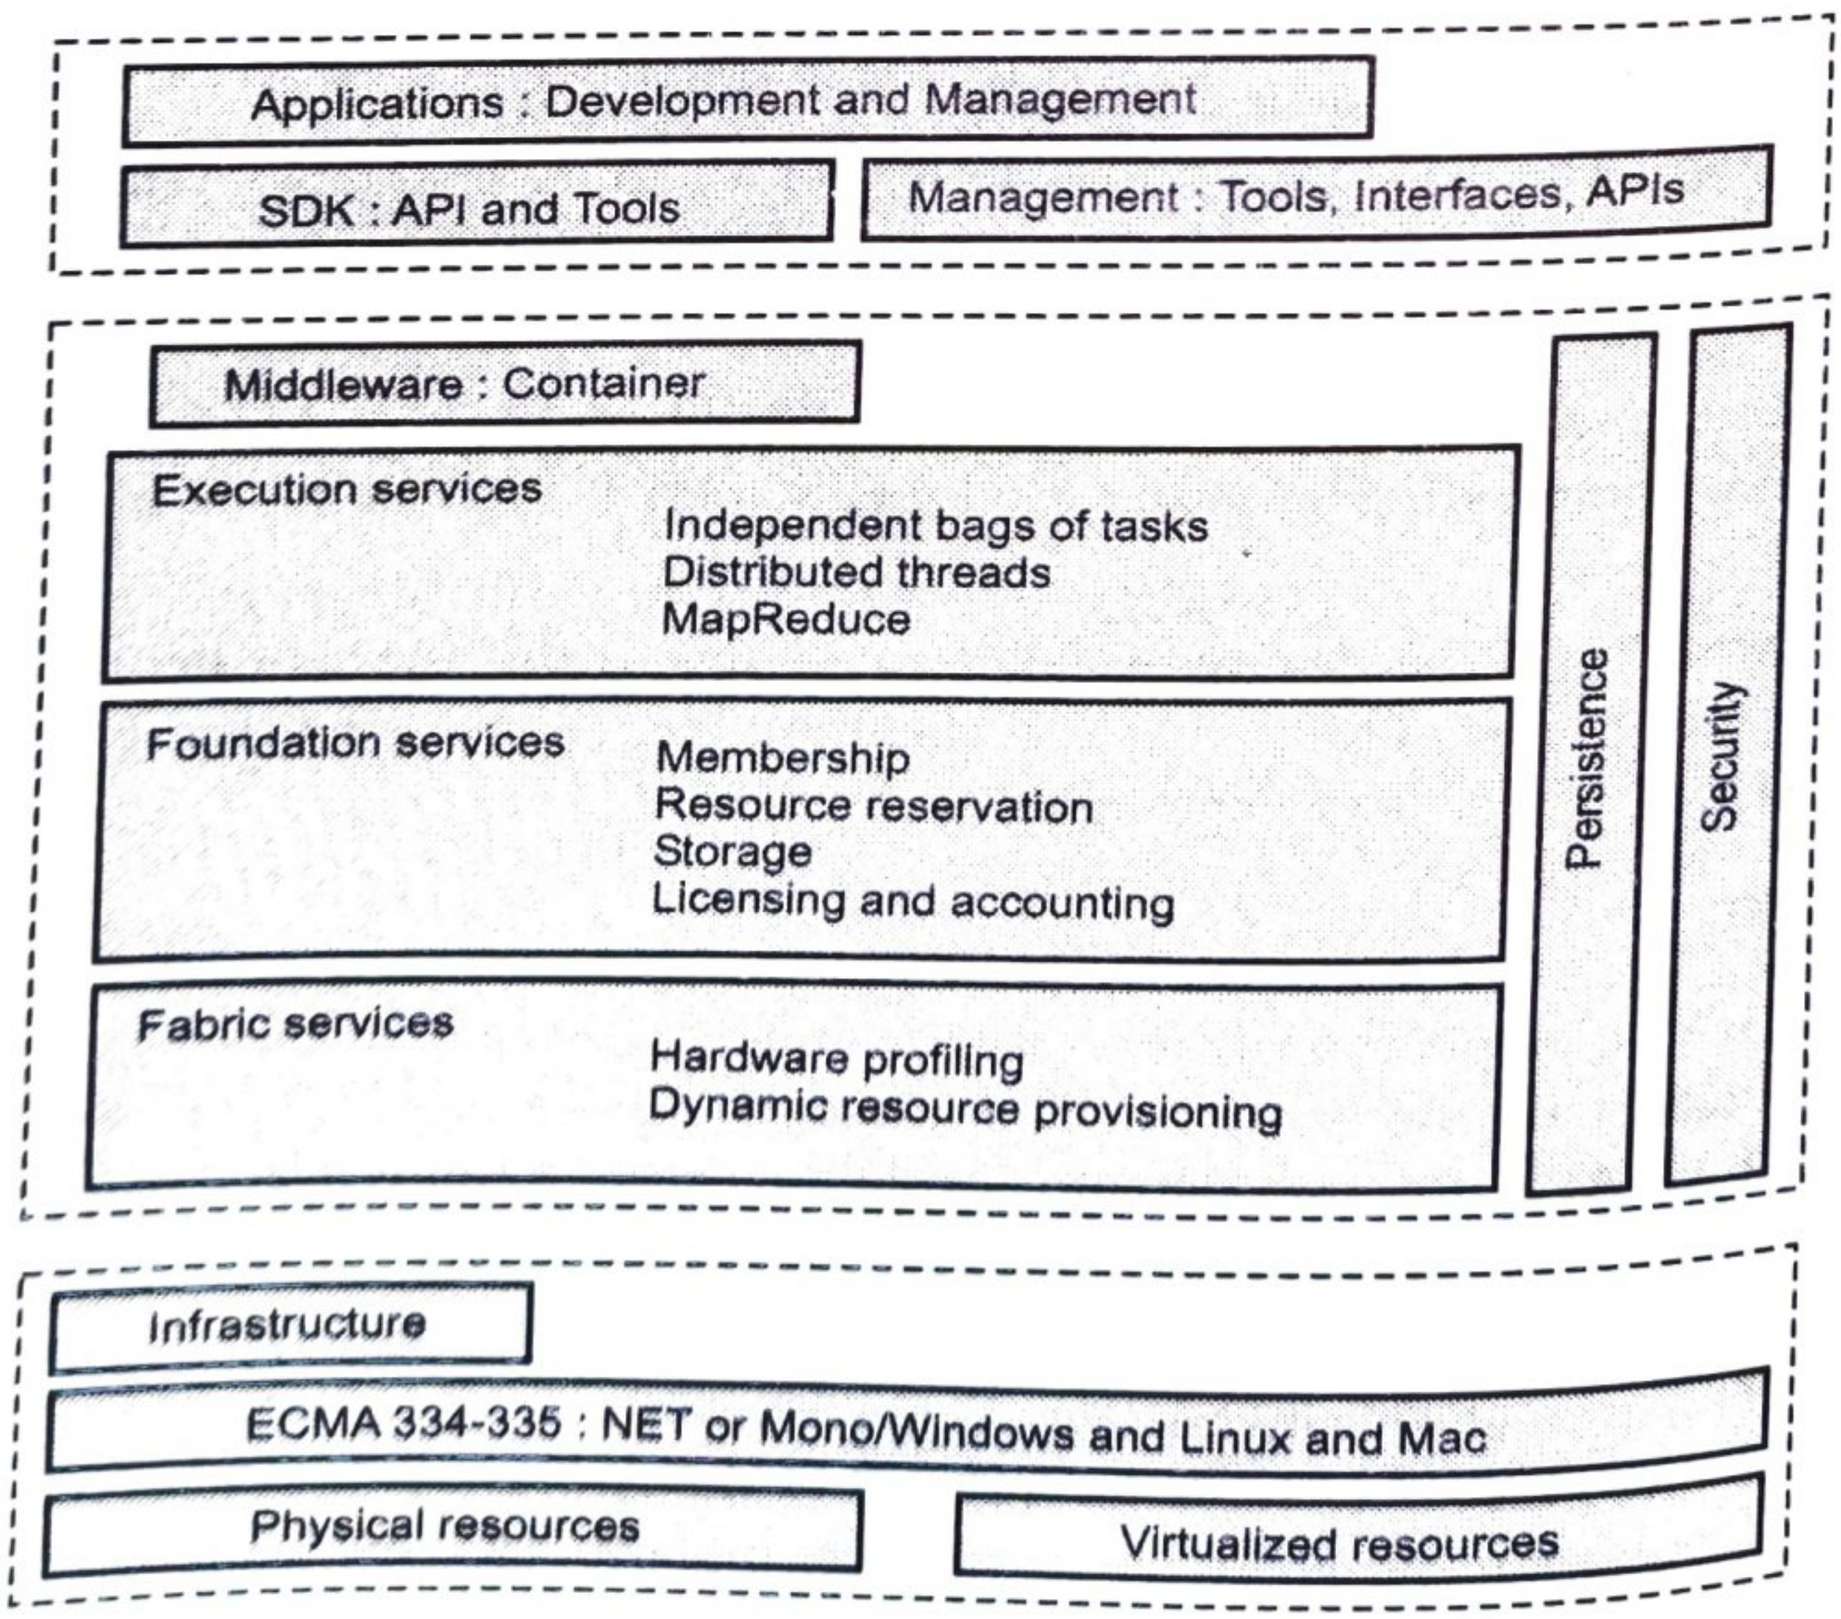

Fig. 6.5.1 Aneka framework architecture

- The container is the building block of the middleware and represents the runtime environment for executing applications; it contains the core functionalities of the system and is built up from an extensible collection of services that allow administrators to customize the Aneka cloud.

- Aneka implements a service-oriented architecture and services are the fundamental components of an Aneka Cloud. Services operate at container level.
- Aneka container can be classified into three major categories : Fabric services, Foundation services and Application services.
- There are three classes of services that characterize the container :
  1. **Execution services** : They are responsible for scheduling and executing applications. Each of the programming models supported by Aneka defines specialized implementations of these services for managing the execution of a unit of work defined in the model.
  2. **Foundation services** : These are the core management services of the Aneka container. They are in charge of metering applications, allocating resources for execution, managing the collection of available nodes and keeping the services registry updated.
  3. **Fabric services** : They constitute the lowest level of the services stack of Aneka and provide access to the resources managed by the cloud.
- Aneka also provides a tool for managing the cloud, allowing administrators to easily start, stop and deploy instances of the Aneka container on new resources and then reconfigure them dynamically to alter the behavior of the cloud.
- Applications managed by the Aneka container can be dynamically mapped to heterogeneous resources, which can grow or shrink according to the application's needs. This elasticity is achieved by means of the resource provisioning framework, which is composed primarily of services built into the Aneka fabric layer.

6.6 Google App Engine

- Google App Engine (GAE) is a Platform as a Service cloud computing platform for developing and hosting web applications in Google-managed data centers.
- Google App Engine is a way to write your own web applications and have them hosted on Google servers. It enables developers to build their web applications on the same scalable system that power Google applications.
- An app is a piece of software which can run on the computer, internet, phone or any other electronic device. Google refers to their online services as Apps. They also sell a specific suite of services known as Google Apps.

- Google's providing both SaaS and PaaS solutions in cloud computing. Some of the examples for SaaS solutions including Google Apps which including Gmail, Doc, etc. and PaaS includes Google App engine.
- Services provided by App engine includes :
  a) Platform as a Service (PaaS) to build and deploy scalable applications
  b) Hosting facility in fully-managed data centers
  c) A fully-managed, flexible environment platform for managing application server and infrastructure.
  d) Support in the form of popular development languages and developer tools.
- **Major feature of Google App Engine :**
  1. Automatic scaling and load balancing.
  2. Authentication using Google Accounts API.
  3. Provides dynamic web services based on common standards.
  4. Integration with other Google Cloud Services and API.
  5. Support persistent storage, with query access sorting and transaction management features.
- Google App engine offers users the ability to build and host web applications on Google's infrastructure.

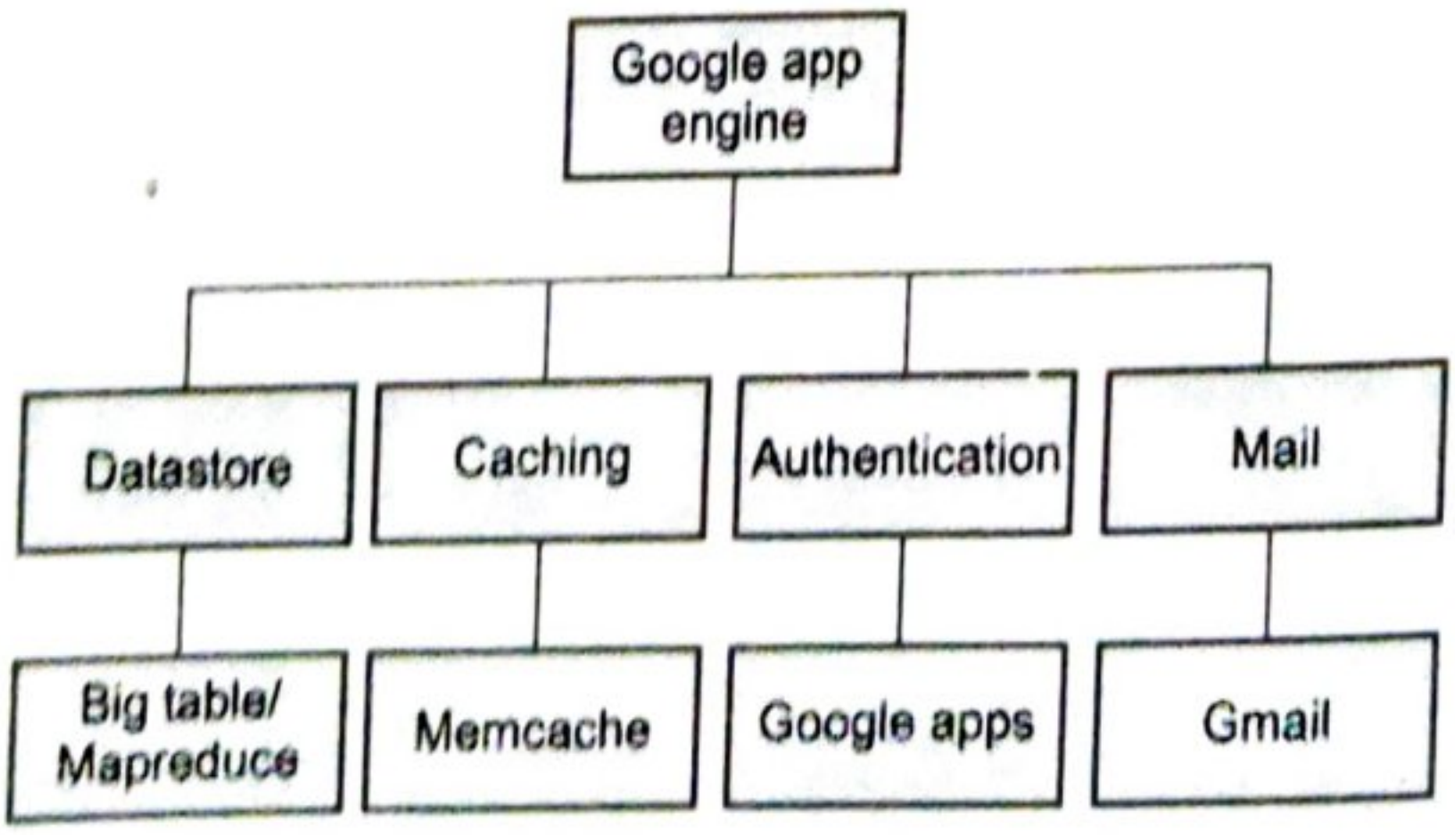

Fig. 6.6.1

- The App Engine offers a number of services that enable you to perform several common operations when managing your application. The following APIs are available to access these services :
  1. **Mail :** Using the mail API, the developers can send email messages.
  2. **Memcache :** The Memcache service gives the users the benefit of working efficiently by providing high retrieval speed, even when multiple users access the same application at the same instance of time.

3. **Image manipulation :** The Image service allows you to manipulate images of your application. With the use of this API, you can resize, crop, rotate and flip images in JPEG and PNG formats.

- In the PaaS space Google is a key player. App Engine is a platform to create, store and run applications on Google's servers using development languages as java and python.
- App Engine includes tools for managing the data store, monitoring the site and its resource consumption and debugging and logging. A user can serve the app from his own domain name using Google Apps.
- **Key features of GAE programming mode using java and python.**
- The Google App engine Software Development Kit (SDK) provides Java and Python programming languages.
- The languages have their own web server application that contains all Google App Engine services on a local computer. The web server also simulates a secure sandbox environment.
- The Google App engine SDK has APIs and libraries including the tools to upload applications. The architecture defines the structure of applications that run on the Google App engine.

1. **Python :**
   - The Google App engine allows implementation of applications using python programming language and running them on its interpreter.
   - The Google App engine provides rich APIs and tools for designing web applications, data modeling, managing, accessing apps data, support for mature libraries and frameworks like Django.
   - The main characteristics of Google App engine are its DataStore, configuration file app.yaml and how it serves an application.

2. **Java :**
   - The Google App engine provides tools and APIs required for the development of web applications that run on the Google App engine Java run time.
   - The application interacts with the environment using servlets and web technologies like Java Server Pages (JSPs) which can be developed using Java6.
   - The GAE environment uses Java SE Runtime JRE platform 6 and libraries which the applications can access using APIs.
   - Java SDK has implementations for Java Data Objects (JDO) and Java Persistence (JPA) interface.

- To exchange email messages with Google App engine, it provides the Google App Engine mail service through the Java Mail API.
- Support for other languages like JavaScript, Ruby or Scalar is also provided by Google App engine with the use of JVM compatible compilers and interpreters.
- When Google App engine gets a web request that corresponds to the URL mentioned in the applications deployment descriptor it invokes a servlet corresponding to the request and uses Java Servlets API to provide requested data and accepts response data.
- Google App engine makes it easy to build an applications that runs reliably, even under heavy load and with large amounts of data.
- App engine includes the below features :
  a) Dynamic web serving, with full support for common web technologies.
  b) Persistent storage with queries, sorting and transactions.
  c) Automatic scaling and load balancing.
  d) APIs for authenticating users and sending email using Google accounts.
  e) Scheduled tasks for triggering events at specified times and regular intervals.

6.7 Multiple Choice Questions

Q.1 OpenStack is a recently open-sourced, _____ cloud-computing platform founded by Rack space Hosting and NASA, and is used widely in industry.

a IaaS
b SaaS
c PaaS
d IIaaS

Q.2 The Google App engine Software Development Kit provides _____ and _____ programming languages.

a R, Java
b Java, Python
c Object oriented, SQL
d C++, Java

Q.3 Google App engine is a _____ cloud computing platform for developing and hosting web applications in Google-managed data centers.

a Software as a Service
b Infrastructure as a Service
c Host as a Service
d Platform as a Service

Q.4 EyeOS is an open-source _____ desktop following the cloud computing concept. It is mainly written in PHP, XML and JavaScript.

a Personal
b URL
c web
d single

Q.5 CloudSim is a _____ for cloud computing simulation in Java language.

a package
b library
c source file
d service

Q.6 Windows Azure is Microsoft's application platform for the _____ cloud.

a hybrid
b private
c public
d all of these

Q.7 Aneka container can be classified in to following major categories :

a Fabric services
b Foundation services
c Application services
d All of these

Answer Keys for Multiple Choice Questions :

Q.1	a	Q.2	b
Q.3	d	Q.4	c
Q.5	b	Q.6	c
Q.7	d		

❑❑❑

7 Cloud Based Case-Studies

Syllabus

Overview of Cloud services, Designing Solutions for the Cloud, Implement & Integrate Solutions, Emerging Markets and the Cloud, Tools for Building Private Cloud: IaaS using Eucalyptus, PaaS on IaaS - AppScale.

Contents

7.1 Overview of Cloud Services

- Choosing the right service model is a critical success factor for delivering cloud-based solutions. In order to choose the right service model or combination of service models, one must fully understand what each service model is and what responsibilities the cloud service providers assume versus the responsibilities the cloud service consumer assumes.
- Cloud based services provide information technology as a service over the Internet or dedicated network, with delivery on demand, and payment based on usage. Cloud based services range from full applications and development platforms, to servers, storage and virtual desktops.
- Corporate and government entities utilize cloud-based services to address a variety of application and infrastructure needs such as CRM, database, compute, and data storage.
- Clouds provide services at different levels (IaaS, PaaS, SaaS). The amount of control available to users and developers decreases with the level of abstraction. According to their deployment model, clouds can be categorized into public and private clouds.
- Fig. 7.1.1 shows cloud computing services.

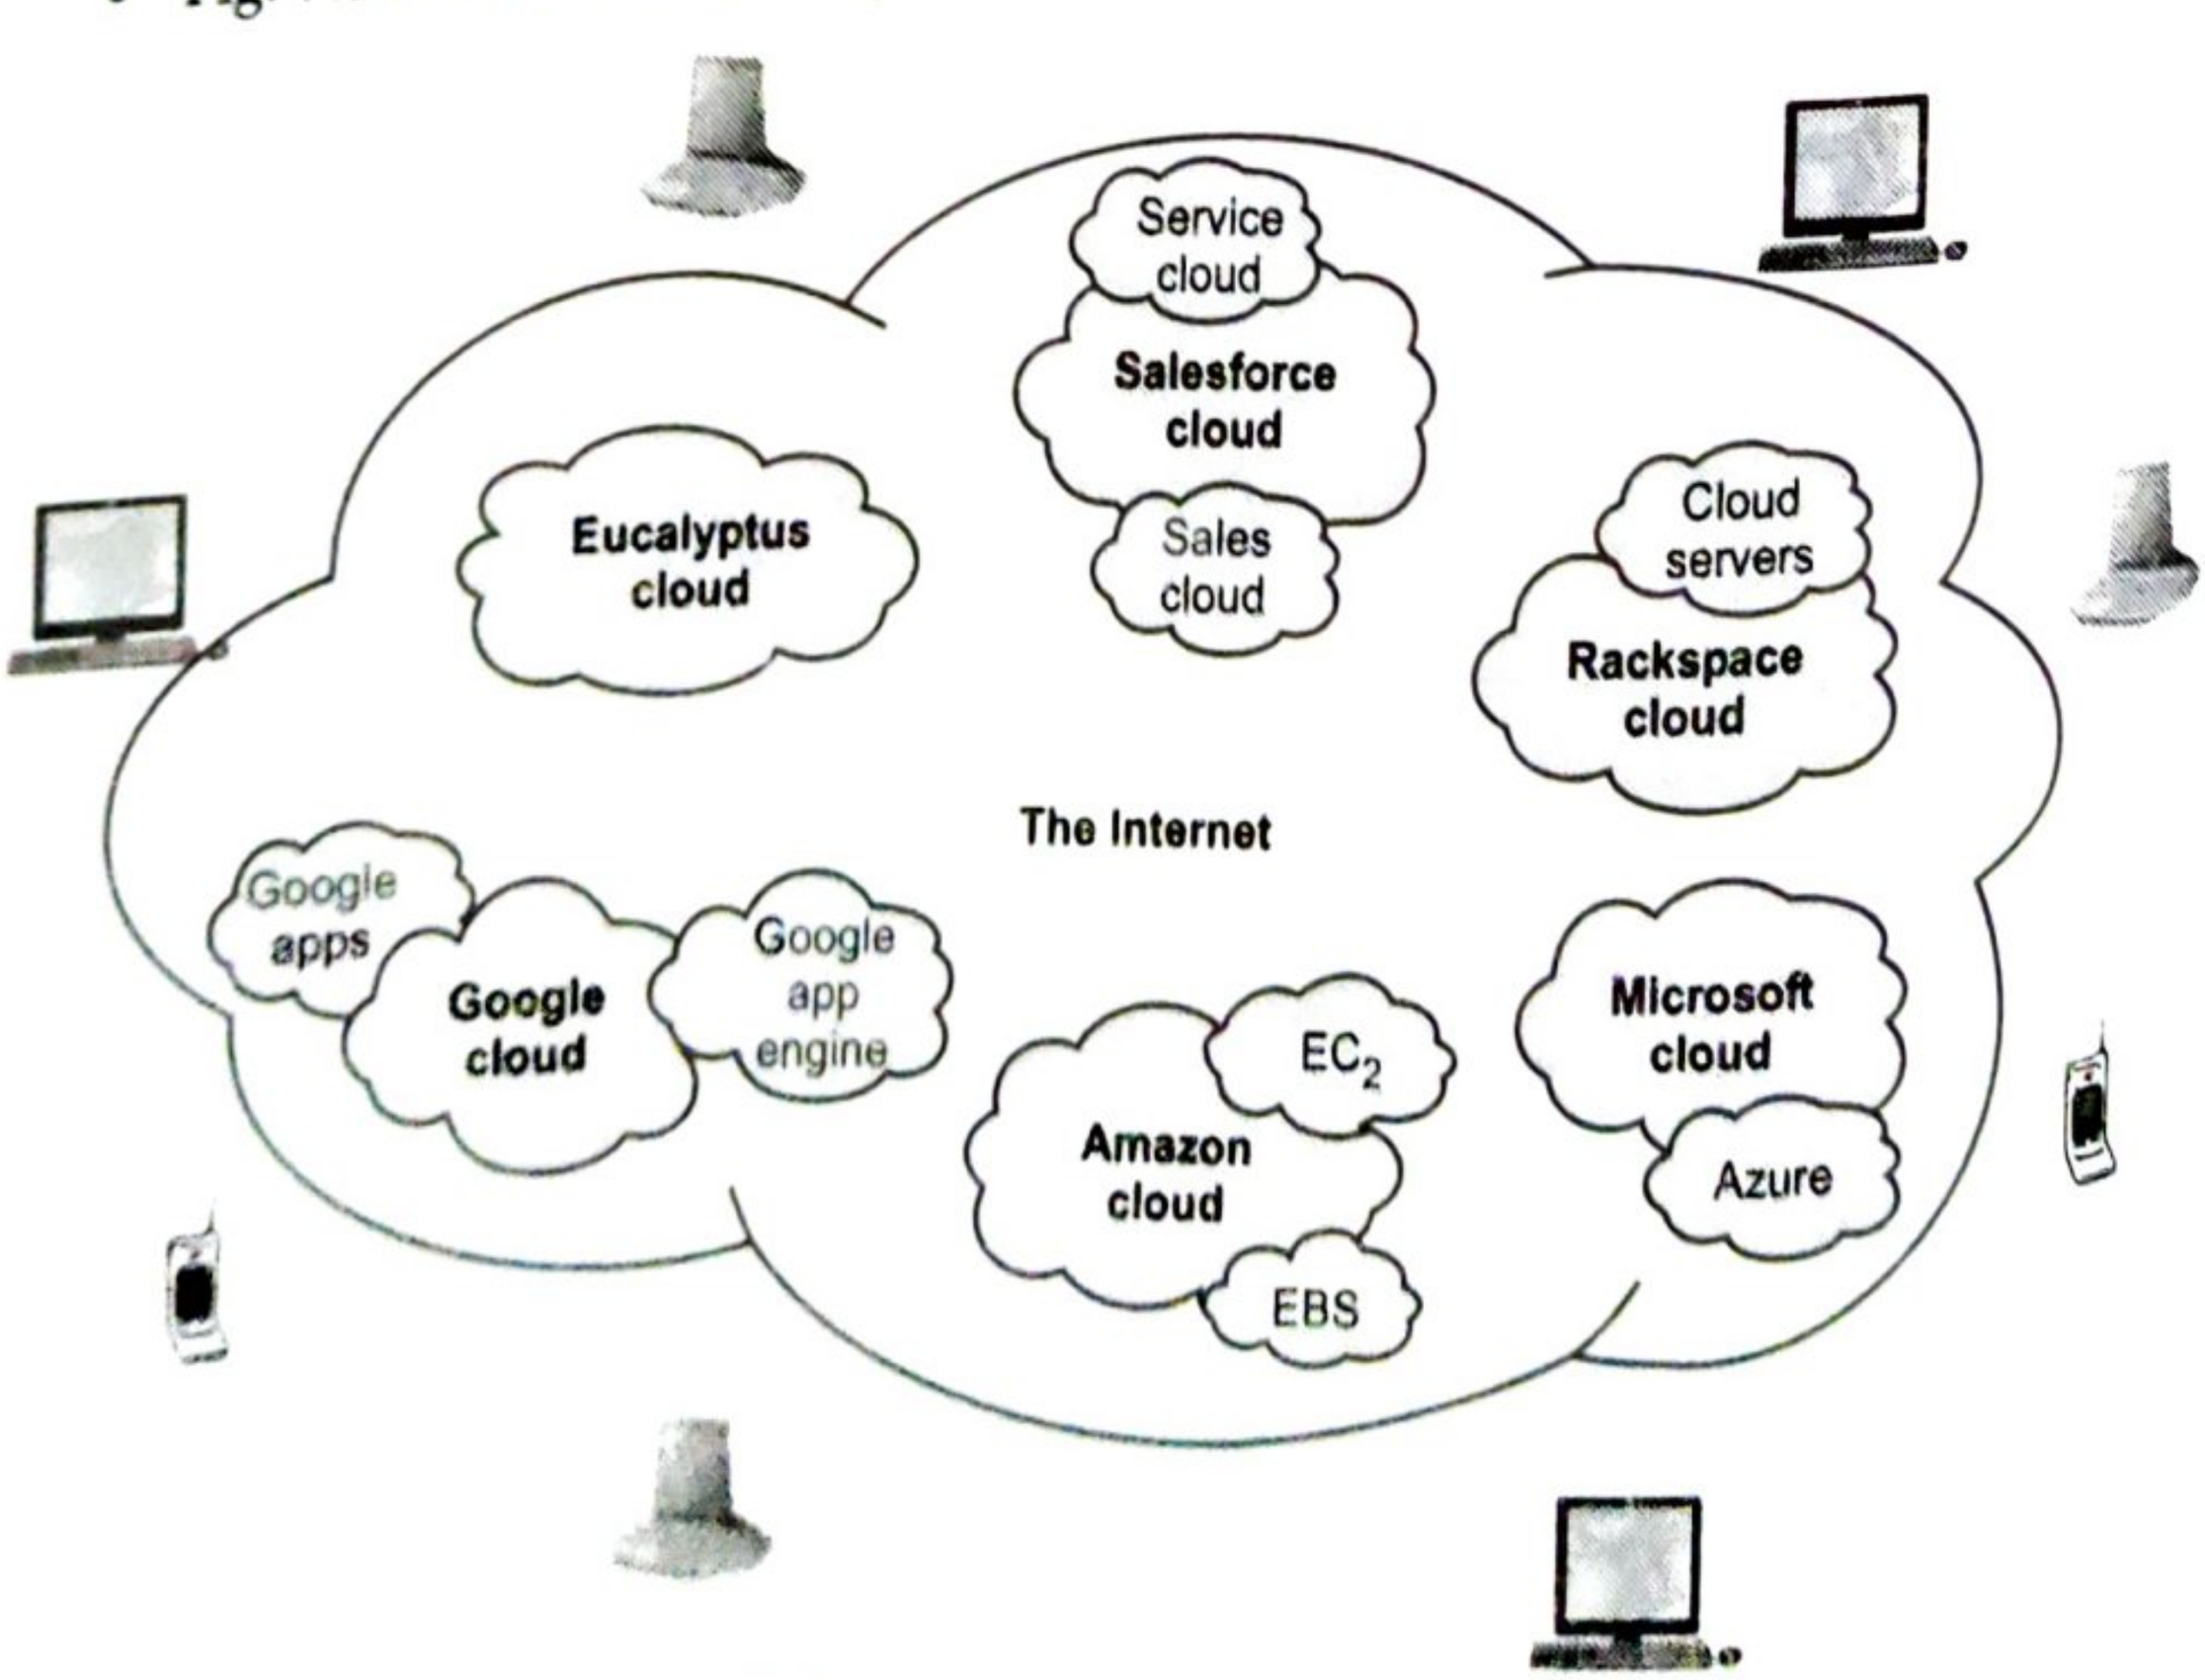

Fig. 7.1.1 Design of cloud computing services

- There are three cloud service models : Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS). Each cloud service model provides a level of abstraction that reduces the efforts required by the service consumer to build and deploy systems.
- Fig. 7.1.2 shows main categories of cloud computing services.

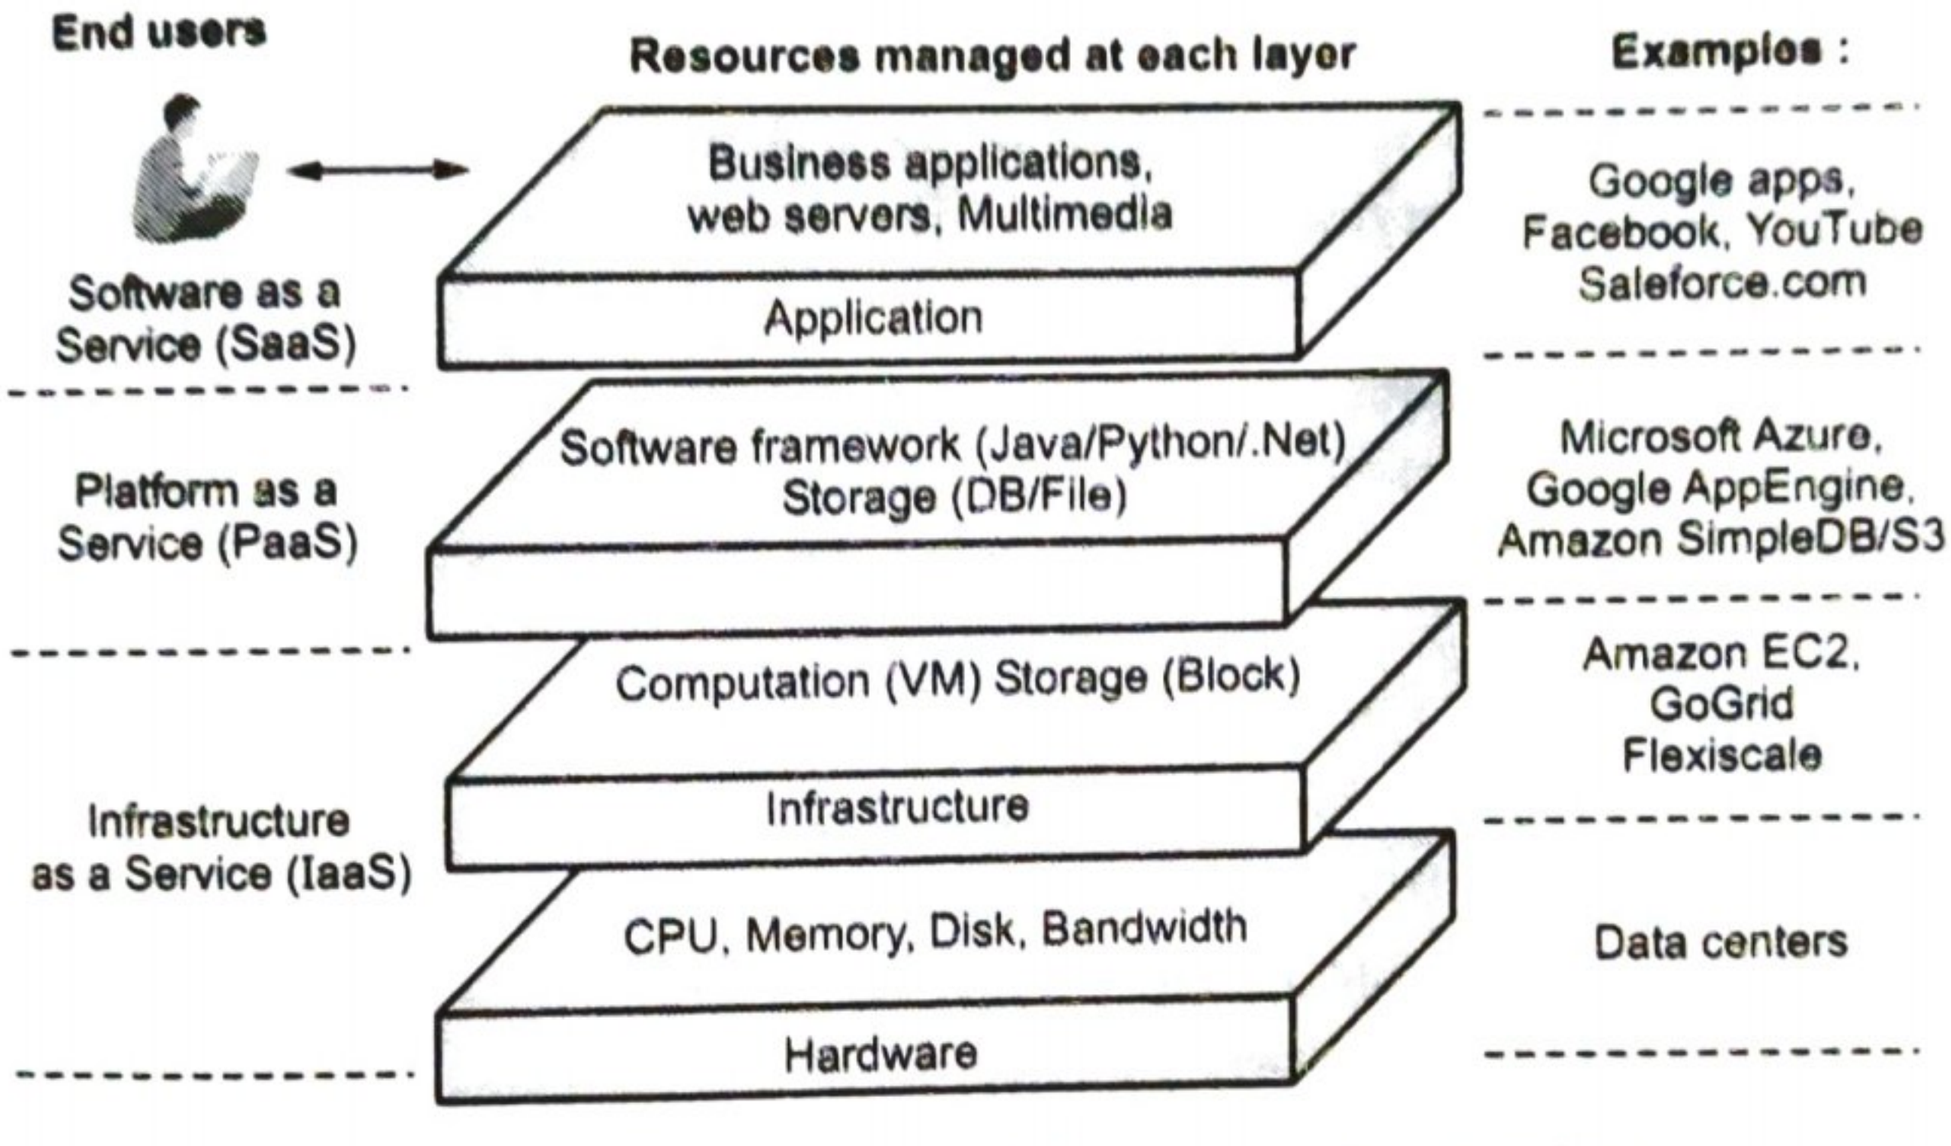

Fig. 7.1.2 Main categories of cloud computing services

- In a traditional on-premises data center, the IT team has to build and manage everything. Whether the team is building proprietary solutions from scratch or purchasing commercial software products, they have to install and manage one-to-many servers, develop and install the software, ensure that the proper levels of security are applied, apply patches routinely and much more.
- Each cloud service model provides levels of abstraction and automation for these tasks, thus providing more agility to the cloud service consumers so they can focus more time on their business problems and less time on managing infrastructure.
- Cloud based services have several common attributes :

a) **Virtualization** : Cloud computing utilizes server and storage virtualization extensively to allocate / reallocate resources rapidly.

b) **Multi-tenancy** : Resources are pooled and shared among multiple users to gain economies of scale.

c) **Network-access** : Resources are accessed via web-browser or thin client using a variety of networked devices (computer, tablet, smartphone).

d) **On demand** : Resources are self-provisioned from an online catalogue of pre-defined configurations.

e) **Elastic** : Resources can scale up or down automatically.

f) **Metering/chargeback** : Resource usage is tracked and billed based on service arrangement.

7.1.1 Implement Cloud Services

- Moving an application to the cloud is not a simple task, so its best to work with an experienced service provider that can help ensure the migration, implementation and ongoing support of your solution.
- **Steps for implementation :**

1. **Define your project :** Some applications and infrastructures should never be put on a cloud. Decide what you want to move to the cloud and whether or not it's feasible.
2. **Select the platform :** Choose a platform that is fast, easy and safe to deploy. Ensure you have a flexible platform that scales to support your evolving business model and future growth.
3. **Understand security policies :** Many service providers believe that data security is your responsibility, not theirs. Make sure you have a clear understanding of who is responsible and ensure that the right resources are in place.
4. **Select your cloud computing service provider :** Partner with a service provider that has success with businesses similar to yours and knows your technology.
5. **Determine service level agreements :** In addition to uptime, be very clear with your service provider when it comes to SLAs and exactly what they do and do not cover, such as data availability or data protection.
6. **Understand who owns recovery :** Outages will happen, so know in advance if you or your service provider is responsible for recovery.
7. **Migrate in phases :** Roll out a phased migration that allows you to gradually increase the load and gives you time to fine tune and minimize risks while maintaining business continuity.
8. **Think ahead :** Your business requirements can change at any time, so choose a cloud solution that allows you to move between on-premise and cloud as needed, and one that allows you to move to a different cloud service provider if necessary.

7.1.2 Emerging Markets and the Cloud

- Cloud computing represents a paradigm shift, a transition from computing-as-a-product to computing-as-a-service. Instead of buying hardware and software products, which require installation, configuration, and maintenance, cloud computing lets you use applications and computing infrastructures in the cloud as a service, so you pay only for resources used.
- Cloud computing can be applied to a range of areas, including e-commerce, education, healthcare, governance, telecommuting, community building, and emergency response. Not having legacy systems or applications requiring migration gives enterprises and individuals in emerging markets an optimal opportunity for cloud solutions.
- Furthermore, new developments are flourishing in emerging markets, making them attractive to both global and local cloud providers in search of new revenue opportunities.
- The accessibility of the cloud may become a chief factor in the ability of these markets to expand their global trade capabilities and enhance trade with other emerging markets. This will also impact small and medium sized businesses by driving job creation and increasing access to new products and business configurations.
- In the case of developing governments, the cloud can support efforts to enhance their ability to provide services in an economical and effective manner to citizens in areas such as healthcare, education, telecommunications, etc.
- Cloud services offer practically unlimited potential for emerging markets in nations like India and China that have already embarked on the process of implementing new technology.
- **Cloud in education** : The cloud has been a bit slower to catch on in education, with a small number of organizations saying the cloud currently has a pervasive presence in either type of economy. Educational organizations in developing economies, however, believe that in the next one to five years, the cloud will be a major factor in education.
- **Cloud in retail** : The sky's the limit for the cloud in the retail sector in developing economies. Three-quarters of retail organizations already say the cloud has a strong presence in developing economies. Most impressively, 80 percent of organizations in developing economies say the cloud will be a major factor for retail overall in their country.
- **Manufacturing** : Manufacturing organizations in both developed and developing economies say the cloud has a significant presence now. Cloud computing is being

used to reduce supply chain costs, connect suppliers globally, and to support partnerships between customers and suppliers. But the challenges for manufacturers to work in the cloud are not insignificant. Embedding cloud into a factory requires the design of new sensors, ensuring common standards across machines, communications protocols and a host of other cyber-physical challenges to be met.

- **Cloud in banking services :** One industry where the cloud will truly enable organizations to leapfrog in their technology adoption is in banking. The impact of the cloud will be felt on the banking industry overall, but particularly in enabling customers to make digital payments.

7.2 Tools for Building Private Cloud : IaaS using Eucalyptus

GTU : Summer-18, Winter-18

- Eucalyptus stands for "Elastic Utility Computing Architecture for Linking Your programs to Useful Systems". It is used to build private, public and hybrid clouds. It can also produce your own data center into a private cloud and allow you to extend the functionality to many other organizations.
- Eucalyptus in cloud computing is an open-source software platform for carrying out Infrastructure-as-a-Service in a hybrid cloud computing or private cloud computing environment.
- Eucalyptus is open-source software for building AWS-compatible private and hybrid clouds. As an Infrastructure as a Service (IaaS) product, Eucalyptus allows your users to provision your compute and storage resources on-demand.
- Eucalyptus has the following key features :

a) Support for multiple users with the help of a single cloud

b) Support for linux and windows virtual machines

c) Accounting reports

d) Use of WS-security to ensure secure communication between internal resources and processes

e) The option to configure policies and service level agreements based on users and the environment

f) Provisions for group, user management and security groups

- Challenges

a) **Extensibility :** Simple architecture and open internal APIs

b) **Client-side interface :** Amazon's EC2 interface and functionality (familiar and testable)

c) **Networking :** Virtual private network per cloud and must function as an overlay

d) **Security** : Must be compatible with local security policies

e) Packaging, installation, maintenance: system administration staff is an important constituency for uptake

- Fig. 7.2.1 shows Eucalyptus architecture

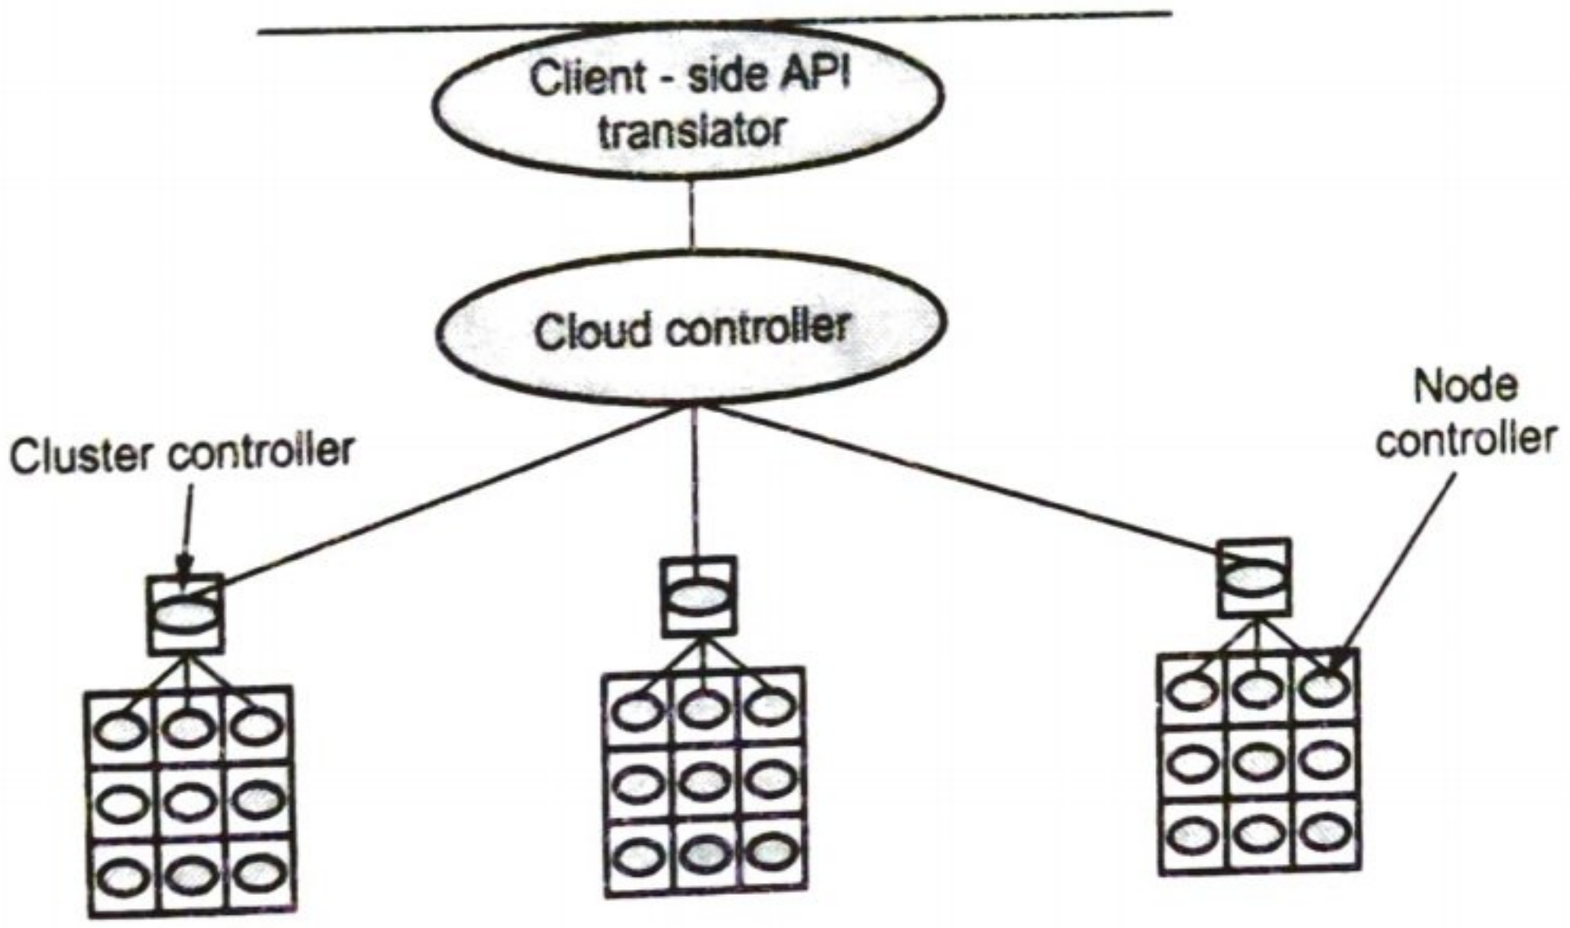

Fig. 7.2.1 Eucalyptus architecture

- **Components of eucalyptus in cloud computing :**

1. **Node controller** : The Node Controller (NC) is the component that executes on the physical resources that host VM instances and is responsible for instance start up, inspection, shutdown and clean-up.
2. **Cluster controller** : A collection of NCs that logically belong together report to a single Cluster Controller (CC) that typically executes on a cluster head node or server that has access to both private and public networks. The CC is responsible for gathering state information from its collection of NCs, scheduling incoming VM instance execution requests to individual NCs, and managing the configuration of public and private instance networks.
3. **Cloud controller** : Each Eucalyptus installation includes a single Cloud Controller (CLC) that is the user-visible entry point and global decision-making component of a Eucalyptus installation. The CLC is responsible for processing incoming user-initiated or administrative requests, making high-level VM instance scheduling decisions, processing service-level agreements (SLAs) and maintaining persistent system and user metadata.

- The CLC itself is composed of a collection of services that handle user requests and authentication, persistent system and user metadata, and the management and monitoring of VM instances.

4. **Client interface :** The CLC's client interface service essentially acts as a translator between the internal Eucalyptus system interfaces and some defined external client interface.
   - For example, Amazon provides a WSDL document that describes a Web-service SOAP based client interface to their service as well as a document describing an HTTP Query-based interface, both of which can be translated by the CLC user interface service into Eucalyptus internal objects.
5. **Administrative interface :** In addition to supporting primary tasks, such as starting and stopping instances, a cloud infrastructure must support administrative tasks, such as adding and removing users and disk images.
   - Eucalyptus supports such tasks though a Web based interface, implemented by the cloud controller, and command line tools. Unlike the client interface, however, the administrative interface is unique to Eucalyptus.
6. **Instance control :** Creation of virtual machine instance metadata in Eucalyptus is managed by a component of the CLC named the VmControl service.
7. **SLA implementation and management :** Service-level agreements (SLAs) are implemented as extensions to the message handling service which can inspect, modify, and reject the message, as well as the state stored by VmControl
   - Eucalyptus does not assume that all worker nodes will have publicly routable IP addresses. Each cloud allocation will have one or more public IP addresses. All cloud images have access to a private network interface. Two types of networks internal to a cloud allocation.

7.2.1 Eucalyptus Installation

- To install Eucalyptus, perform the following tasks :

1. **Plan your installation :** In order to get the most out of a Eucalyptus deployment.
2. **Configure dependencies :** Before you install Eucalyptus, ensure you have the appropriate dependencies installed and configured.
3. **Install repositories :** Downloads RPM packages.
4. Configure eucalyptus
5. Start eucalyptus
6. Register eucalyptus services
7. Configure the runtime environment

- Features of eucalyptus in cloud computing are :

a) Supports both Windows and Linux virtual machines.

b) API is viable with the Amazon EC2 platform.

c) Viable with Simple Storage Service (S3) and Amazon Web Services (AWS).

Installing the node controller :

- There are two main ways of going about installing Eucalyptus.
- The first way is to download the required RPMs onto your machine, install each of them and then manually configure the cloud as per your needs.
- The second way is much faster and will get your Eucalyptus cloud up and running in a matter of minutes.
- Installing the node controller is a very simple process. Once your machine boots from the Eucalyptus Faststart DVD, select the option 'Install CentOS 6 with Eucalyptus Node Controller' from the boot screen

Installing the cloud controller :

- Installation of the cloud controller is very similar to the nodes, with a few exceptions. Once your machine boots from the Eucalyptus Faststart DVD, select the option 'Install CentOS 6 with Eucalyptus Frontend' from the boot screen.
- Again, select the appropriate 'Language' and 'Keyboard settings' according to your needs.
- Provide a 'Static IP' and a suitable 'Host Name' to your cloud controller in the 'Network Configuration' wizard.
- Once done, you will be provided with an interface to supply a 'Public IP Range/ List' for your Eucalyptus cloud. You need to enter a valid IP address range here. These public IPs will be mapped to individual Eucalyptus instances (virtual machines) once they are launched in the cloud.

University Question

1. *What is eucalyptus ? Explain in brief.* **GTU : Summer-18, Winter-18, Marks 4**

7.3 PaaS on IaaS : AppScale

- AppScale is an open source distributed software system that implements a cloud platform as a service (PaaS). The goal of AppScale is to provide developers with a rapid, API-driven development platform that can run applications on any cloud infrastructure
- AppScale makes cloud applications easy to deploy and scale over disparate cloud fabrics, implementing a set of APIs and architecture that also makes apps portable across the services they employ.
- AppScale is API-compatible with Google App Engine (GAE) and thus executes GAE applications on-premise or over other cloud infrastructures, without modification.

- AppScale implements the App Engine programming model for Web-based application development by implementing each API that GAE defines and supporting all the GAE programming languages.
- AppScale is a platform that allows users to deploy and host their own Google App engine applications. It executes automatically over Amazon EC2 and Eucalyptus as well as Xen and KVM. It supports the Python, Java, and Go Google App Engine platforms.
- The AppScale platform virtualizes, abstracts, and multiplexes cloud and system services across multiple applications, enabling write-one, run-anywhere (WORA) program development for the cloud.
- AppScale decouples application logic from its service ecosystem to give developers and cloud administrators control over application deployment, data storage, resource use, backup, and migration
- Fig. 7.3.1 shows design of the AppScale cloud platform.

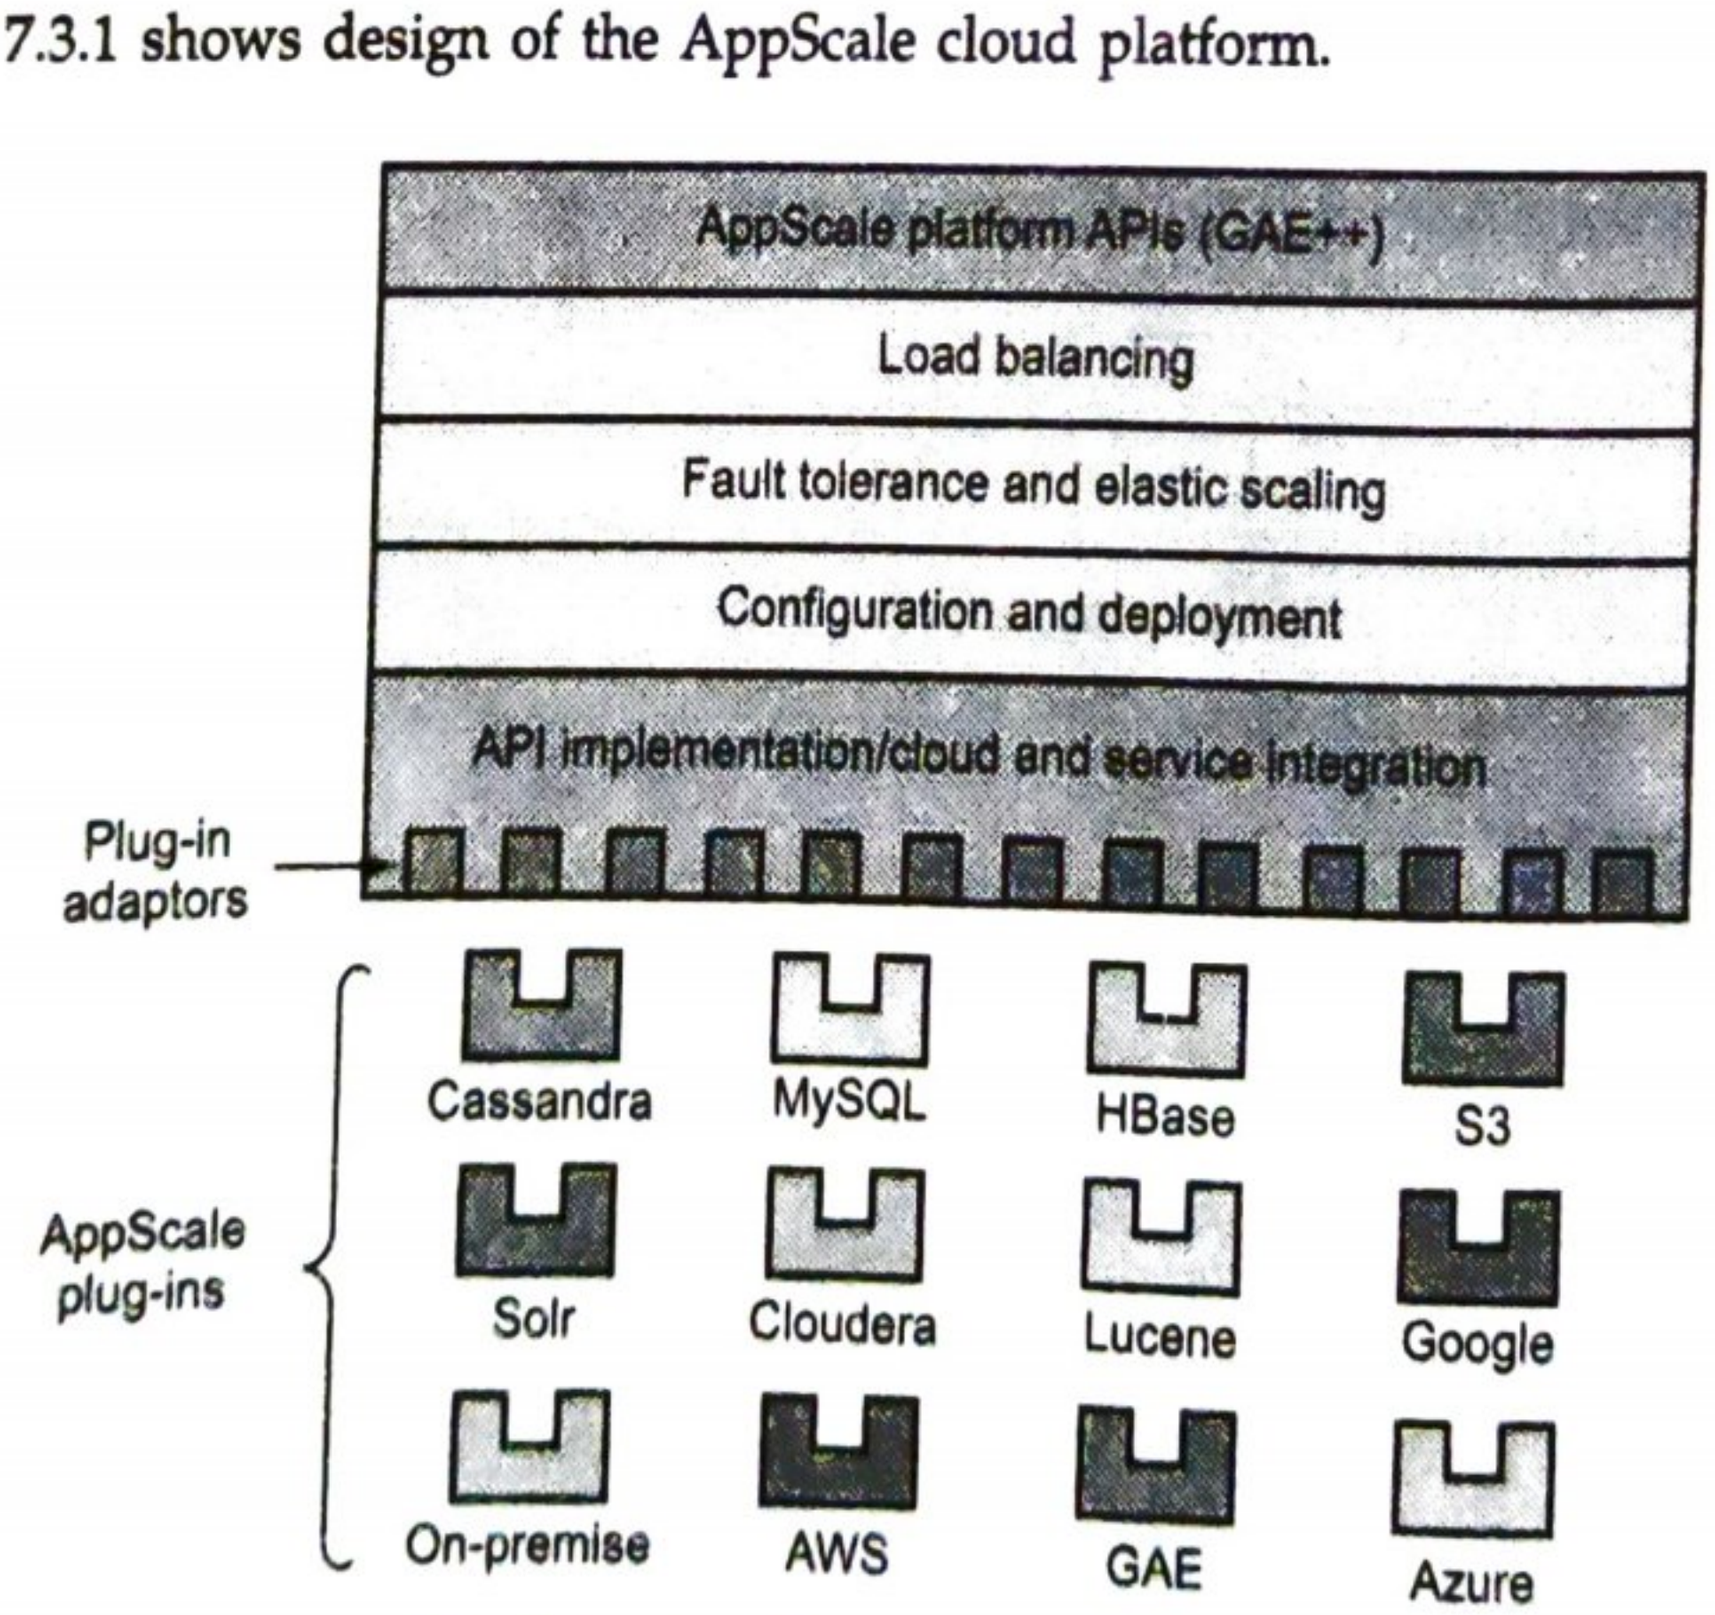

Fig. 7.3.1 AppScale cloud platform

- AppScale implements a multitier distributed web service stack with automatic deployment, load balancing and scaling, along with API adaptors for alternatives for each service API.
- The AppScale platform also provides the scalability, ease of use and high availability that users have come to expect from public cloud platforms and infrastructures. This includes elasticity and fault detection/ recovery,

authentication and user control, monitoring and logging, cross-cloud data and application migration, hybrid cloud multi-tasking, and offline analytics and disaster recovery

7.4 Multiple Choice Questions

Q.1 The Eucalyptus framework was one of the first open-source projects to focus on building _____ clouds.

a IaaS
b PaaS
c SaaS
d All of these

Q.2 _____ is used to build private, public and hybrid clouds.

a AppScale
b Eucalyptus
c Amazon
d Google

Q.3 AppScale is an open source distributed software system that implements a cloud _____.

a infrastructure as a service
b software as a service
c data as a service
d platform as a service

Q.4 AppScale is a platform that allows users to deploy and host their own _____ engine applications.

a Microsoft App
b Amazon App
c Google App
d Youtube App

Q.5 Eucalyptus provides a platform for _____ interface so that users can calculate the resources available in private clouds and the resources available externally in public cloud services.

a single
b multiple
c single and multiple
d two

Q.6 Which controller is not part of Eucalyptus ?

a Cluster controller
b Node controller
c Object controller
d Storage controller

Q.7 Eucalyptus is example of _____

a IaaS
b PaaS
c SaaS
d CaaS

Answer Keys for Multiple Choice Questions :

Q.1	a	Q.2	b
Q.3	d	Q.4	c
Q.5	a	Q.6	c
Q.7	a		

❑❑❑