PNS SCHOOL OF ENGINEERING & TECHNOLOGY

Nishamani Vihar, Marshaghai, Kendrapara

LECTURE NOTES

ON

CLOUD COMPUTING

DEPARTMENT OF COMPUTER SCIENCE & ENGINEERING

6TH SEMESTER

PREPARED BY

Er. JYOTSNAMAYEE BISWAL

LECTURER IN COMPUTER SCIENCE & ENGINEERING

Th-3 CLOUD COMPUTING

Theory	**4 Periods per week**	**Internal Assessment**	**20 Marks**
Total Periods	**60 Periods**	**End Sem Exam**	**80 Marks**
Examination	**3hours**	**Total Marks**	**100Marks**

A. Topic wise distribution of periods

Sl. No.	Topics	Periods
1	INTRODUCTION TO CLOUD COMPUTING	05
2	CLOUD COMPUTING ARCHITECTURE	08
3	SCALABILITY AND FAULT TOLERANCE	08
4	CLOUD MANAGEMENT AND VIRTUALISATION TECHNOLOGY	08
5	VIRTUALISATION	08
6	CLOUD SECURITY	08
7	CLOUD COMPUTING SECURITY ARCHITECTURE	05
8	MARKET BASED MANAGEMT OF CLOUDS	05
9	HADOOP	05
	TOTAL	**60**

B.RATIONALE:

Cloud computing is one of the emerging topics in Information Technology. It is the biggest buzz in the computer world. Cloud computing means you can deliver applications to your end user faster than ever, without investing in new infrastructure, training new personnel or licensing new software. It is a practical approach to experience direct cost benefits and easy to use for the users.

C. Objective : After completion of this course the student will be able to**:**

- Understand the basic concepts of cloud and cloud architecture.
- Learn about different cloud computing technology.
- Learn about the service levels for cloud applications.
- Provides a practical exposure to professionals intending to work in cloud computing environment.
- Understand the map reduce model and its application.
- Learn about basic concepts of software productivity in a cloud.
- Understand web services and platforms.

D. DETAIL CONTAINS:

1. Introduction To Cloud Computing

1.1. Historical development
1.2. Vision of Cloud Computing
1.3. Characteristics of Cloud computing
1.4. Cloud computing Reference model
1.5. Cloud computing environment
1.6. Cloud Service requirements
1.7. Cloud and Dynamic Infrastructure
1.8. Cloud Adoption
1.9. Cloud applications

2. Cloud Computing Architecture

2.1. Introduction
2.2. Cloud Reference Model
2.3. Types of Clouds
2.4. Cloud Interoperability and standards
2.5. Cloud computing Interoperability use cases
2.6. Role of standards in Cloud Computing environment

3. Scalability and Fault Tolerance

3.1. Introduction
3.2. Scalability and Fault Tolerance
3.3. Cloud solutions
3.4. Cloud Ecosystem
3.5. Cloud Business process management
3.6. Portability and Interoperability
3.7. Cloud Service management
3.8. Cloud Offerings
3.9. Testing under Control
3.10. Cloud service Controls
3.11. Virtual desktop Infrastructure

4. Cloud Management and Virtualisation Technology

4.1. Create a virtualised Architecture
4.2. Data Centre
4.3. Resilience
4.4. Agility
4.5. Cisco Data Centre Network architecture
4.6. Storage
4.7. Provisioning
4.8. Asset Management
4.9. Concept of Map Reduce
4.10. Cloud Goverance
4.11. Load Balancing
4.12. High Availability
4.13. Disaster Recovery

5. Virtualisation

5.1. Virtualisation
5.2. Betwork Virtualisation
5.3. Desktop and Application Virtualisation
5.4. Desktop as a service
5.5. Local desktop Virtualisation
5.6. Virtualisation benefits

5.7. Server Virtualisation
5.8. Block and File level Storage Virtualisation
5.9. Virtual Machine Monitor
5.10. Infrastructure Requirements
5.11. VLAN and VSAN

6. Cloud Security

6.1. Cloud Security Fundamentals
6.2. Cloud security services
6.3. Design Principles
6.4. Secure Cloud software requirements
6.5. Policy Implementation
6.6. Cloud Computing Security Challenges

7. Cloud Computing Security Architecture

7.1. Architectural Considerations
7.2. Information Classification
7.3. Virtual Private Networks
7.4. Public Key and Encryption Key management
7.5. Digital certificates
7.6. Key management
7.7. Memory Cards
7.8. Implementing Identity Management
7.9. Controls and Autonomic System

8. Market Based Management of Clouds

8.1. Cloud Information security vendors
8.2. Cloud Federation, charactrization
8.3. Cloud Federation stack
8.4. Third Party Cloud service
8.5. Case study

9. Hadoop

9.1. Introduction
9.2. Data Source
9.3. Data storage and Analysis
9.4. Comparison with other system

Coverage of Syllabus upto Internal Exams (I.A.) Chapter 1,2,3,4

BOOKS Recommended:-

Sl .No.	Name of the Author	Title of the Book	Name of the Publisher
1	Pankaj Sharma	Cloud Computing	Katson Books
1	Dr. U.S. Pandey , Dr. KavitaChoudhary	Cloud Computing	S. Chand
2	PrasantkumarPattnaik, ManasRanjanKabat , Souvik Pal	Fundamentals of Cloud Computing	Vikas

UNIT-1
Introduction To Cloud Computing

HISTORICAL DEVELOPMENT

- The history of cloud computing starts from the 1950's and the work done by AT & T in the area of telephone networking .
- At that time AT & T had already begun to develop an architecture and system where data would be located centrally.
- The IT services progressed over the decades with the adoption of technologies such as Internet Service Providers (ISP) Application service Providers.
- One of the main principles of cloud computing from SAAS (Software as a service) to provide storage on demand, is that the computing capacity varies immediately and transparently with the customer's need.

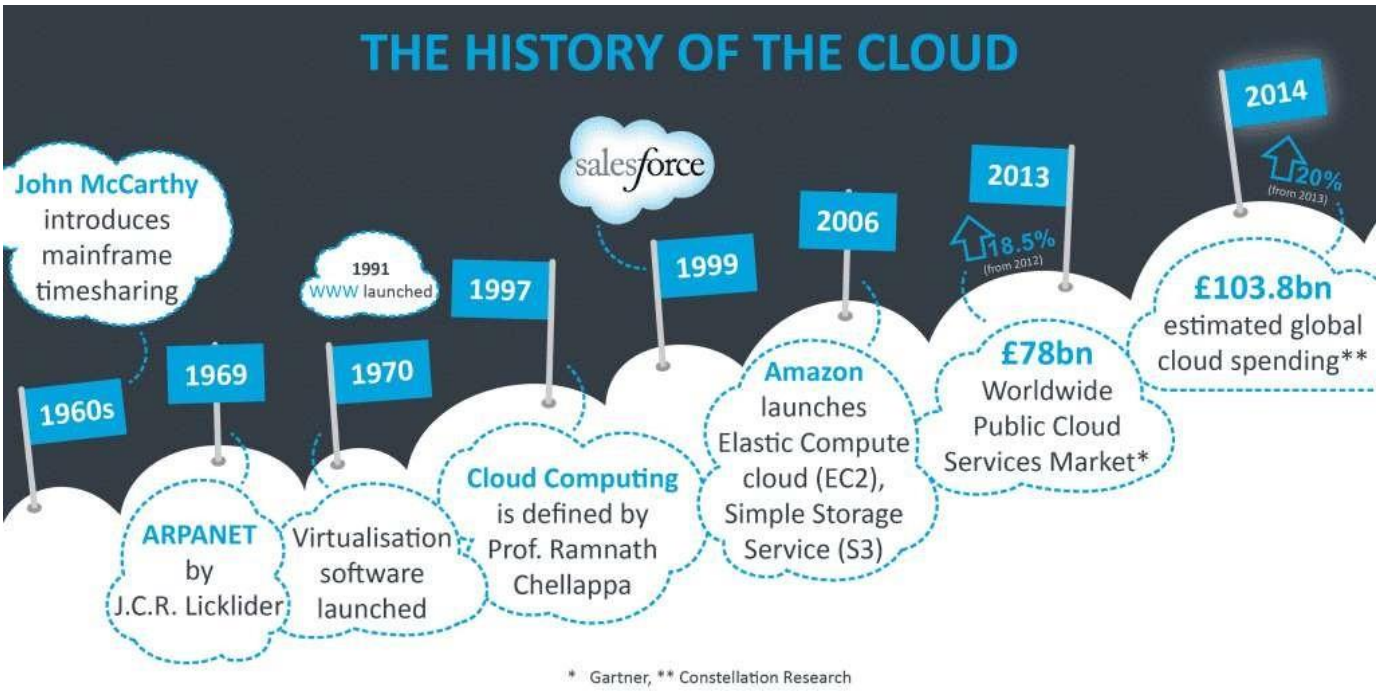

Evolution of cloud technologies

Following types are today's cloud implementations:-

Distributed Systems

- A distributed system is a collection of independent computers that appears to its users as a single system and also it acts as a single computer.
- The main and primary motive of distributed systems is to share resources and to utilize them better.
- This is absolutely true in case of cloud computing because in cloud computing we are sharing the single resource by paying rent.
- The resource is single because the definition of cloud computing clearly states that in cloud computing the single central copy of a particular software is stored in a sever .

Mainframes and thin client computing

It is highly reliable, powerful, centrally located form of computing service. A user of a mainframe system may access applications using a thin client**.**

Each mainframe system is designed to run at a high level of utilization without failure, and to support hardware up gradation.

The mainframes can host multiple virtual instances of operating system and this is a crucial requirement for supporting scalability within cloud computing

Utility Computing

Computing services that can be metered and billed to customers in the same way that electricity or telephony system operate, are known as utility computing services.

The concept of utility computing is also associated with the commercialization of problem solving in supercomputing systems.

Grid and Super Computing

The use of specialist supercomputers, or large number of computers configured to run inparallel in a 'grid' to solve the complex problems such as predicting the weather or decrypting data encrypted with strong encrypting algorithms is known as Grid and Super Computing.

Scalability and on demand processing power

The use of a supercomputer or grid computing service provides a level of scalability to those needing resources that may be too cost prohibitive to purchase in house.

The processing power within these systems can be shared and provided to multiple users concurrently to execute complex software programs.

Web 2.0

The global presence of the internet and the introduction of wireless networking and mobile devices featuring always on internet connectivity has raised expectations of users and demand for services over the internet.

Web 2.0 is the term given to describe a second generation of the World Wide Web that isfocused on the ability for people to collaborate and share information online.

Web 2.0 basically refers to the transition from static HTML Web pages to a more dynamic Web that is more organized and is based on serving Web applications to users.

Other improved functionality of Web 2.0 includes open communication with an emphasis on Web-based communities of users, and more open sharing of information.

Over time Web 2.0 has been used more as a marketing term than a computer-science-based term. wikis, and Web services are all seen as components of Web 2.0.

VISION OF CLOUD COMPUTING

- Cloud computing provides the facility to provision virtual hardware, runtime environment and services to a person having money.
- These all things can be used as long as they are needed by the user, there is no requirement for the upfront commitment.
- The whole collection of computing system is transformed into a collection of utilities, which can be provisioned and composed together to deploy systems in hours rather than days, with no maintenance costs.
- The long term vision of a cloud computing is that IT services are traded as utilities in an open market without technological and legal barriers.
- In the near future we can imagine that it will be possible to find the solution that matches with our requirements by simply entering our request in a global digital market that trades with cloud computing services.
- The existence of such market will enable the automation of the discovery process and its integration into its existing software systems.
- Due to the existence of a global platform for trading cloud services will also help serviceproviders to potentially increase their revenue.
- A cloud provider can also become a consumer of a competitor service in order to fulfill its promises to customers.

CHARACTERISTICS OF CLOUD COMPUTING

National Institute of Standards and Technology (NIST) is an agency under the scope of US Department of Commerce. NIST is responsible for defining standards in Science and Technology.

The Computer Security Division of NIST has provided a formal Definition and Characteristics of Cloud computing.

NIST five essential characteristics of Cloud Computing

1. On demand self-service
2. Broad network access
3. Resource pooling
4. Rapid Elasticity
5. Measured service

ISO 17788 six essential characteristics of Cloud Computing

1. On demand self-service
2. Broad network access
3. Resource pooling
4. Rapid Elasticity
5. Measured service
6. Multi-tenancy

1. On Demand Self service

Computer services such as Email, Application Network, or Server service can be provided without requiring interaction with each service provider.

Self-service means that the consumer performs all the actions needed to acquire the service himself, instead of going through an IT department. For example – The consumer's request is

then automatically processed by the cloud infrastructure, without human intervention on the provider's side.

2. Broad Network Access

Cloud capabilities are available over the network and accessed through standard mechanism that promote use by heterogeneous client such as mobile phone, laptop

3. Resource pooling

– The providers computing resources are pooled together to serve multiple customers, with different physical and virtual resources dynamically assigned and reassigned according to the customers demand.
– There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g. country, state, or datacenter).
– Example of resources includes storage, processing, memory, and network bandwidth.

4. Rapid elasticity

– Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand.
– To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time.

5. Measured service

– Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g. storage, processing, bandwidth, and active use account).
– Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

6. Multi-tenancy

In a private cloud, the customers are also called tenants, can have different business divisions inside the same company. In a public cloud, the customers are often entirely different organizations.

Most public cloud providers use the multi-tenancy model. Multi-tenancy allows customers to run one server instance, which is less expensive and makes it easier to deploy updates to a large number of customers.

Components of Cloud Computing
Components in a cloud refers to the platforms, like front end, back end and cloud based delivery and the network that used

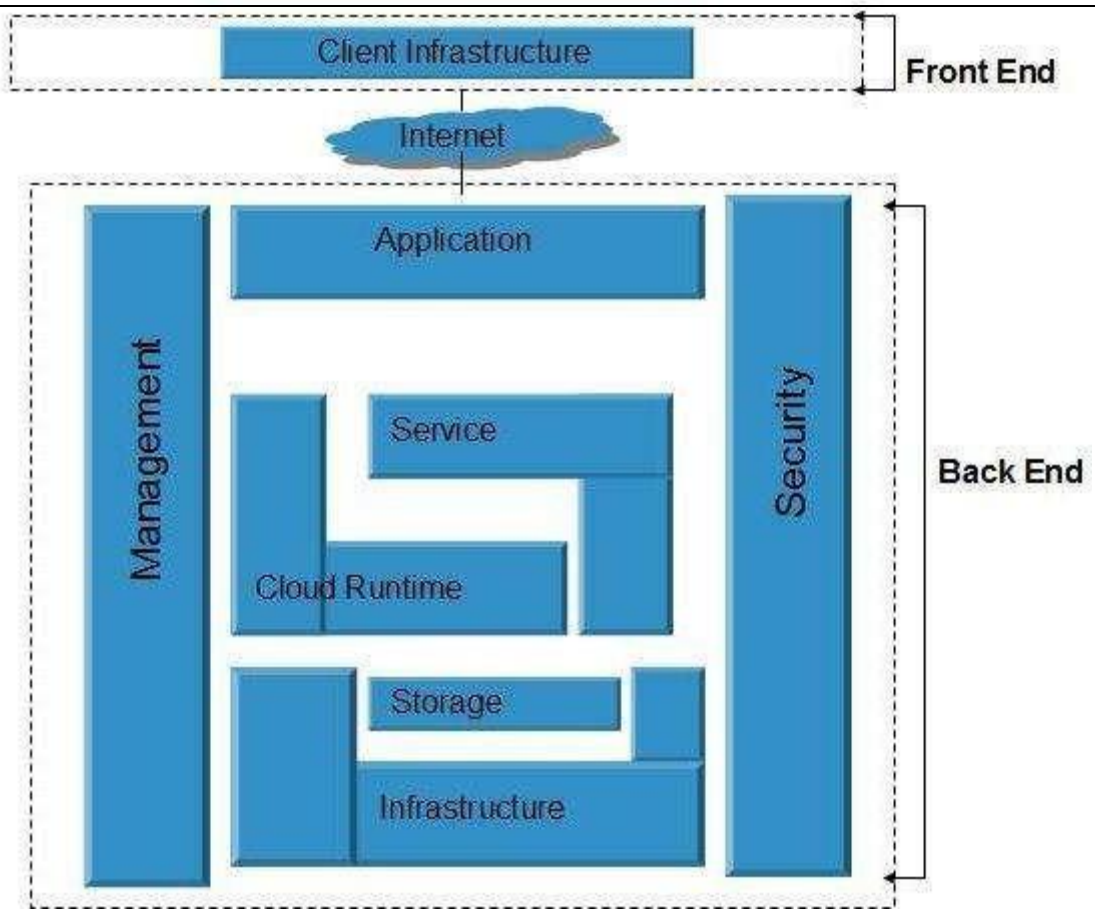

i. Front End

The front end is the client part of Cloud Computing which uses as per the requirement of the user. Front-end comprises of the applications and the interfaces which help to access the cloud computing. Example- Browser or an app created by the company itself.

ii. Back End

The back end is a part which manages by the allotted authorities of the company and their back end has large data storage facilities, Virtual machines, security system, and servers. They are also engaged in traffic management along with security management.

The basic components of cloud computing in a simple topology are divided into 3 (three) parts, namely clients, datacenter, and distributed servers. The three basic components have specific goals and roles in running cloud computing operations. The concept of the three components can be described as follows:

Clients on cloud computing architecture are said to be the exact same things that are plain, old, everyday local area networks (LANs). They are, typically, the computers that just sit on your desk. But they might also be laptops, tablet computers, mobile phones, or PDAs - all big drivers for cloud computing because of their mobility. Clients are interacting with to manage their information on the cloud.

Datacenter is collection of servers where the application to which you subscribe is housed. It could be a large room in the basement of your building full of servers on the other side of the world that you access via the Internet. A growing trend in the IT world is virtualizing servers. That is, software can be installed allowing multiple instances of virtual servers to be used. In this way, you can have half a dozen virtual servers running on one physical server.

Distributed Servers is a server placement in a different location. But the servers don't have to be housed in the same location. Often, servers are in geographically disparate locations. But to you, the cloud subscribers, these servers act as if they're humming away right next to each other.

Another component of cloud computing is Cloud Applications cloud computing in terms of software architecture. So that the user does not need to install and run applications using a computer. Cloud Platform is a service in the form of a computing platform that contains hardware infrastructure and software. Usually have certain business applications and use services PaaS as its business application infrastructure. Cloud Storage involves processes delivering data storage as a service. Cloud Infrastructure is the delivery of computing infrastructure as a service.

Cloud Computing services have several components required, namely:
a. Cloud Clients, a computer or software specifically designed for the use of cloud computing based services.
Example :
Mobile - Windows Mobile, Symbian
Thin Client - Windows Terminal Service, CherryPal
Thick Client - Internet Explorer, FireFox, Chrome
b. Cloud Services, products, services and solutions that are used and delivered real-time via internet media.
Example :
Identity - OpenID, OAuth, etc.
Integration - Amazon Simple Queue Service.
Payments - PayPal, Google Checkout.
Mapping - Google Maps, Yahoo! Maps.
c. Cloud Applications, applications that use Cloud Computing in software architecture so that users don't need to install but they can use the application using a computer.
Example :
Peer-to-peer - BitTorrent, SETI, and others.
Web Application - Facebook.
SaaS - Google Apps, SalesForce.com, and others
d. Cloud Platform, a service in the form of a computing platform consisting of hardware and infrastructure software. This service is a service in the form of a computing platform which contains infrastructure hardware and software. Usually has an application certain businesses and use PaaS services as application infrastructure his business
Example :
Web Application Frameworks - Python Django, Rubyon Rails, .NET
Web Hosting
Propietary - Force.com
e. Cloud Storage, involves the process of storing data as a service.
Example :
Database - Google Big Table, Amazon SimpleDB.
Network Attached Storage - Nirvanix CloudNAS, MobileMe iDisk.
f. Cloud Infrastructure, delivery of computing infrastructure as a service.
Example:
Grid Computing - Sun Grid.
Full Virtualization - GoGrid, Skytap.

Compute - Amazon Elastic Compute Cloud

THE CLOUD COMPUTING REFERENCE MODEL

The cloud computing reference model is an abstract model that characterizes and standardizes the functions of a cloud computing environment by partitioning it into abstraction layers and cross- layer functions. This reference model groups the cloud computing functions and activities into five logical layers and three cross-layer functions.

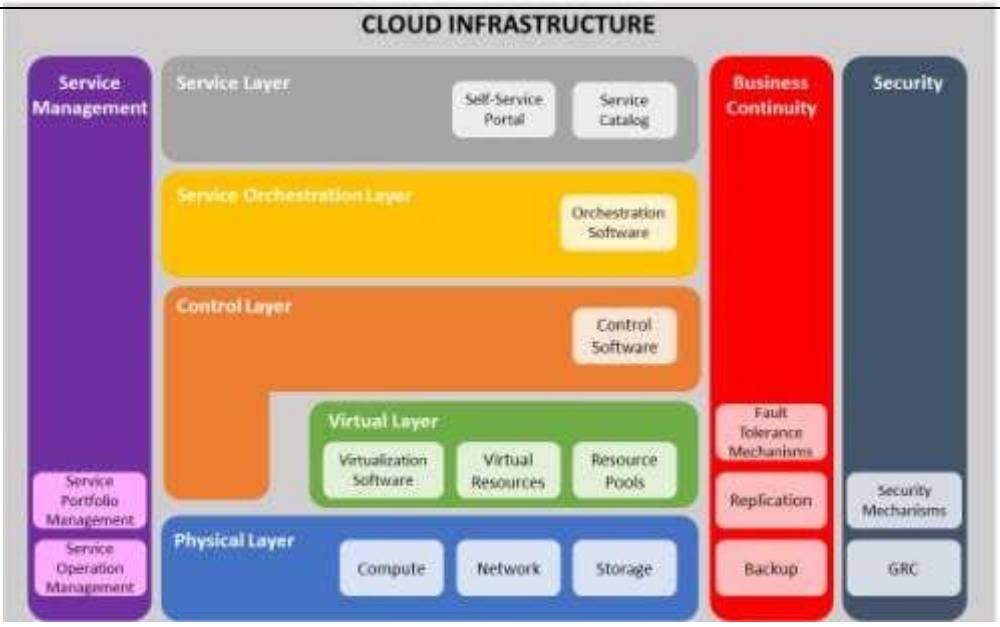

The five layers are physical layer, virtual layer, control layer, service orchestration layer, and service layer. Each of these layers specifies various types of entities that may exist in a cloud computing environment, such as compute systems, network devices, storage devices, virtualization software, security mechanisms, control software, orchestration software, management software, and so on. It also describes the relationships among these entities.

The three cross-layer functions are business continuity, security, and service management. Business continuity and security functions specify various activities, tasks, and processes that are required to offer reliable and secure cloud services to the consumers. Service management functionspecifies various activities, tasks, and processes that enable the administrations of the cloud infrastructure and services to meet the provider's business requirements and consumer's expectations.

Cloud computing layers

Physical Layer

- Foundation layer of the cloud infrastructure.
- Specifies entities that operate at this layer : Compute systems, network devices and storage devices. Operating environment, protocol, tools and processes.
- Functions of physical layer : Executes requests generated by the virtualization and control layer.

Virtual Layer

- Deployed on the physical layer.
- Specifies entities that operate at this layer : Virtualization software, resource pools, virtual resources.
- Functions of virtual layer : Abstracts physical resources and makes them appear as virtual resources (enables multitenant environment). Executes the requests generated by control layer.

Control Layer

- Deployed either on virtual layer or on physical layer
- Specifies entities that operate at this layer : control software
- Functions of control layer : Enables resource configuration, resource pool configuration and resource provisioning. Executes requests generated by service layer. Exposes resources to and supports the service layer. Collaborates with the virtualization software and enables resource pooling and creating virtual resources, dynamic allocation and optimizing utilization of resources.

Service Orchestration Layer

- Specifies the entites that operate at this layer : Orchestration software.
- Functions of orchestration layer : Provides workflows for executing automated tasks. Interacts with various entities to invoke provisionning tasks.

Service Layer

- Consumers interact and consume cloud resources via thos layer.
- Specifies the entities that operate at this layer : Service catalog and self-service portal.
- Functions of service layer : Store information about cloud services in service catalog and presents them to the consumers. Enables consumers to access and manage cloud services via a self-service portal.

Cross-layer function

Business continuity

- Specifies adoption of proactive and reactive measures to mitigate the impact of downtime.
- Enables ensuring the availability of services in line with SLA.
- Supports all the layers to provide uninterrupted services.

Security

- Specifies the adoption of : Administrative mechanisms (security and personnel policies, standard procedures to direct safe execution of operations) and technical mechanisms (firewall, intrusion detection and prevention systems, antivirus).
- Deploys security mechanisms to meet GRC requirements.
- Supports all the layers to provide secure services.

Service Management

Specifies adoption of activities related to service portfolio management and service operation management.

Service portfolio management :

• Define the service roadmap, service features, and service levels

• Assess and prioritize where investments across the service portfolio are most needed

• Establish budgeting and pricing

• Deal with consumers in supporting activities such as taking orders, processing bills, and collecting payments

Service operation management :

- Enables infrastructure configuration and resource provisioning
- Enable problem resolution
- Enables capacity and availability management
- Enables compliance conformance
- Enables monitoring cloud services and their constituent elements

SERVICE MODEL:

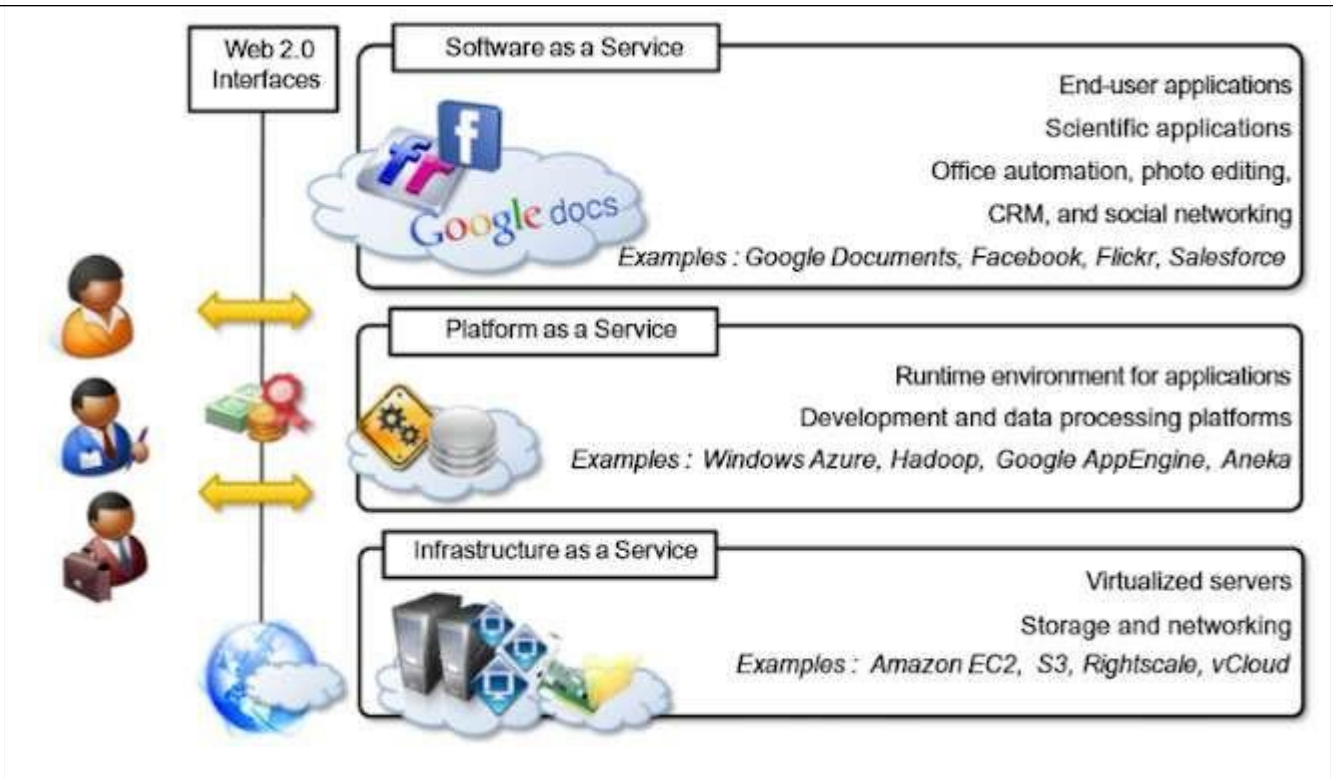

If we look in to the reference model as seen in above image we will find classification of Cloud Computing services:

1. Infrastructure-as-a-Service (IaaS),
2. Platform-as-a-Service (PaaS), and
3. Software-as-a-Service (SaaS).
4. Web 2.0

1. Infrastructure as a service (IaaS) is a cloud computing offering in which a vendor provides users access to computing resources such as servers, storage and networking.

2. Platform as a service (PaaS) is a cloud computing offering that provides users with a cloud environment in which they can develop manage and deliver applications.

3. Software as a service (SaaS) is a cloud computing offering that provides users with access to a vendor's cloud-based software. Users do not install applications on their local devices. Instead, the applications reside on a remote cloud network accessed through the web or an API. Through the application, users can store and analyze data and collaborate on project

4.Web 2.0 is the term used to describe a variety of web sites and applications that allow anyone to create and share online information or material they have created. A key element of the technology is that it allows people to create, share, collaborate & communicate.

THE FIVE MOST POPULAR CLOUD DEPLOYMENT MODELS

Deploying to the cloud provides organizations with flexible and scalable virtual computing resources.

A cloud deployment model is the type of architecture a cloud system is implemented on. These models differ in terms of management, ownership, access control, and security protocols.

The five most popular cloud deployment models are **public**, **private**, **virtual private (VPC)**, **hybrid**, and **community cloud**.
The following sections explain cloud deployment models in further detail.

Public Cloud

The public cloud model is the most widely used cloud service. This cloud type is a popular option for web applications, file sharing, and non-sensitive data storage.

The service provider owns and operates all the hardware needed to run a public cloud. Providers keep devices in massive data centers.

The public cloud delivery model plays a vital role in development and testing. Developers often use public cloud infrastructure for development and testing purposes. Its virtual environment is cheap and can be configured easily and deployed quickly, making it perfect for test environments.

Advantages of Public Cloud

Benefits of the public cloud include:

- **Low cost:** Public cloud is the cheapest model on the market. Besides the small initial fee, clients only pay for the services they are using, so there is no unnecessary overhead.
- **No hardware investment:** Service providers fund the entire infrastructure.
- **No infrastructure management:** A client does not need a dedicated in-house team to make full use of a public cloud.

Disadvantages of Public Cloud

The public cloud does have some drawbacks:

- **Security and privacy concerns:** As anyone can ask for access, this model does not offer ideal protection against attacks. The size of public clouds also leads to vulnerabilities.
- **Reliability:** Public clouds are prone to outages and malfunctions.
- **Poor customization:** Public offerings have little to no customization. Clients can pick the operating system and the sizing of the VM (storage and processors), but they cannot customize ordering, reporting, or networking.
- **Limited resources:** Public clouds have incredible computing power, but you share the resources with other tenants. There is always a cap on how much resource you can use, leading to scalability issues.

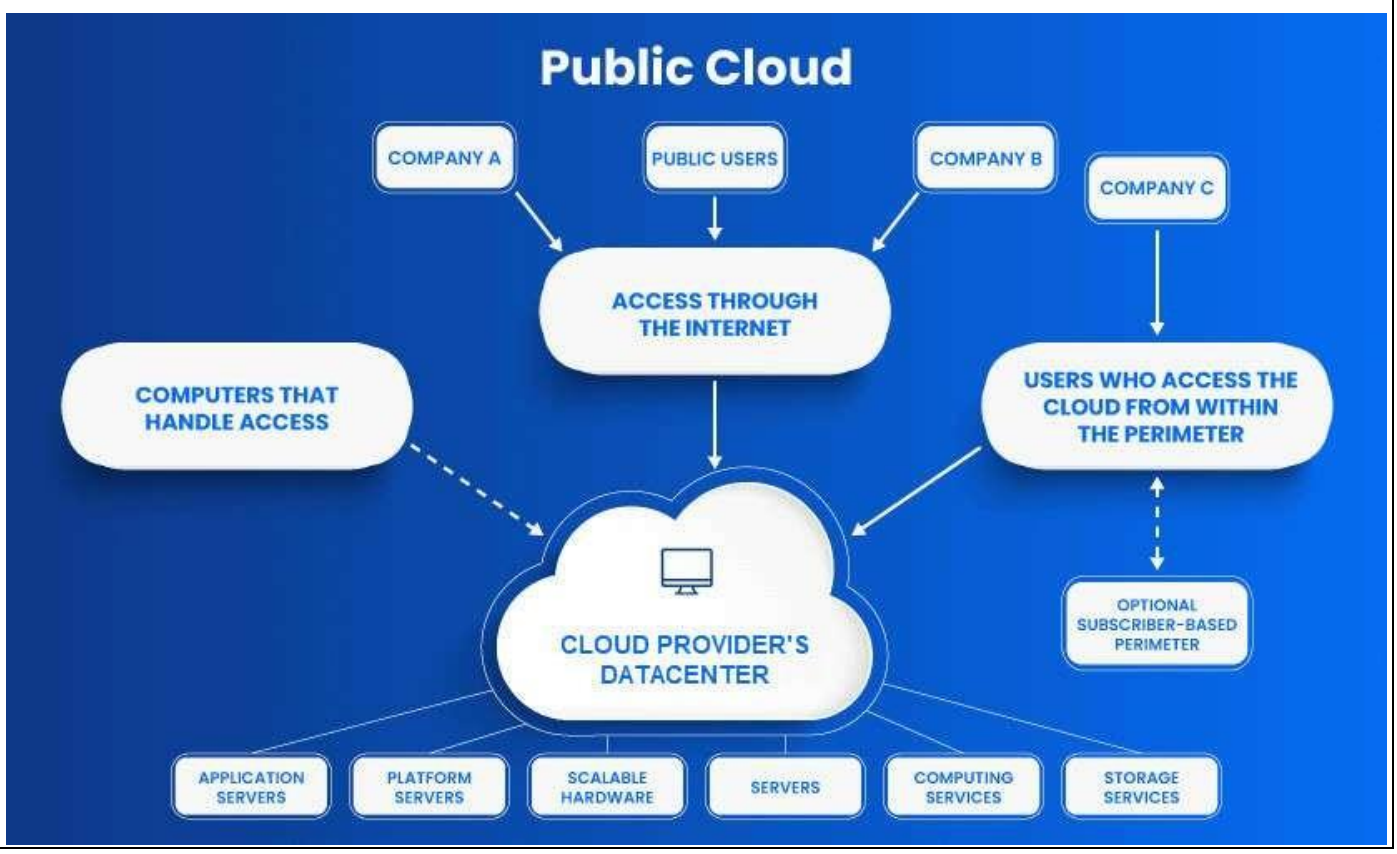

Private Cloud

Whereas a public model is available to anyone, a private cloud belongs to a specific organization. That organization controls the system and manages it in a centralized fashion. While a third party (e.g., service provider) can host a private cloud server (a type of collocation), most companies choose to keep the hardware in their on-premises data center. From there, an in-house team can oversee and manage everything.

The private cloud deployment model is also known as the internal or corporate model.

Advantages of Private Cloud

Here are the main reasons why organizations are using a private cloud:

- **Customization:** Companies get to customize their solution per their requirements.
- **Data privacy:** Only authorized internal personnel can access data. Ideal for storing corporate data.
- **Security:** A company can separate sets of resources on the same infrastructure. Segmentation leads to high levels of security and access control.
- **Full control:** The owner controls the service integrations, IT operations, rules, and user practices. The organization is the exclusive owner.
- **Legacy systems:** This model supports legacy applications that cannot function on a public cloud.

Disadvantages of Private Cloud

- **High cost:** The main disadvantage of private cloud is its high cost. You need to invest in hardware and software, plus set aside resources for in-house staff and training.
- **Fixed scalability:** Scalability depends on your choice of the underlying hardware.
- **High maintenance:** Since a private cloud is managed in-house, it requires high maintenance.

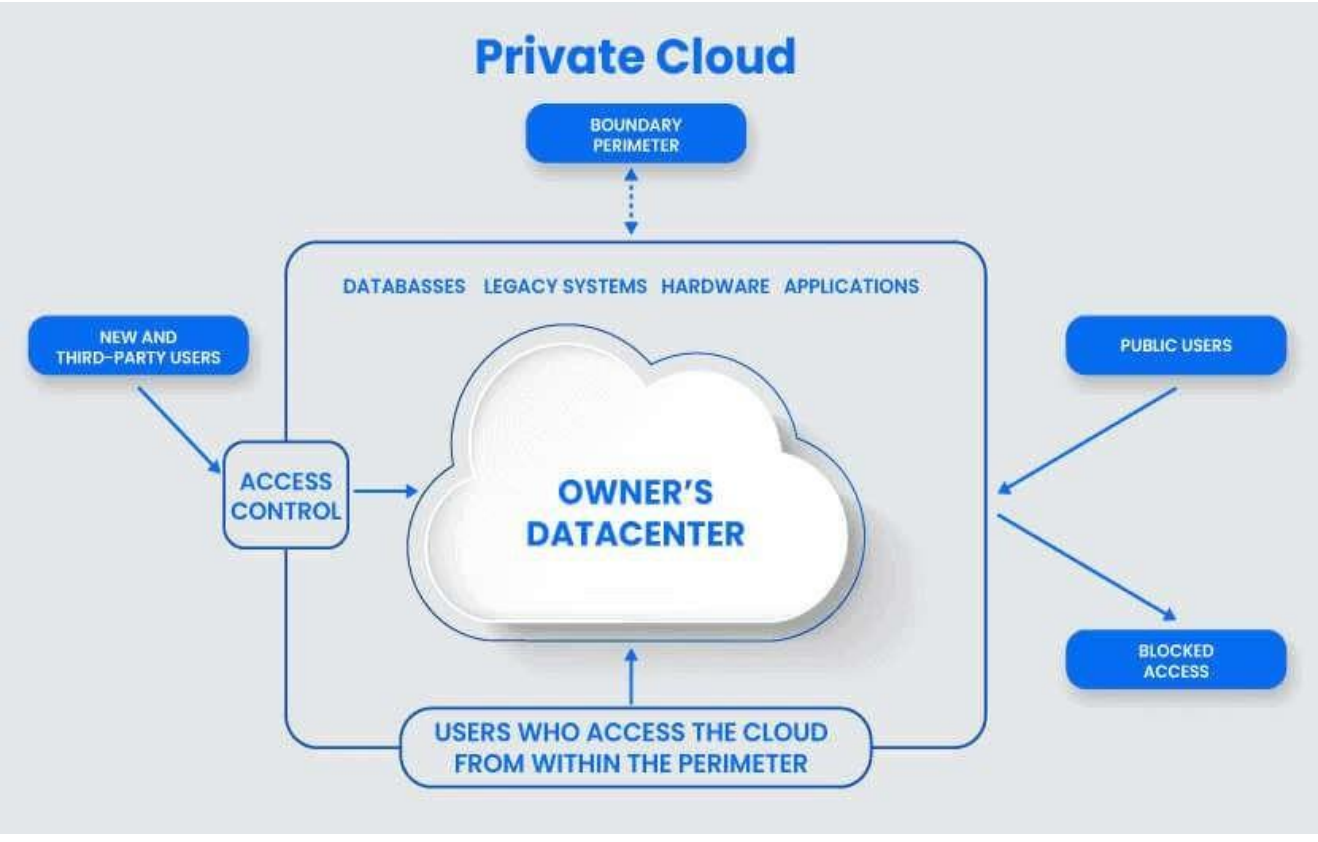

Virtual Private Cloud (VPC)

A VPC customer has exclusive access to a segment of a public cloud. This deployment is a compromise between a private and a public model in terms of price and features.

Access to a virtual private platform is typically given through a secure connection (e.g., VPN). Access can also be restricted by the user's physical location by employing firewalls and IP address whitelisting.

See phoenixNAP's Virtual Private Data Center offering to learn more about this cloud deployment model.

Advantages of Virtual Private Cloud

Here are the positives of VPCs:

- **Cheaper than private clouds:** A VPC does not cost nearly as much as a full-blown private solution.
- **More well-rounded than a public cloud:** A VPC has better flexibility, scalability, and security than what a public cloud provider can offer.
- **Maintenance and performance:** Less maintenance than in the private cloud, more security and performance than in the public cloud.

Disadvantages of Virtual Private Cloud

The main weaknesses of VPCs are:

- **It is not a private cloud:** While there is some versatility, a VPC is still very restrictive when it comes to customization.
- **Typical public cloud problems:** Outages and failures are commonplace in a VPC setup.

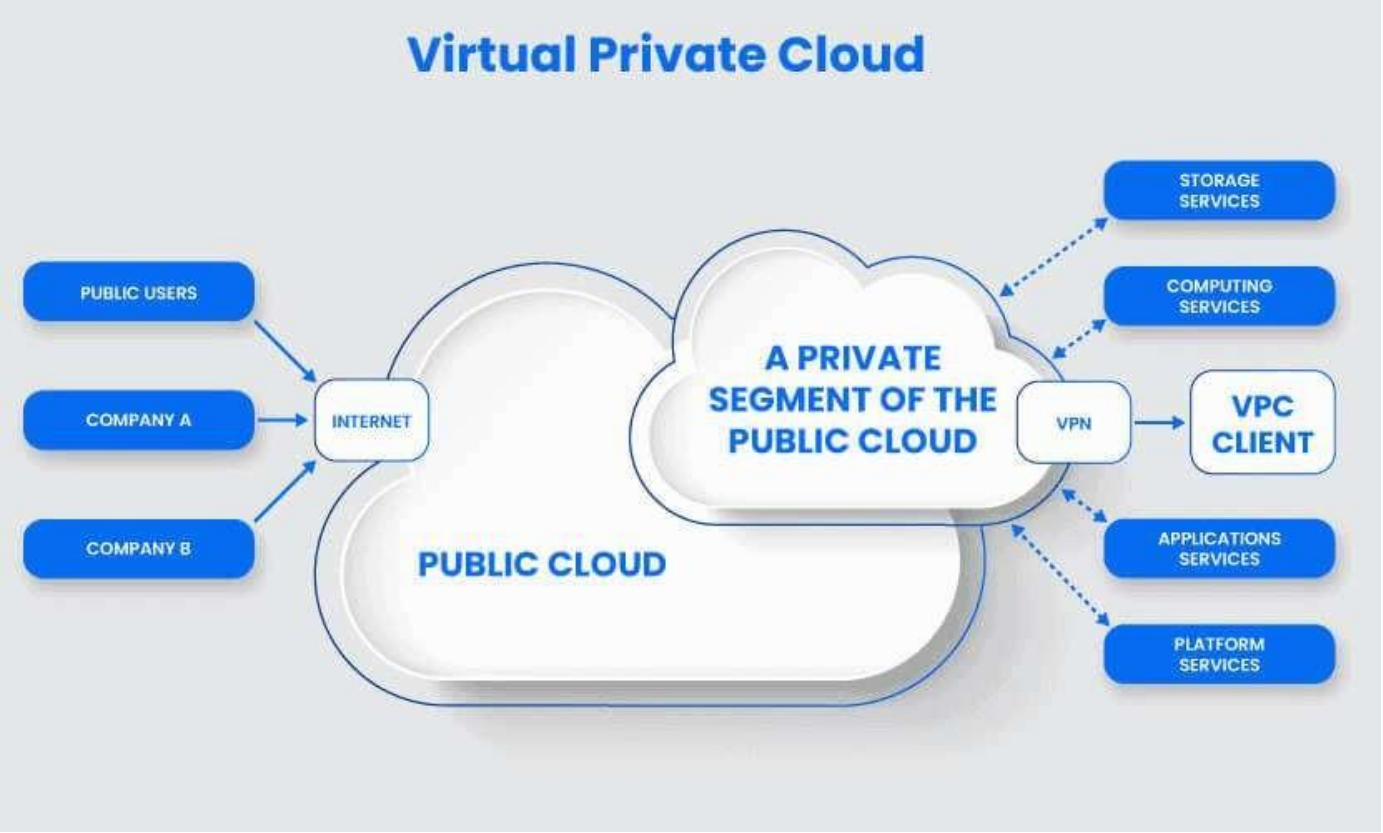

Community Cloud

The community cloud deployment model operates as a public cloud. The difference is that this system only allows access to a specific group of users with shared interests and use cases.

This type of cloud architecture can be hosted on-premises, at a peer organization, or by a third-party provider. A combination of all three is also an option.

Typically, all organizations in a community have the same security policies, application types, and legislative issues.

Advantages of Community Cloud

Here are the benefits of a community cloud solution:

- **Cost reductions:** A community cloud is cheaper than a private one, yet it offers comparable performance. Multiple companies share the bill, which additionally lowers the cost of these solutions.
- **Setup benefits:** Configuration and protocols within a community system meet the needs of a specific industry. A collaborative space also allows clients to enhance efficiency.

Disadvantages of Community Cloud

The main disadvantages of community cloud are:

- **Shared resources:** Limited storage and bandwidth capacity are common problems within community systems.
- **Still uncommon:** This is the latest deployment model of cloud computing. The trend is still catching on, so the community cloud is currently not an option in every industry.

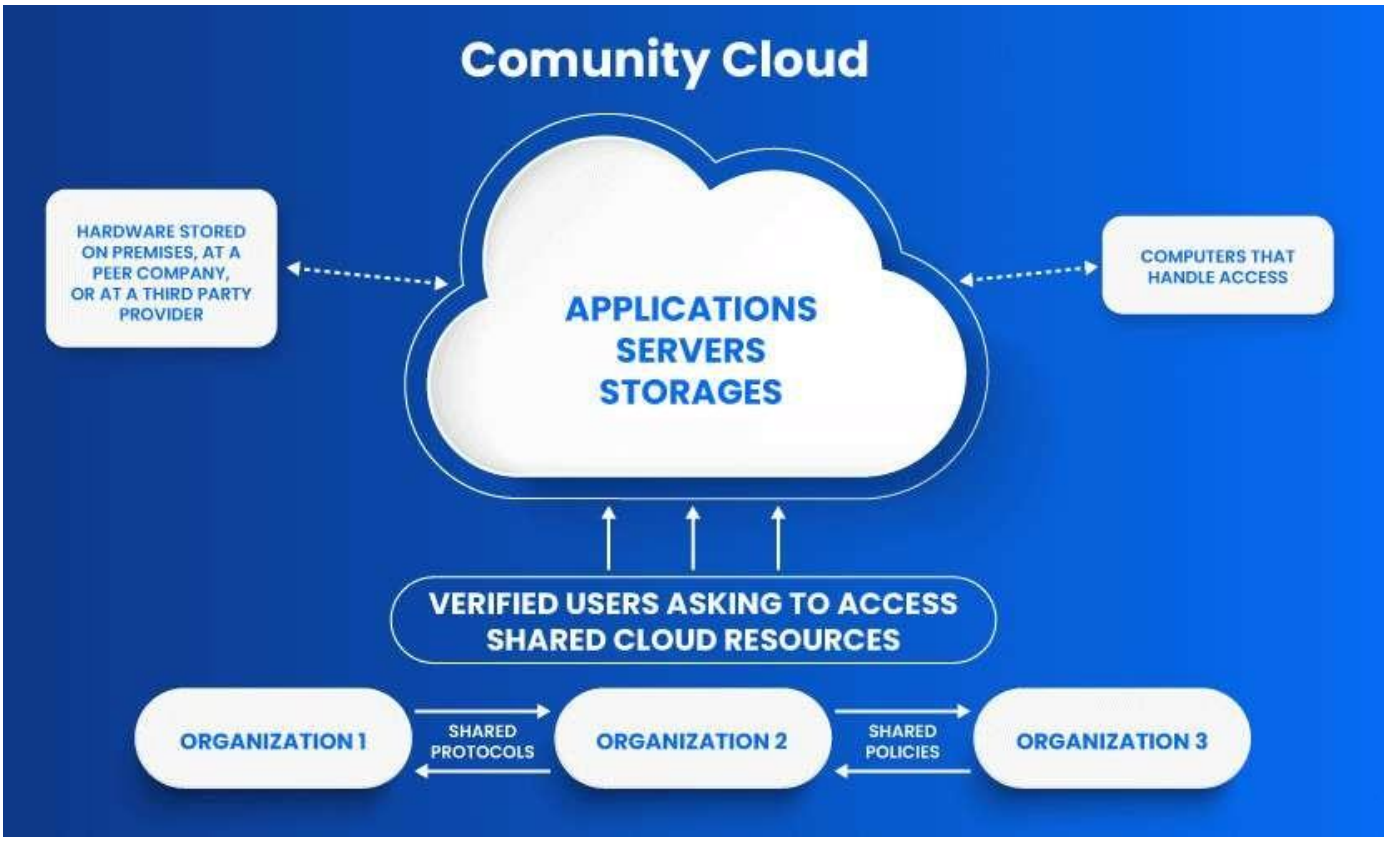

Hybrid Cloud

A hybrid cloud is a combination of two or more infrastructures (private, community, VPC, public cloud, and dedicated servers). Every model within a hybrid is a separate system, but they are all a part of the same architecture.

A typical deployment model example of a hybrid solution is when a company stores critical data on a private cloud and less sensitive information on a public cloud. Another use case is when a portion of a firm's data cannot legally be stored on a public cloud.

The hybrid cloud model is often used for cloud bursting. Cloud bursting allows an organization to run applications on-premises but "burst" into the public cloud in times of heavy load. It is an excellent option for organizations with versatile use cases.

Advantages of Hybrid Cloud

Here are the benefits of a hybrid cloud system:

- **Cost-effectiveness:** A hybrid solution lowers operational costs by using a public cloud for most workflows.
- **Security:** It is easier to protect a hybrid cloud from attackers due to segmented storage and workflows.
- **Flexibility:** This cloud model offers high levels of setup flexibility. Clients can create custom-made solutions that fit their needs entirely.

Disadvantages of Hybrid Cloud

The disadvantages of hybrid solutions are:

- **Complexity:** A hybrid cloud is complex to set up and manage as you combine two or more different cloud service models.
- **Specific use case:** A hybrid cloud makes sense only if an organization has versatile use cases or need to separate sensitive and non-sensitive data.

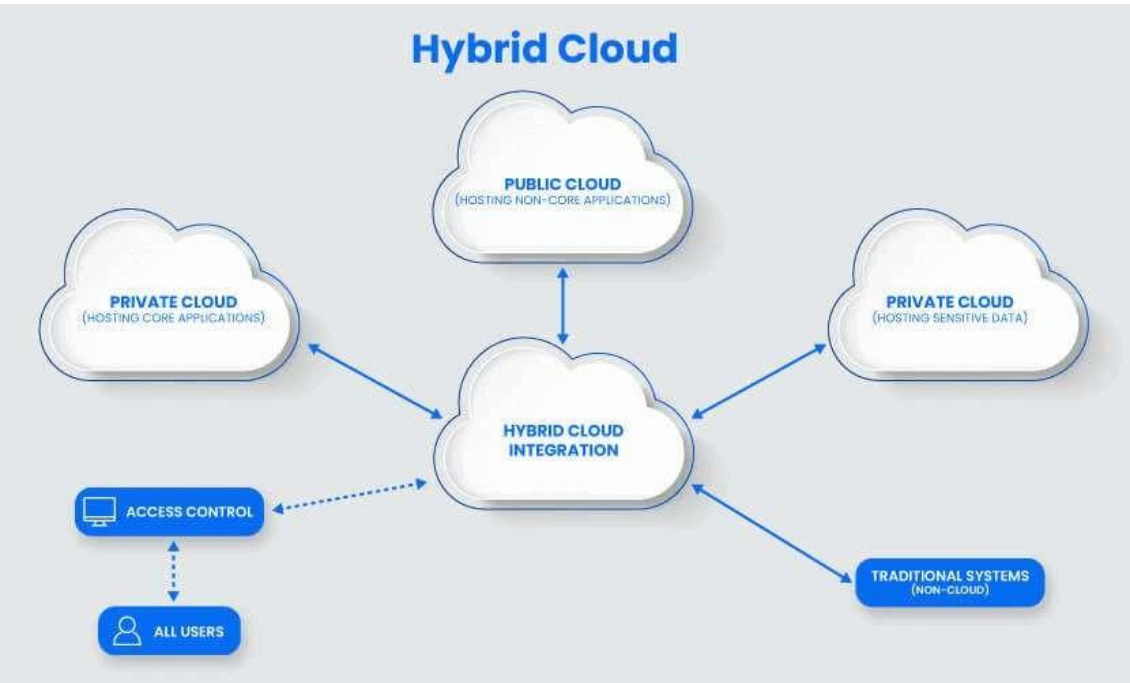

Comparison of Cloud Deployment Models

Here is a comparative table that provides an overview of all five cloud deployment models:

	Public	**Private**	**VPC**	**Community**	**Hybrid**
Ease of setup	Very easy to set up, the provider does	Very hard to set up as your team creates the	Easy to set up, the provider does	Easy to set up because of	Very hard to set up due to interconnected systems

	most of the work	system	most of the work (unless the client asks other wise)	community practices	
Ease of use	Very easy to use	Complex and requires an in-house team	Easy to use	Relatively easy to use as members help solve problems and establish protocols	Difficult to use if the system was not set up properly
Data control	Low, the provider has all control	Very high as you own the system	Low, the provider has all control	High (if members collaborate)	Very high (with the right setup)
Reliability	Prone to failures and outages	High (with the right team)	Prone to failures and outages	Depends on the community	High (with the right setup)

Scalability	Low, most providers offer limited resources	Very high as there are no other system tenants	Very high as there are no other tenants in your segment of the cloud	Fixed capacity limits scalability	High (with the right setup)
Security and privacy	Very low, not a good fit for sensitive data	Very high, ideal for corporate data	Very low, not a good fit for sensitive data	High (if members collaborate on security policies)	Very high as you keep the data on a private cloud
Setup flexibility	Little to no flexibility, service providers usually offer only predefined setups	Very flexible	Less than a private cloud, more than a public one	Little flexibility, setups are usually predefined to an extent	Very flexible
Cost	Very Inexpensive	Very expensive	Affordable	Members share the costs	Cheaper than a private model, pricier than a public one

Demand for in-house hardware	No	In-house hardware is not a must but is preferable	No	No	In-house hardware isnot a must but is preferable

How to Choose Between Cloud Deployment Models

To choose the best cloud deployment model for your company, start by defining your requirements for:

- **Scalability:** Is your user activity growing? Does your system run into sudden spikes in demand?
- **Ease of use:** How skilled is your team? How much time and money are you willing to invest in staff training?
- **Privacy:** Are there strict privacy rules surrounding the data you collect?
- **Security:** Do you store any sensitive data that does not belong on a public server?
- **Cost:** How much resources can you spend on your cloud solution? How much capital can you pay upfront?
- **Flexibility:** How flexible (or rigid) are your computing, processing, and storage needs?
- **Compliance:** Are there any notable laws or regulations in your country or industry? Do you need to adhere to compliance standards?

1.5. Cloud computing environment

1.6. Cloud Service requirements

1.5. Cloud computing environment

- **Personal cloud computing environment**
  Single computer is used by single person. All the hardware devices are present at single location and packed as single unit.
- **Time sharing environment**
  Multiple computers are connected to a single large computers called server. Server share resources to all other system.
- **Client-server environment**
  The environment is similar to time sharing environment. Here server provides website.
- **Distributed computing environment**
  A single task is divided into parts and each part is executed by a different system. After execution the result is merged to get the final result

1.6. Cloud Service requirements

Requirements for Building a Cloud Infrastructure

When building out a cloud strategy, there are several in-depth steps that must be taken to ensure a robust infrastructure.

Requirement 1: Service and Resource Management

A cloud infrastructure virtualizes all components of a data center. Service management is a measured package of applications and services that end users can easily deploy and manage via a public and/or private cloud vendor. And a simplified tool to outline and gauge services is vital for cloud administrators to market functionality. Service management needs to contain resource maintenance, resource guarantees, billing cycles, and measured regulations. Once deployed, management services should help create policies for data and workflows to make sure it's fully efficient and processes are delivered to systems in the cloud.

Requirement 2: Data Center Management Tools Integration

Most data centers utilize a variety of IT tools for systems management, security, provisioning, customer care, billing, and directories, among others. And these work with cloud management services and open APIs to integrate existing operation, administration, maintenance, and provisioning (OAM&P) systems. A modern cloud service should support a data center's existing infrastructure as well as leveraging modern software, hardware, and virtualization, and other technology.

Requirement 3: Reporting, Visibility, Reliability, a Security

Data centers need high levels of real-time reporting and visibility capabilities in cloud environments to guarantee compliance, SLAs, security, billing, and chargebacks. Without robust reporting and visibility, managing system performance, customer service, and other processes are nearly impossible. And to be wholly reliable, cloud infrastructures must operate regardless of oneor more failing components. For to safeguard the cloud, services must ensure data and apps are secure while providing access to those who are authorized.

Requirement 4: Interfaces for Users, Admins, and Developers

Automated deployment and self-service interfaces ease complex cloud services for end users, helping lower operating costs and deliver adoption. Self-service interfaces offer customers the ability to effectively launch a cloud service by managing their own data centers virtually, designing and driving templates, maintaining virtual storage, networking resources, and utilizing libraries. Administrator interfaces present better visibility to all resources, virtual machines, templates, service offers, and various cloud users. And all of these structures integrate by way of APIs for developers.

Advantages of Using Cloud Infrastructure

The arguments in favour of using the cloud are only getting stronger as the technology continues to improve. So, there are some obvious key benefits to migrating to a cloud infrastructure that helps companies streamline business processes.

Cost: First and foremost, the cloud removes or greatly reduces the operating expense of a company setting up and managing its own data center. Taking on this process begins to add up with all the various hardware, software, servers, energy bills, IT experts, and the updates that come along with this multi-faceted set-up. With cloud infrastructure, a company simply pays for it all to be managed while paying only for as-needed services.

Agility and flexibility: Most cloud service infrastructures are offered as self-managed, where service changes can be made within minutes. This improves the uptime and efficiency of business systems while allowing off-site co-workers and partners to access shared data on mobile devices whenever and wherever. And with a cloud infrastructure managing processes, a company becomes more business-focused than IT-focused.

Security: There's a common misconception that cloud services are generally not secure and that data can easily be compromised. There is some truth in that, however, the risks are often blown out of proportion at least in terms of enterprise-level cloud infrastructure and services. Cloud infrastructure technologies and providers are always improving protection against hackers, viruses, and other data breaches with stronger firewalls, advanced encryption keys, and a hybrid approach that stores sensitive data in a private cloud and other data, even apps, in a public cloud.

Disadvantages of Using Cloud Infrastructure

That being said, not all cloud infrastructures are perfect. And while there are far more advantages, there are still some drawbacks.

Vendor overturn: The cloud is still an evolving, albeit improving, technology that rapidly fluctuates. Meaning, some cloud services companies get it right and some don't. If a company goes out of business or sees a massive overhaul, that could be destructive to a business that relies on justone infrastructure for its entire database.

Connection reliance: A cloud infrastructure is only as good as its network connection. Therefore, the cloud can't stay afloat without a dependable connection. Any glitches in an internet or intranet connection due to a technical outage or storm mean the cloud goes down along with all the data, software, and/or applications in it. A reliable network means business promises and SLAs are delivered.

Control: Since a company's cloud infrastructure is generally controlled by its service provider, there are times organizations have limited access to data. And business customers have even less control than they might want, with limited access to applications, data, and tools stored on a server.

BASIC REQUIREMENTS OF A CLOUD COMPUTING SERVICE

Cloud computing is here to rule. Right now, most of the small, medium enterprises have gone 100% on cloud. I have seen several start-ups - which are using cloud services for all their computing needs. But large enterprises are reluctant to move to cloud services and rightly so. Many companiesare just testing waters and have held back on full scale deployment of cloud IT services.

Today the cloud services have several deficiencies - which from an enterprise perspective are the basic requirements for them to consider cloud services. In this article I have written about 6 basic requirements for enterprises to adapt cloud services in a big way.

1. Availability - with loss less DR

Customers want their IT services be up and available at all times. But in reality, computers sometimes fail. This implies that the service provider should have implemented a reliable disaster recovery (DR) mechanism - where in the service provider can move the customer from one data center to another seamlessly and the customer does not even have to know about it.

As a cloud service provider, there will be enormous pressure to minimise costs by optimally utilizing the entire IT infrastructure. The traditional Active-Passive DR strategy is very expensive and cost inefficient. Instead, service providers will have to create an Active-Active disaster recovery mechanism - where more than one data centre will be active at all times and ensures that the data and services can be accessed by the customer from either of the data centres seamlessly.

Today, there are several solutions that are available to do just that. EMC VPLEX solution to maintain an Active-Active data centre. Another approach will be implementing Hadoop/Hive stack for data intensive applications such as emails, messaging, data store, services.

In an ideal scenario, the customer on the cloud services should not even notice any change at all and the movement of all his data & applications from one data center to another must be transparent to the end user.

2. Portability of Data & Applications

Customers hate to be locked into a service or a platform. Ideally a cloud offering must be able to allow customers to move out their data & applications from one service provider to another - just like customers can switch from one telephone service provider to another.

As applications are being written on standard platforms - Java, PHP, Python, etc. It should be possible to move the customer owned applications from one service provider to another. Customers should also take care to use only the open standards and tools, and avoid vendor specifictools. Azure or Google services offers several tools/applications/utilities which are valuable - but it also creates a customer locking - as the customer who uses these vendors specific tools cannot migrate to another service provider without rewriting the applications.

To illustrate this, today in India, customers can move from one cell phone service provider to another without changing their handsets, but in US, if one were to move from AT&T to Verizon, one need to pay for the handset - which forms a customer lock in instrument.

With public cloud services, customers should be able to move their data & applications from one cloud to another - without disrupting the end user's IT services. This movement should be transparent to the end user.

The Cloud Computing Interoperability Forum (CCIF) was formed by organizations such as Intel, Sun, and Cisco in order to enable a global cloud computing ecosystem whereby organizations are able to seamlessly work together for the purposes for wider industry adoption of cloud computing technology. The development of the Unified Cloud Interface (UCI) by CCIF aims at creating a standard programmatic point of access to an entire cloud infrastructure.

Recently in EMC world 2011, EMC demonstrated moving several active VMs & applications from EMC data center to CSC data center without disruption of service. This was just a proof of concept, but to make this a common place, some amount of regulation and business coordination will be required.

However, in their current form, most of cloud computing services and platforms do not employ standard methods of storing user data and applications. Consequently, they do not interoperate and user data are not portable.

3. Data Security

Security is the key concern for all customers - since the applications and the data is residing in the public cloud, it is the responsibility of the service provider for providing adequate security. In my opinion security for customer data/applications becomes a key differentiator when it comes to selecting the cloud service provider. When it comes to IT security, customers tend to view the cloud service providers like they view banks. The service provider is totally responsible for user security, but there are certain responsibilities that the customer also needs to take.

The service provider must a robust Information Security Risk Management process - which is well understood by the customer, and customer must clearly know his responsibilities as well. As there are several types of cloud offerings (SaaS, PaaS, IaaS etc), there will be different sets of responsibility for the customer and the service provider depending on the cloud service offering.

When it comes to security, the cloud service providers offer better security than what the customer's own data center security. This is a kin to banks - where banks can offer far greater security than any individual or company. The security in cloud is much higher due to: Centralized monitoring enhanced incidence detection/forensics, logging of all activity, greater security/venerability testing, centralized authentication testing (aka password protection/assurance), secure builds & testing patches before deployment and lastly better securitysoftware/systems.

Cloud service providers know that the security is the key to their success - and hence invest more on security. The amount of efforts/money invested by cloud service providers will always be greater than the amount an individual company (most) can spend.

Security issues will also be addressed through legal & regulatory systems. Despite the best IT security, breaches can happen and when it happens, the laws and rules of the land - where the data resides play an important role. For example, specific cryptography techniques could not be used because they are not allowed in some countries. Similarly, country laws can impose that sensitive data, such as patient health records, are to be stored within national borders. Therefore customer needs to pay attention to Legal and regulatory issues when selecting the service providers.

4. Manageability

Managing the cloud infrastructure from the customer perspective must be under the control of the customer admin. Customers of Cloud services must be able to create new accounts, must be able to provision various services, do all the user account monitoring - monitoring for end user usage, SLA breaches, data usage monitoring etc. The end users would like to see the availability, performance and configuration/provisioning data for the set of infrastructure they are using in the cloud.

Cloud service provider will have various management tools for Availability management, performance management, configuration management and security management of applications and infrastructure (storage, servers, and network). Customers want to know how the entire infrastructure is being managed - and if possible can that management information be shared with them, and alert the customer on any outage, slow service, or breach of SLA as it happens. This allows customer to take corrective actions - either move the applications to another cloud or enabletheir contingency plans.

Sharing the application performance and resource management information will help improve utilization and consequently optimize usage by customers. This will result in improving ROI for the customers and encourage customers to adapt cloud services.

As customers buy cloud services from multiple vendors, it will become a necessity to have a unified management system to manage all the cloud services they have. This implies that cloud service providers must embrace an XML based reporting formats to provide management information to customers and customers then can build their own management dashboards.

5. Elasticity

Customer on Cloud computing have a dynamic computing loads. At times of high load, they need greater amount of computing resources available to them on demand, and when the work loads are low, the computing resources are released back to the cloud pool. Customer expects the service provider to charge them for what they have actually used in the process.

Customers also want a self service on-demand resource provisioning capability from the service provider. 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. This eliminates the need for more time-consuming, labour-intensive, human driven procurement processes familiar to many in IT.

This implies that the dynamic provisioning system should be the basic part of cloud management software - through which users can easily interact with the system.

To provide elastic computing resources, the service provider must be able to dynamically provision resources as needed and have adequate charge back systems to bill the customer.

In reality, it may not be possible for any single cloud service provider to build an infinitely scalable infrastructure and hence customers will have to rely on a federated system of multiple cloudservice providers sharing the customer loads. (Just like a power grid, where the load gets distributed to other power plants during peak loads)

6. Federated System

There are several reasons as to why customers will need a Federated cloud system. Customers may have to buy services from several cloud service providers for various services - email from Google, online sales transaction services from Amazon and ERP from another vendor etc. In such cases customer want their cloud applications to interact with other services from several vendors to provide a seamless end to end IT services.

This implies that each of the cloud services must have an interface with other cloud services for load sharing & application interoperability.

In a federated environment there is potentially an infinite pool of resources. To build such a system, there should be inter-cloud framework agreements between multiple service providers, and adequate chargeback systems in place.

Having a federated system helps customers to move their data/applications across different cloud service providers and prevents customer locking.

Interoperability of applications across different cloud services has led to creations of standard APIs. But these APIs are cumbersome to use and that has led to creation of Cloud Integration Bus - based on Enterprise Service Bus (ESB).

1.7. Cloud and Dynamic Infrastructure
1.8. Cloud Adoption
1.9. Cloud applications

CLOUD AND DYNAMIC INFRASTRUCTURE

What is cloud infrastructure?

Cloud infrastructure is a term used to describe the components needed for cloud computing, which includes hardware, abstracted resources, storage, and network resources. Think of cloud infrastructure as the tools needed to build a cloud. In order to host services and applications in the cloud, you need cloud infrastructure.

Build a cloud infrastructure using IT resources you already have

How does cloud infrastructure work?

An abstraction technology or process—like virtualization—is used to seperate resources from physical hardware and pool them into clouds; automation software and management tools allocate these resources and provision new environments so users can access what they need—when they need it.

What's included in cloud infrastructure?

Cloud infrastructure is made up of several components, each integrated with one another into a single architecture supporting business operations. A typical solution may be composed of hardware, virtualization, storage, and networking components.

As a term, cloud infrastructure can be used to describe a complete cloud computing system—once all the pieces are put together—as well as the individual technologies themselves.

Components of cloud infrastructure

Hardware

Although you probably think of clouds as being virtual, they require hardware as part of the infrastructure.

A cloud network is made up of a variety of physical hardware that can be located at multiple geographical locations.

The hardware includes networking equipment, like switches, routers, firewalls, and load balancers, storage arrays, backup devices, and servers.

Virtualization connects the servers together, dividing and abstracting resources to make them accessible to users.

Virtualization

Virtualization is technology that separates IT services and functions from hardware.

Software called a hypervisor sits on top of physical hardware and abstracts the machine's resources, such as memory, computing power, and storage.

Once these virtual resources are allocated into centralized pools they’re considered clouds.

With clouds, you get the benefits of self-service access, automated infrastructure scaling, and dynamic resource pools.

Storage

Within a single datacenter, data may be stored across many disks in a single storage array. Storage management ensures data is correctly being backed up, that outdated backups are removed regularly, and that data is indexed for retrieval in case any storage component fails.

Virtualization abstracts storage space from hardware systems so that it can be accessed by users as cloud storage.

When storage is turned into a cloud resource, you can add or remove drives, repurpose hardware, and respond to change without manually provisioning separate storage servers for every new initiative.

Network

The network is composed of physical wires, switches, routers, and other equipment. Virtual networks are created on top of these physical resources.

A typical cloud network configuration is composed of multiple subnetworks, each with varying levels of visibility. The cloud permits the creation of virtual local area networks (VLANs) andassigns static and/or dynamic addresses as needed for all network resources.

The cloud resources are delivered to users over a network, such as the internet or an intranet, so you can access cloud services or apps remotely on demand.

PUBLIC, PRIVATE, AND HYBRID CLOUD INFRASTRUCTURE

The basic elements of cloud infrastructure are the same whether you have a private cloud, public cloud, or a combination.

To get started with any of the cloud computing types, you need a cloud infrastructure. You can create a private cloud by building it yourself using resources dedicated solely to you, or you can usea public cloud by renting the cloud infrastructure from a cloud provider so you don't have to set it up yourself.

What is Dynamic Infrastructure?

Dynamic Infrastructure is an IT system wherein the design of **data centers** are such that the underlying hardware and software layers can respond dynamically to changing levels of demand in more fundamental and efficient ways. It is also known as **Infrastructure** 2.0 and Next Generation **Data Center.** The basic premise of **Dynamic Infrastructures** is that leverage pooled IT resources can provide flexible IT capacity, enabling the seamless, real-time allocation of IT resources in line with demand from business processes. This is achieved by using server virtualization technology to pool computing resources wherever possible, and allocating these resources on-demand using automated tools. This allows for **load balancing** and is a more efficient approach than keeping massive computing resources in reserve to run tasks that take place.

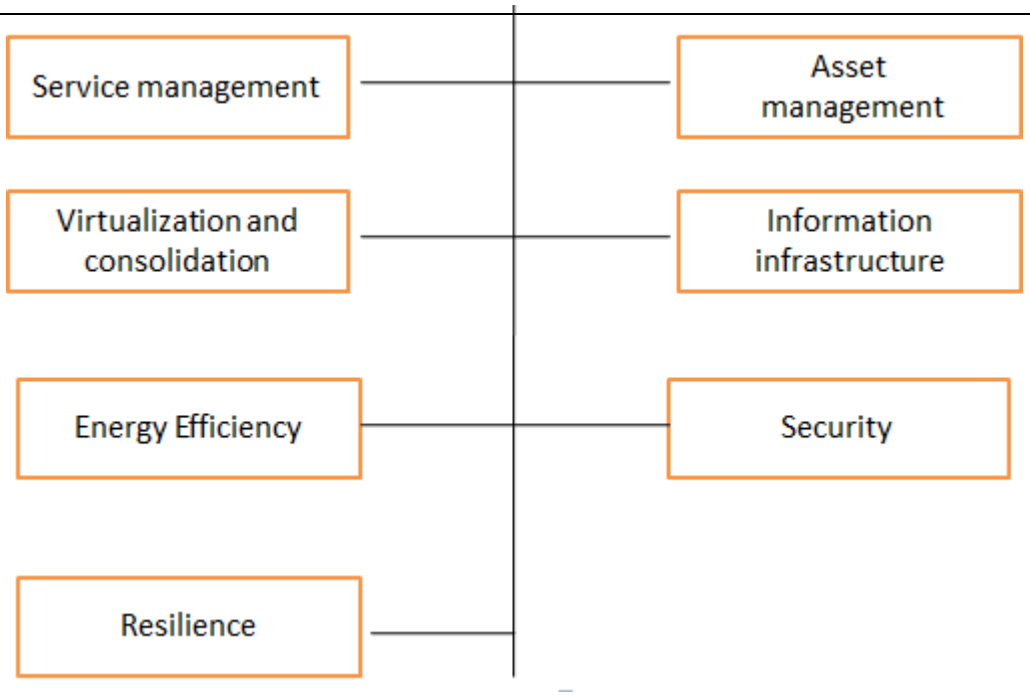

SERVICE MANAGEMENT :

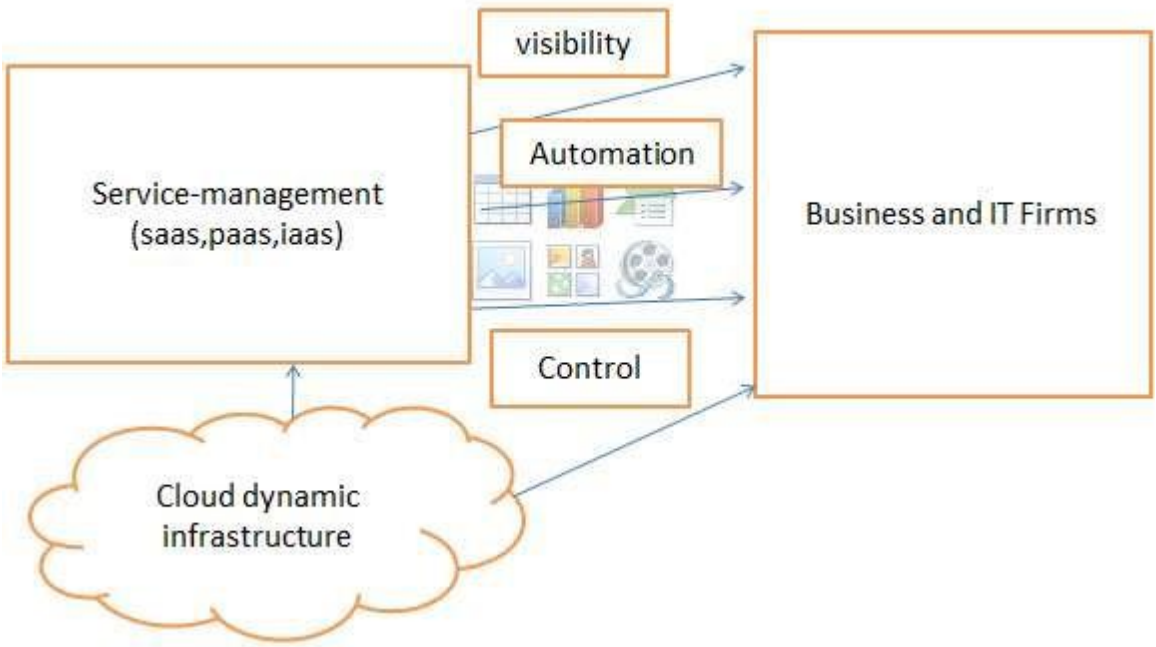

- This type of special facility or a functionality is provided to the cloud IT services by the cloud service providers.
- This facility includes visibility, automation and control to delivering the first class IT services.

ASSET MANAGEMENT:

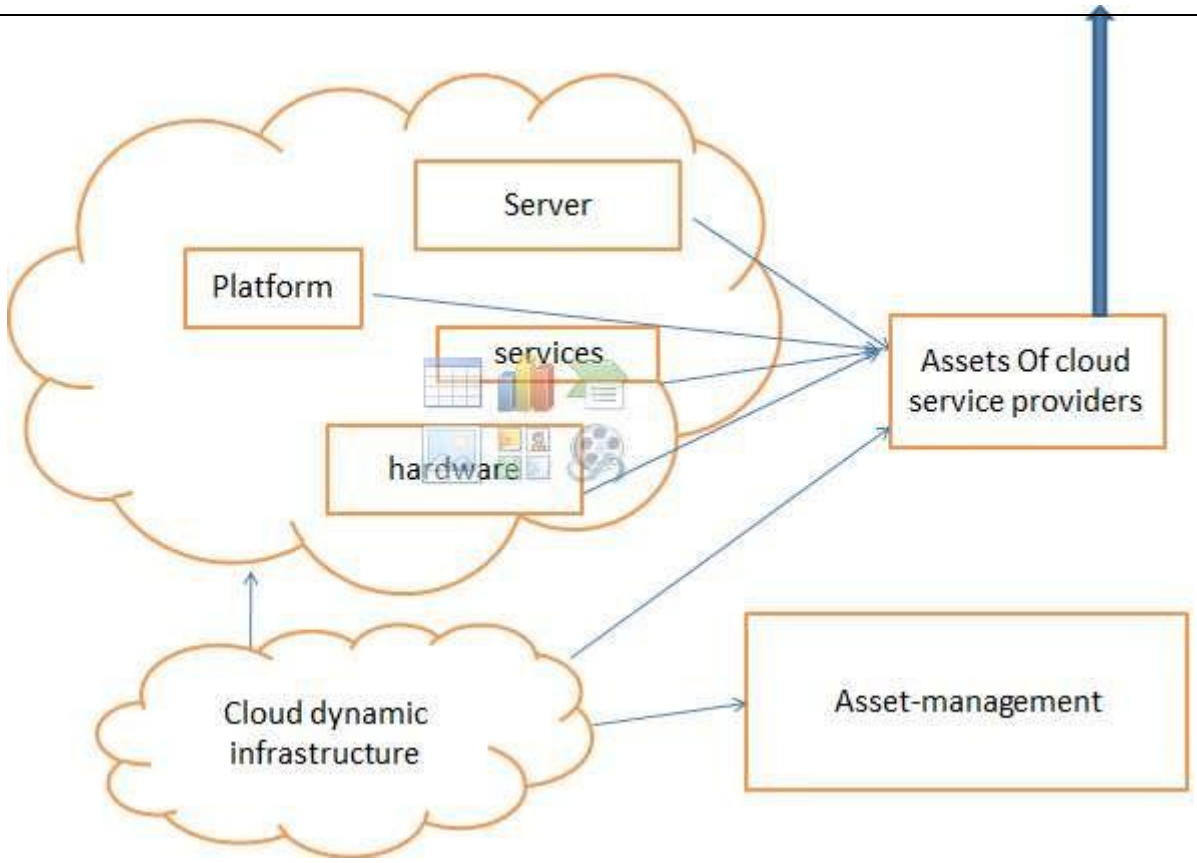

- In this actually the assets or the property which is involved in providing the cloud services are getting managed.
- They are getting managed in such a way so that their value will got maximized.

VIRTUALIZATION AND CONSOLIDATION:

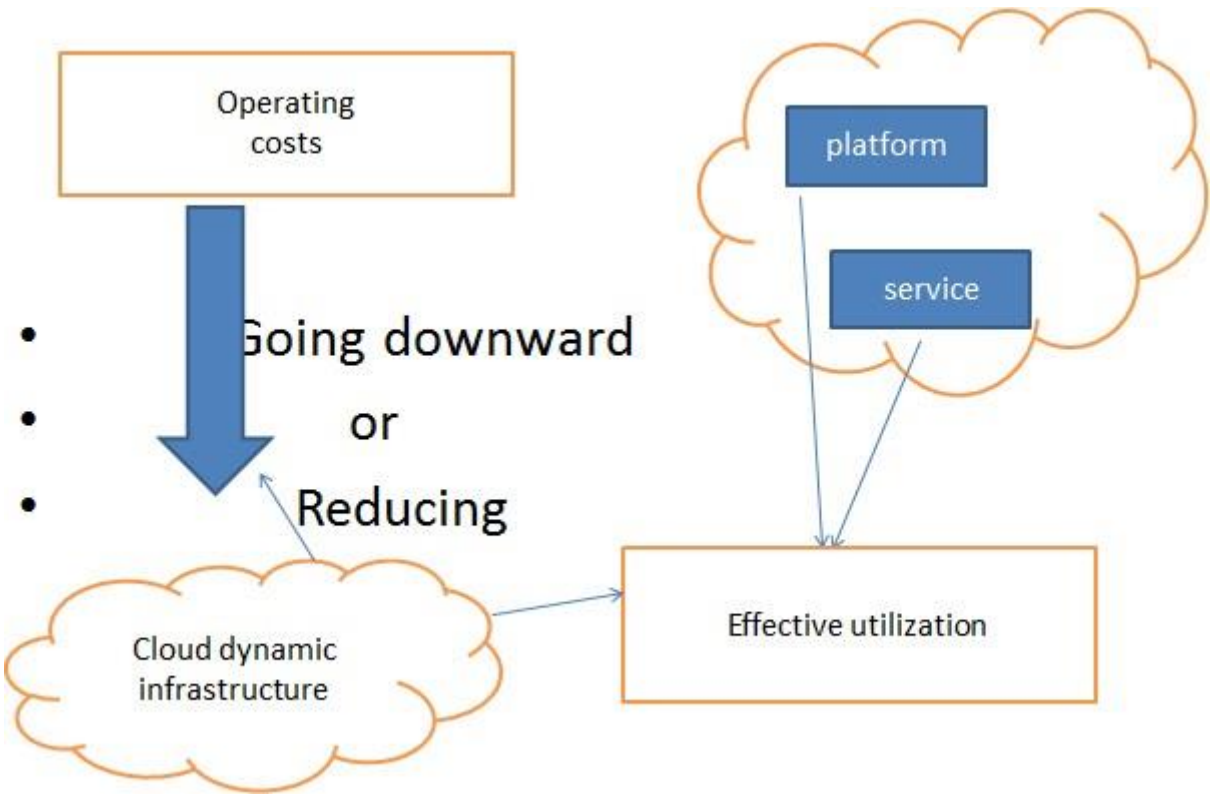

- Here in the above fig it is clearly stated that ”resources are getting utilized more and more efficiently.

- Also the operating cost of the systems is getting down.

INFORMATION STRUCTURE:

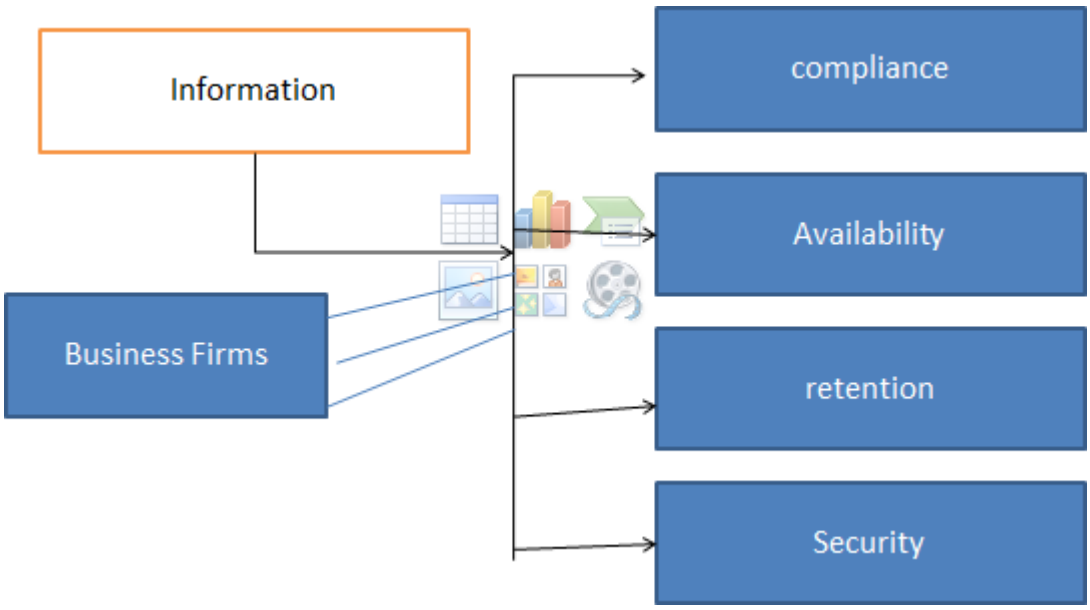

It helps the business organizations to achieve the following :

- Information compliance,
- availability of resources
- retention and security objectives.

SECURITY:

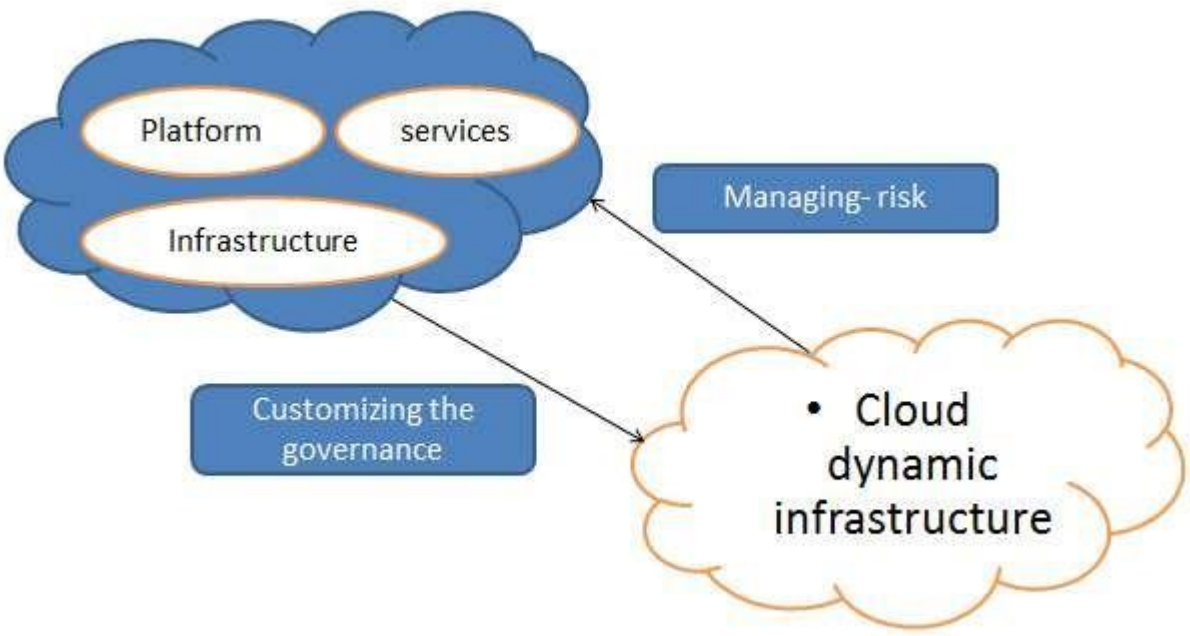

- This cloud infrastructure is responsible for the risk management, customizing the governance.
- Risk management Refers to the risks involves in the services which are being provided by the cloud-service providers.

- Customization of governance implies that the features of the governing body or admin body can be changed but these changes are totally depends on the providers wish.

RESILIENCE :

- This infrastructure provides the feature of resilience means the services are resilient.
- It means the infrastructure is safe from all side.
- The IT operations will not be easily get affected.

ENERGY-EFFICIENCY:

- Here the IT infrastructure or organization sustainable.
- It means it is not likely to damage or effect any other thing.

CLOUD ADOPTION

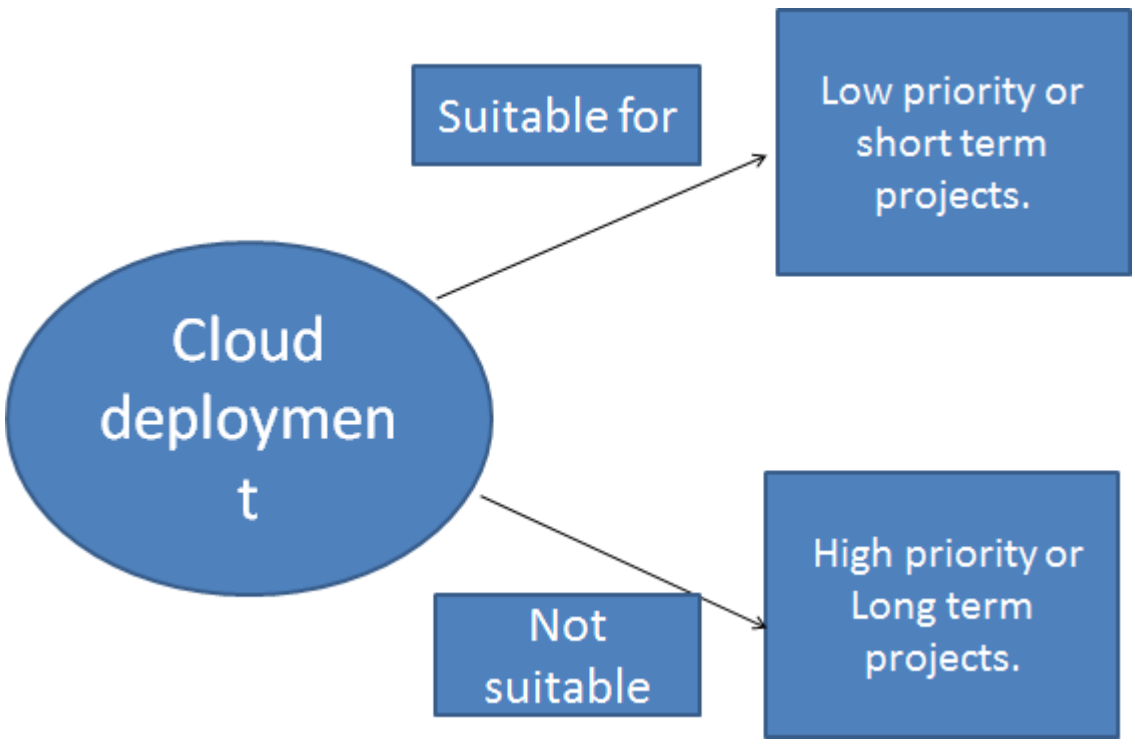

- Here Cloud means The environment of cloud where the cloud services are being operated.

- Adoption term states that accepting the services of new Technology.

- Adoption means following some kind of new trend or existing trend or a technology.

- This Cloud adoption is suitable for low priority business applications.

- Because as we have already discussed that the cloud computing is not beneficial for long term projects.

- It supports some interactive applications that combines two or more data sources.

- These applications must having low availability requirements and short life spans.

- For example:-if a marketing company requires to grow his business in the whole country in a short span of time then it must need a quick promotion or short promotion across the country.

- Cloud Adoption is useful when the recovery management, backup recovery based implementations are required.

- By considering the above key points we conclude that it is only suitable for the applications that are modular and loosely coupled.

- It will works well with research and development projects.

- It means the testing of new services ,design models and also the applications that can be get adjusted on small servers.

- Applications which requires different level of infrastructure throughout the day or throughout the month should be deployed Through the cloud.

- The applications whose demand is unknown can also be deployed using clouds.

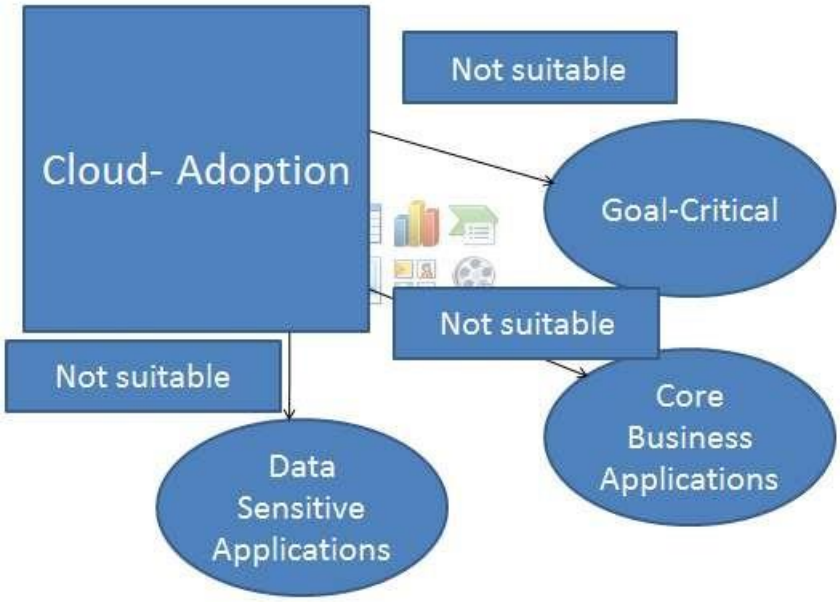

Cloud Computing Applications

Cloud service providers provide various applications in the field of art, business, data storage and backup services, education, entertainment, management, social networking, etc.

The most widely used cloud computing applications are given below -

1. Art Applications

Cloud computing offers various art applications for quickly and easily design **attractive cards, booklets,** and **images**. Some most commonly used cloud art applications are given below:

i Moo

Moo is one of the best cloud art applications. It is used for designing and printing business cards, postcards, and mini cards.

ii. Vistaprint

Vistaprint allows us to easily design various printed marketing products such as business cards, Postcards, Booklets, and wedding invitations cards.

iii. Adobe Creative Cloud

Adobe creative cloud is made for designers, artists, filmmakers, and other creative professionals. It is a suite of apps which includes PhotoShop image editing programming, Illustrator, InDesign, TypeKit, Dreamweaver, XD, and Audition.

2. Business Applications

Business applications are based on cloud service providers. Today, every organization requires the cloud business application to grow their business. It also ensures that business applications are 24*7 available to users.

There are the following business applications of cloud computing -

i. MailChimp

MailChimp is an **email publishing platform** which provides various options to **design, send,** and **save** templates for emails.

iii. Salesforce

Salesforce platform provides tools for sales, service, marketing, e-commerce, and more. It also provides a cloud development platform.

iv. Chatter

Chatter helps us to **share important information** about the organization in real time.

v. Bitrix24

Bitrix24 is a **collaboration** platform which provides communication, management, and social collaboration tools.

vi. Paypal

Paypal offers the simplest and easiest **online payment** mode using a secure internet account. Paypal accepts the payment through debit cards, credit cards, and also from Paypal account holders.

vii. Slack

Slack stands for **Searchable Log of all Conversation and Knowledge**. It provides a **user- friendly** interface that helps us to create public and private channels for communication.

viii. Quickbooks

Quickbooks works on the terminology "**Run Enterprise anytime, anywhere, on any device**." It provides online accounting solutions for the business. It allows more than 20 users to work simultaneously on the same system.

3. Data Storage and Backup Applications

Cloud computing allows us to store information (data, files, images, audios, and videos) on the cloud and access this information using an internet connection. As the cloud provider is responsiblefor providing security, so they offer various backup recovery application for retrieving the lost data.

A list of data storage and backup applications in the cloud are given below -

i. Box.com

Box provides an online environment for **secure content management, workflow,** and **collaboration**. It allows us to store different files such as Excel, Word, PDF, and images on the cloud. The main advantage of using box is that it provides drag & drop service for files and easily integrates with Office 365, G Suite, Salesforce, and more than 1400 tools.

ii. Mozy

Mozy provides powerful **online backup solutions** for our personal and business data. It schedules automatically back up for each day at a specific time.

iii. Joukuu

Joukuu provides the simplest way to **share** and **track cloud-based backup files**. Many users use joukuu to search files, folders, and collaborate on documents.

iv. Google G Suite

Google G Suite is one of the best **cloud storage** and **backup** application. It includes Google Calendar, Docs, Forms, Google+, Hangouts, as well as cloud storage and tools for managing cloud apps. The most popular app in the Google G Suite is Gmail. Gmail offers free email services to users.

4. Education Applications

Cloud computing in the education sector becomes very popular. It offers various **online distance learning platforms** and **student information portals** to the students. The advantage of using cloud in the field of education is that it offers strong virtual classroom environments, Ease of accessibility, secure data storage, scalability, greater reach for the students, and minimal hardware requirements for the applications.

There are the following education applications offered by the cloud -

i. Google Apps for Education

Google Apps for Education is the most widely used platform for free web-based email, calendar, documents, and collaborative study.

ii. Chromebooks for Education

Chromebook for Education is one of the most important Google's projects. It is designed for the purpose that it enhances education innovation.

iii. Tablets with Google Play for Education

It allows educators to quickly implement the latest technology solutions into the classroom and make it available to their students.

iv. AWS in Education

AWS cloud provides an education-friendly environment to universities, community colleges, and schools.

5. Entertainment Applications

Entertainment industries use a **multi-cloud strategy** to interact with the target audience. Cloud computing offers various entertainment applications such as online games and video conferencing.

i. Online games

Today, cloud gaming becomes one of the most important entertainment media. It offers various online games that run remotely from the cloud. The best cloud gaming services are Shaow, GeForce Now, Vortex, Project xCloud, and PlayStation Now.

ii. Video Conferencing Apps

Video conferencing apps provides a simple and instant connected experience. It allows us to communicate with our business partners, friends, and relatives using a cloud-based video conferencing. The benefits of using video conferencing are that it reduces cost, increases efficiency, and removes interoperability.

6. Management Applications

Cloud computing offers various cloud management tools which help admins to manage all types of cloud activities, such as resource deployment, data integration, and disaster recovery. These management tools also provide administrative control over the platforms, applications, and infrastructure.

Some important management applications are -

i. Toggle

Toggle helps users to track allocated time period for a particular project.

ii. Evernote

Evernote allows you to sync and save your recorded notes, typed notes, and other notes in one convenient place. It is available for both free as well as a paid version.

It uses platforms like Windows, macOS, Android, iOS, Browser, and Unix.

iii. Outright

Outright is used by management users for the purpose of accounts. It helps to track income, expenses, profits, and losses in real-time environment.

iv. GoToMeeting

GoToMeeting provides **Video Conferencing** and **online meeting apps**, which allows you to start a meeting with your business partners from anytime, anywhere using mobile phones or tablets. Using GoToMeeting app, you can perform the tasks related to the management such as join meetings in seconds, view presentations on the shared screen, get alerts for upcoming meetings,etc.

7. Social Applications

Social cloud applications allow a large number of users to connect with each other using social networking applications such as **Facebook, Twitter, Linkedln,** etc.

There are the following cloud based social applications -

i. Facebook

Facebook is a **social networking website** which allows active users to share files, photos, videos, status, more to their friends, relatives, and business partners using the cloud storage system. On Facebook, we will always get notifications when our friends like and comment on the posts.

ii. Twitter

Twitter is a **social networking** site. It is a **microblogging** system. It allows users to follow high profile celebrities, friends, relatives, and receive news. It sends and receives short posts called tweets.

iii. Yammer

Yammer is the **best team collaboration** tool that allows a team of employees to chat, share images, documents, and videos.

iv. LinkedIn

LinkedIn is a **social network** for students, freshers, and professionals.

Real World Applications of Cloud Computing

Applications of Cloud computing in real-world

Cloud Service Providers (CSP) are providing many types of cloud services and now if we sill cloud computing has touched every sector by providing various cloud applications. Sharing and managing resources is easy in cloud computing that's why it is one of the dominant fields of computing. These properties have made it an active component in many fields. Now let's know some of the real-world applications of cloud computing.

1. Online Data Storage

Cloud computing allows storing data like files, images, audios, and videos, etc on the cloud storage. The organization need not set physical storage systems to store a huge volume of business data which costs so high nowadays. As they are growing technologically, data generation is also growing with respect to time, and storing that becoming problem. In that situation, Cloud storage is providing this service to store and access data any time as per requirement.

2. Backup and Recovery

Cloud vendors provide security from their side by storing safe to the data as well as providing a backup facility to the data. They offer various recovery application for retrieving the lost data. In the traditional way backup of data is a very complex problem and also it is very difficult sometimes impossible to recover the lost data. But cloud computing has made backup and recovery applications very easy where there is no fear of running out of backup media or loss of data.

3. Bigdata Analysis

We know the volume of big data is so high where storing that in traditional data management system for an organization is impossible. But cloud computing has resolved that problem by allowing the organizations to store their large volume of data in cloud storage without worrying about physical storage. Next comes analyzing the raw data and finding out insights or useful information from it is a big challenge as it requires high-quality tools for data analytics. Cloud computing provides the biggest facility to organizations in terms of storing and analyzing big data.

4. Testing and development

Setting up the platform for development and finally performing different types of testing to check the readiness of the product before delivery requires different types of IT resources and infrastructure. But Cloud computing provides the easiest approach for development as well as testing even if deployment by using their IT resources with minimal expenses. Organizations find it more helpful as they got scalable and flexible cloud services for product development, testing, and deployment.

5. Anti-Virus Applications

:

Previously, organizations were installing antivirus software within their system even if we will see we personally also keep antivirus software in our system for safety from outside cyber threats. But nowadays cloud computing provides cloud antivirus software which means the software is stored in the cloud and monitors your system/organization's system remotely. This antivirus software identifies the security risks and fixes them. Sometimes also they give a feature to download the software.

6. E-commerce Application :

Cloud-based e-commerce allows responding quickly to the opportunities which are emerging. Users respond quickly to the market opportunities as well as the traditional e-commerce responds to the challenges quickly. Cloud-based e-commerce gives a new approach to doing business with the minimum amount as well as minimum time possible. Customer data, product data, and other operational systems are managed in cloud environments.

7. Cloud computing in education

Cloud computing in the education sector brings an unbelievable change in learning by providing e-learning, online distance learning platforms, and student information portals to the students. It is a new trend in education that provides an attractive environment for learning, teaching, experimenting, etc to students, faculty members, and researchers. Everyone associated with the field can connect to the cloud of their organization and access data and information from there.

8. E-Governance Application

Cloud computing can provide its services to multiple activities conducted by the government. It can support the government to move from the traditional ways of management and service providers to an advanced way of everything by expanding the availability of the environment, making the environment more scalable and customized. It can help the government to reduce the unnecessary cost in managing, installing, and upgrading applications and doing all these with help of could computing and utilizing that money public service.

9. Cloud Computing in Medical Fields

In the medical field also nowadays cloud computing is used for storing and accessing the data as it allows to store data and access it through the internet without worrying about any physical setup. It facilitates easier access and distribution of information among the various medical professional and the individual patients. Similarly, with help of cloud computing offsite buildings and treatment facilities like labs, doctors making emergency house calls and ambulances information, etc can be easily accessed and updated remotely instead of having to wait until they can access a hospital computer.

10. Entertainment Applications :

Many people get entertainment from the internet, in that case, cloud computing is the perfect place for reaching to a varied consumer base. Therefore different types of entertainment industries reach near the target audience by adopting a multi-cloud strategy. Cloud-based entertainment provides various entertainment applications such as online music/video, online games and video conferencing, streaming services, etc and it can reach any device be it TV, mobile, set-top box, or any other form. It is a new form of entertainment called On-Demand Entertainment (ODE).

With respect to this as a cloud, the market is growing rapidly and it is providing various services day by day. So other application of cloud computing includes social applications, management application, business applications, art application, and many more. So in the future cloud computing is going to touch many more sectors by providing more applications and services.

UNIT-2
CLOUD COMPUTING ARCHITECTURE

2.1. Introduction
2.2. Cloud Reference Model
2.3. Types of Clouds

2.4. Cloud Interoperability and standards
2.5. Cloud computing Interoperability use cases
2.6. Role of standards in Cloud Computing environment

Introduction

WHAT IS CLOUD COMPUTING TECHNOLOGY ARCHITECTURE?

Cloud Architecture refers to the various components in terms of databases, software capabilities, applications, etc. engineered to leverage the power of cloud resources to solve business problems. Cloud architecture defines the components as well as the relationships between them.

The various components of Cloud Architecture are:

- On premise resources
- Cloud resources
- Software components and services
- Middleware

The entire cloud architecture is aimed at providing the users with high bandwidth, allowing users to have uninterrupted access to data and applications, on-demand agile network with possibility to move quickly and efficiently between servers or even between clouds and most importantly network security

The various cloud based services have their own distinct and unique cloud architectures:

- Software as a Service (SaaS) involves software hosted and maintained on internet. With SaaS, users do not have to install the software locally.
- Development as a Service (DaaS) involves web based development tools shared across communities.
- Platform as a Service (PaaS) provides users with application platforms and databases, equivalent to middleware services.
- Infrastructure as a Service (IaaS) provides for infrastructure and hardware such as servers, networks, storage devices, etc. running in the cloud, available to users against a pay per usage basis.

As we know, cloud computing technology is used by both small and large organizations to **store the information** in cloud and **access** it from anywhere at anytime using the internet connection.

Cloud computing architecture is a combination of **service-oriented architecture** and **event- driven architecture**.

Cloud computing architecture is divided into the following two parts -

- Front End
- Back End

The below diagram shows the architecture of cloud computing -

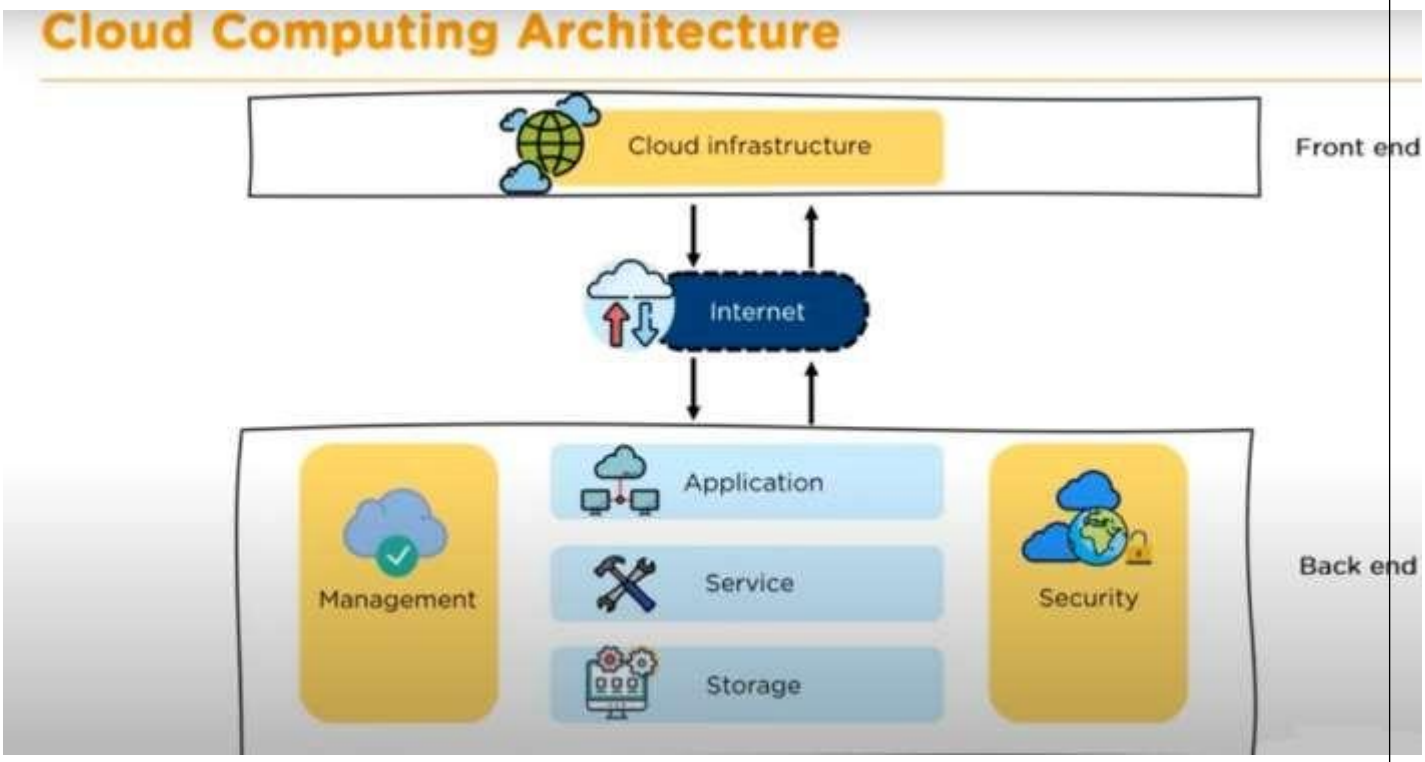

Front End

The front end is used by the client. It contains client-side interfaces and applications that are required to access the cloud computing platforms. The front end includes web servers (including Chrome, Firefox, internet explorer, etc.), thin & fat clients, tablets, and mobile devices.

Back End

The back end is used by the service provider. It manages all the resources that are required to provide cloud computing services. It includes a huge amount of data storage, security mechanism, virtual machines, deploying models, servers, traffic control mechanisms, etc.

Note: Both front end and back end are connected to others through a network, generally using the internet connection.

Components of Cloud Computing Architecture

There are the following components of cloud computing architecture -

1. Client Infrastructure

Client Infrastructure is a Front end component. It provides GUI (Graphical User Interface) to interact with the cloud.

2. Application

The application may be any software or platform that a client wants to access.

3. Service

A Cloud Services manages that which type of service you access according to the client's requirement.

Cloud computing offers the following three type of services:

i. Software as a Service (SaaS) – It is also known as **cloud application services**. Mostly, SaaS applications run directly through the web browser means we do not require to download andinstall these applications. Some important example of SaaS is given below –

Example: Google Apps, Salesforce Dropbox, Slack, Hubspot, Cisco WebEx.

ii. Platform as a Service (PaaS) – It is also known as **cloud platform services**. It is quite similar to SaaS, but the difference is that PaaS provides a platform for software creation, but using SaaS, we can access software over the internet without the need of any platform.

Example: Windows Azure, Force.com, Magento Commerce Cloud, OpenShift.

iii. Infrastructure as a Service (IaaS) – It is also known as **cloud infrastructure services**. It is responsible for managing applications data, middleware, and runtime environments.

Example: Amazon Web Services (AWS) EC2, Google Compute Engine (GCE), Cisco Metapod.

4. Runtime Cloud

Runtime Cloud provides the **execution and runtime environment** to the virtual machines.

5. Storage

Storage is one of the most important components of cloud computing. It provides a huge amount of storage capacity in the cloud to store and manage data.

6. Infrastructure

It provides services on the **host level**, **application level**, and **network level**. Cloud infrastructure includes hardware and software components such as servers, storage, network devices, virtualization software, and other storage resources that are needed to support the cloud computing model.

7. Management

Management is used to manage components such as application, service, runtime cloud, storage, infrastructure, and other security issues in the backend and establish coordination between them.

8. Security

Security is an in-built back end component of cloud computing. It implements a security mechanism in the back end.

9. Internet

The Internet is medium through which front end and back end can interact and communicate with each other.

https://youtu.be/X43KVeWVkSY

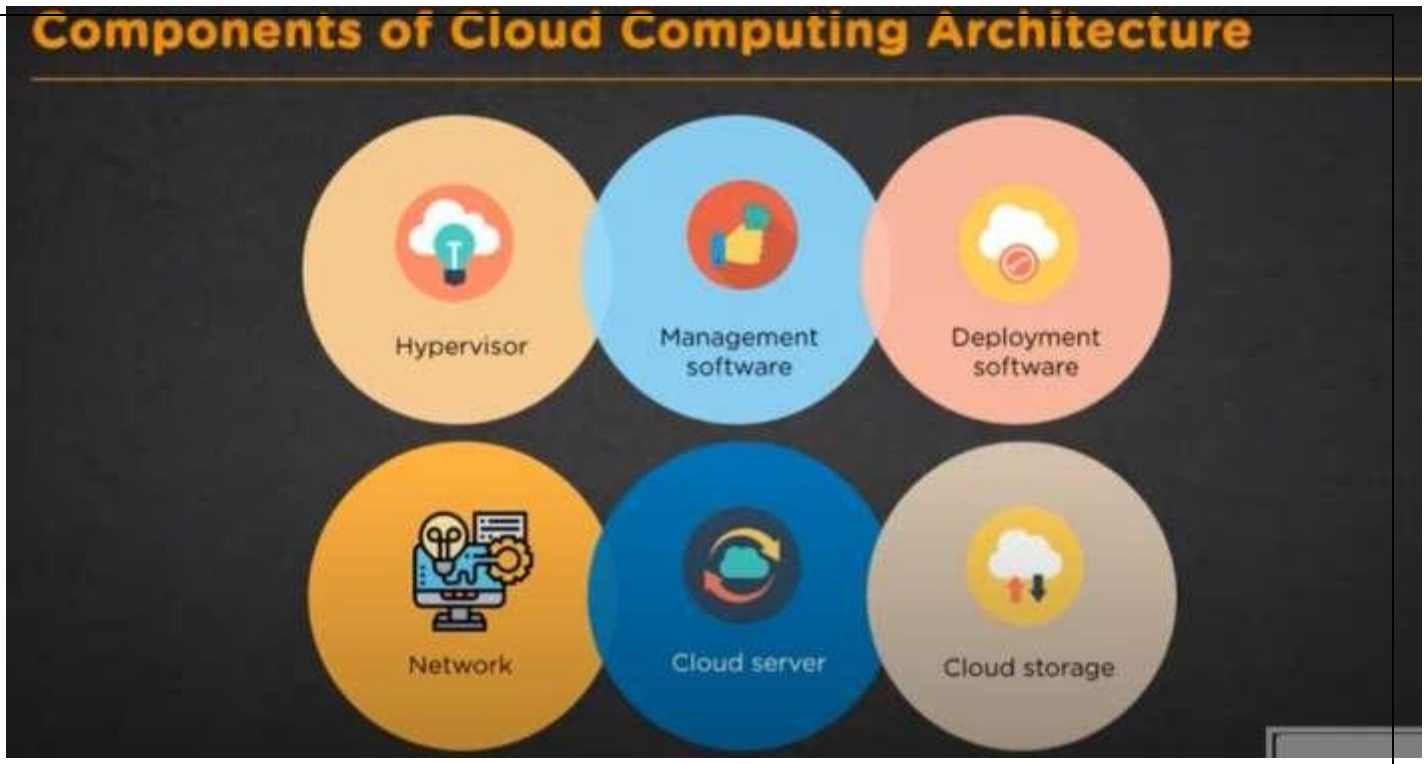

TYPES OF CLOUD

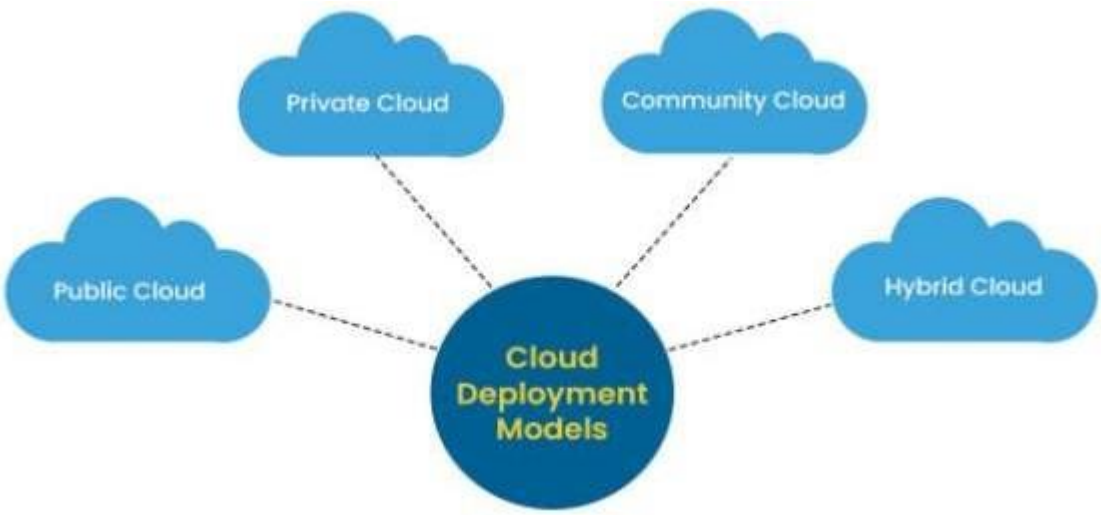

There are 4 types of Cloud deployment models:

- Public Cloud
- Private Cloud
- Hybrid Cloud
- Community Cloud

These deployment models differ on the basis of implementation type, hosting type and who has access to it. All Cloud deployment models are based on the same principle of Virtualization (abstraction of resources from bare metal hardware) but differ in terms of location, storage capacity, accessibility, and more. Depending on the type of data you are working with, you will wantto compare Public, Private, Hybrid, and Community Clouds in terms of different levels of security they offer and the management required.

- **Public Cloud**

  The entire computing infrastructure is located on the premises of the CSP that offers Cloud services over the internet. This is the most economical option for those individuals/organizations that do not wish to invest in IT infrastructure. In a Public Cloud environment, the resources are shared between multiple users who are also called 'Tenants' The cost of using Cloud services is determined through the usage of IT resources consumed.

- **Private Cloud**

  Individuals/organizations that choose Private Cloud gets dedicated infrastructure that is not shared by any other individual/organization. The security and control level is highest while using a private network. The costs are born by an individual/organization and are not shared with any other individual/organization. Management of Private Cloud is taken care of by the user and the CSP does not provide any Cloud management services.

- **Hybrid Cloud**

  This Cloud deployment model includes the characteristics of Public Cloud and Private Cloud. Hybrid Cloud allows the sharing of data and applications between Public and Private Cloud environments. Organizations mainly use Hybrid Cloud when their On-Premise infrastructure needs more scalability, so they make use of scalability on Public Cloud to meet fluctuating business demands. Organizations can keep their sensitive data on their Private Cloud when reaping the power of the Public Cloud.

- **Community Cloud**

  A Community Cloud is a Cloud infrastructure that is shared by users of the same industry or by those who have common goals. This Cloud infrastructure is built after understanding the computingneeds of a community as there are many factors including compliances and security policies which need to be included in the community Cloud infrastructure.

CLOUD INTEROPERABILITY AND STANDARDS

Basic Definition of Interoperability can be defined as a measure of the degree to which diverse systems or components can work together successfully.

More formally, IEEE and ISO define interoperability as the ability for two or more systems or applications to exchange information and mutually use the information that has been exchanged. In the context of cloud computing, interoperability should be viewed as the capability of public cloud services, private cloud services, and other diverse systems within the enterprise to understand each other's application and service interfaces, configuration, forms of authentication and authorization, data formats, etc. in order to work with each other.

Defining interoperability and portability

As with most IT concepts, different perspectives lead to different definitions, especially when consolidation around a particular defination may offer competitive edge. The following definitions, based on ISO/IEC19941, are the product of long debates among industry experts.

Interoperability is the ability of two or more systems or applications to exchange information and to mutually use the information that has been exchanged. ***Cloud interoperability*** is the ability of a customer's system to interact with a cloud service or the ability for one cloud service to interact with other cloud services by exchanging information according to a prescribed method to obtain predictable results.

The two noteworthy dimensions of interoperability – connectivity and usability – have been divided into five layers, as is illustrated in the diagram below.

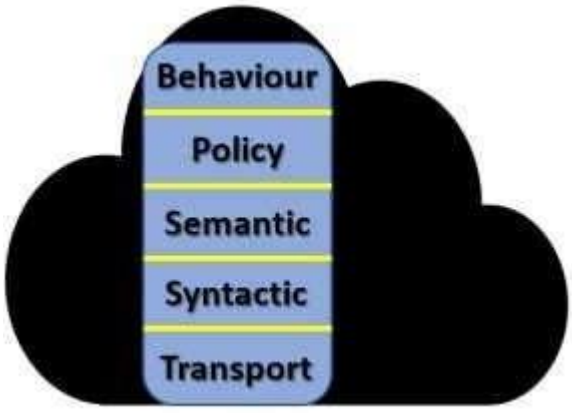

Transport interoperability is the exchange of data using physical networks such as the Internet; syntactic interoperability is concerned with the structure and coding of the data (e.g., English coded into ASCII characters); and semantic interoperability refers to the intended meaning of the data (i.e., the English word "customer" has different implications depending on the context). Interoperability also includes expected behaviour
– did the service behave as expected? Finally, policy interoperability is the ability of clouds to interoperate while complying with any applicable legal, organizational and policy frameworks.

Portability, on the other hand, is moving the data and/or applications from one system to another and having it remain useable or executable. **Cloud data portability** is the ability to easily move data from one cloud service to another without needing to re-enter the data. **Cloud application portability** is the ability to migrate an application from one cloud service to another or between a customer's environment and a cloud service.

The concepts of interoperability and portability can also be treated in terms of relative achievement – it's not an all or nothing proposition. The ideal for interoperability is that all five layers are fully compatible without custom translations, conversions or gateways, but ideal is rarely the case. With portability, the measure would be based on how few changes, especially manual changes, are required during the migration process.

Basic scenarios

The Cloud Standards Customer Council (CSCC) guide to cloud portability and interoperability has identified five major scenarios requiring interoperability and portability:

- Switching cloud service providers – the customer wants to move an application and data from Cloud 1 to Cloud 2;
- Use of multiple cloud service providers – the customer subscribes to the same or different services from two or more clouds (Clouds 1 and 2 in the diagram);
- Directly linked cloud services – the customer needs Cloud 1 to be linked to Cloud 3 to make use of its services;
- **Hybrid cloud configuration** – the customer connects legacy systems to an internal private cloud (e.g., Cloud 1) which is linked to a public cloud service (e.g., Cloud 3); and
- Cloud migration – the customer moves one or more in-house applications and/or data to Cloud 1.

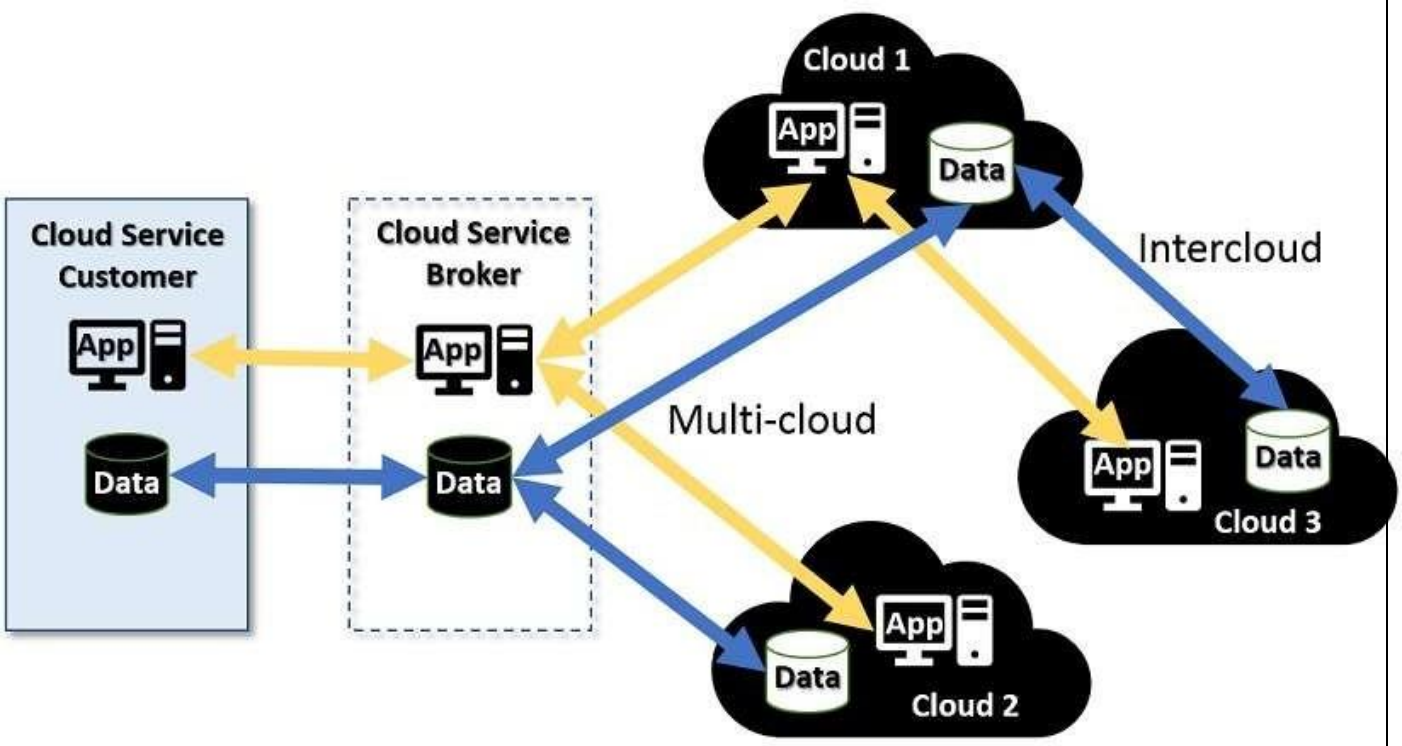

In summary, interoperability and/or portability may be required between legacy systems and cloud services (either public or private, in-house or hosted) and between two or more public and/or private cloud services. A cloud service broker may be a factor through its role as an intermediary between the customer and the cloud service provider(s).

Are interoperability and portability really needed?

The short answer is YES, except perhaps in the use of isolated, independent cloud services (as may be the case for Shadow IT). Any multi-cloud, multi-vendor cloud configuration depends on services that enable interoperability.

The interoperability challenges in each of the cloud service categories (IaaS, PaaS and SaaS) will be different. It is more difficult to achieve behavioural interoperability across two applications than it is to connect a compute service to remote storage, for example. Although no single, simple answer for either interoperability or portability is expected, a common language to discuss the issues and a shared understanding of the customer's requirements are always helpful.

Interoperability reduces technical complexity by eliminating custom gateways, converters and interfaces. One classic example of the interoperability problem can be seen with the transition from IPv4 and IPv6 in the Internet – a lack of interoperability has slowed down the adoption of IPv6.

Portability is often cited as being needed largely to mitigate vendor lock-in. However, moving from one system to another with a minimum of effort, as may be possible with container services, can also improve resilience and scalability.

Ultimately, more flexibility to change service providers, having fewer barriers to the leverage of "best of breed" solutions and greater ability to negotiate lower costs and better service quality are clear benefits for the cloud customer.

CLOUD-COMPUTING INTEROPERABILITY USE CASES

Use cases in the context of cloud computing refer to typical ways in which cloud consumers and providers interact. NIST, OMG, DMTF, and others—as part of their efforts related to standards for data portability, cloud interoperability, security, and management—have developed use cases for cloud computing.

NIST defines 21 use cases classified into three groups: cloud management, cloud interoperability, and cloud security [Badger 2010]. These use cases are listed below [Badger 2010]:

- **Cloud Management Use Cases**
- Open an Account
- Close an Account
- Terminate an Account
- Copy Data Objects into a Cloud
- Copy Data Objects out of a Cloud
- Erase Data Objects on a Cloud
- VM [virtual machine] Control: Allocate VM Instance
- VM Control: Manage Virtual Machine Instance State
- Query Cloud-Provider Capabilities and Capacities
- **Cloud Interoperability Use Cases**
- Copy Data Objects Between Cloud-Providers
- Dynamic Operation Dispatch to IaaS Clouds
- Cloud Burst from Data Center to Cloud
- Migrate a Queuing-Based Application
- Migrate (fully-stopped) VMs from One Cloud Provider to Another
- **Cloud Security Use Cases**
- Identity Management: User Account Provisioning
- Identity Management: User Authentication in the Cloud
- Identity Management: Data Access Authorization Policy Management in the Cloud
- Identity Management: User Credential Synchronization Between Enterprises and the
- **Cloud**
- eDiscovery
- Security Monitoring
- Sharing of Access to Data in a Cloud

CLOUD COMPUTING STANDARDS ORGANIZATIONS

Many cloud computing organizations and informal groups are focused on addressing standards issues in regards to the cloud environment. These standards bodies help maintain a standards and best practices to ensure that different providers and equipment are able to work together.

Cloud Security Alliance

The Cloud Security Alliance was formed to promote a series of best practices to provide security assurance in cloud computing. Its objectives include promoting understanding, researching best practices, launching awareness campaigns with the goal of creating a consensus on ways to ensure cloud security.

Distributed Management Task Force (DMTF)

The DMTF focuses on IaaS (Infrastructure as a Service), and providing standards that enable IaaS to be a flexible, scalable, high-performance infrastructure.

The DMTF is the group that developed the OVF standard that is formally known as DSP0243 Open Virtualization Format (OVF) V1.0.0. It describes an open, secure, and portable format for packaging and distribution of software that will be run in virtual machines.

National Institute of Standards and Technology (NIST)

NIST is a nonregulatory federal agency whose goal is to promote innovation and United States competitiveness by advancing standards, measurement science, and technology. They are focused on helping federal agencies understand cloud computing.

Open Cloud Consortium (OCC)

The OCC goal is to support the development of standards for cloud computing and frameworks for interoperating between clouds. The OCC has a number of different working groups devoted to varying aspects of cloud computing.

Open Grid Forum (OGF)

The OGF is an open community that focuses on driving the adoption and evolution of distributed computing. This includes everything from distributed high-performance computing resources to horizontally scaled transactional systems supporting SOA as well as the cloud.

The Object Management Group (OMG)

The OMG is an international group focused on developing enterprise integration standards for a wide range of industries including government, life sciences, and healthcare. The group provides modeling standards for software and other processes.

Storage Networking Industry Association (SNIA)

The SNIA is focused on developing storage solution specifications and technologies, global standards, and storage education. This organization's mission is "to promote acceptance, deployment, and confidence in storage-related architectures, systems, services, and technologies, across IT and business communities".

Cloud Computing Interoperability Forum (CCIF)

The Cloud Computing Interoperability Forum provides discussion forums to create a cloud computing ecosystem where organizations can work together. A major focus is on creating a framework that enables two or more cloud platforms to exchange information in a unified way.

Vertical industry groups

In addition to these standards groups and discussion groups, vertical industry groups are also beginning to look at cloud standards.

Examples include

- **Telemanagement Forum (TM Forum):** This large group consists of service providers, cable and network operators, software suppliers, equipment suppliers, and systems integrators. It recently began working in the telecommunications initiative for cloud computing.
- **Association for Retail Technology Standards (ARTS):** The goal of this group is to create an open environment where retailers and technology vendors can work together to create international retail technology standards. Recently, this group also started looking at researching this space and developing white papers to address cloud issues for this vertical.

ROLE OF STANDARDS IN CLOUD-COMPUTING ENVIRONMENTS

Cloud users would particularly welcome standards that address the workload migration and data migration use cases because such standards would mitigate vendor lock-in concerns. This requires standardization of virtual-machine image file formats and APIs for cloud
. Standardization for the user-authentication use case has the advantage that user identities based on OpenID or authentication protocols based on OAuth, for example, could be used across multiple providers that support these standards. Similarly, standardization to support the workload-management use case would leverage any existing efforts related to the construction of

workload-management clients and scripts that could be used across multiple providers. However, cloud providers use different types of service models, and some service models stand to benefit more from standardization than others. The remainder of this section looks at how IaaS, PaaS, and SaaS would benefit from standardization.

5.1 20BInfrastructure as a Service (IaaS)

IaaS is the service model that would most benefit from standardization because the main building blocks of IaaS are workloads represented as virtual-machine images and storage units that vary from typed data to raw data [Badger 2011].

For workload migration, standards efforts such as OVF and VHD would allow users to extract an image from one provider and upload it to another provider. Given that most IaaS providers allow consumers to install and run any operating system, a more manual and time-consuming form of migration would be to retrieve the image from the current provider, create a new image on a new provider, and reinstall software [Badger 2011]. This manual migration would not require standards as long as there is a way to retrieve the application state (e.g., application data, files, running processes) from the source image and move it to a new image.

For data migration, standards efforts such as CDMI and the Amazon S3 API, which multiple providers support, would enable users to extract data from one provider and upload it to a different provider. If a provider implements these standard interfaces using SOAP- or REST-based protocols, the cloud will offer the advantages of ease of development and tool availability. However, these standards are more useful for raw data that is not typed (e.g., virtual-machine images, files, blobs) because the cloud resource in this case simply acts as a container and usually does not require data transformation. For typed data, data migration would occur similarly to any other datamigration task: users must extract data from its original source, transform it to a format compatible with the target source, and upload it into the target source, which could be a complex process [Fogarty 2011]. The effort required for transformation will also depend on factors such as the similarity between the target's and source's data-storage technologies (e.g., moving from one SQLcompatible database to another will be easier than moving from an object database to a relational database or vice versa) and the similarity of the interface operations (e.g., two SOAP-based interfaces can have completely different operations).

5.2 21BPlatform as a Service (PaaS)

The PaaS service model benefits less from standardization than IaaS. Organizations that buy into PaaS do it for the perceived advantages of the development platform. The platform provides many capabilities out of the box, such as managed application environments, user authentication, data storage, reliable messaging, and other functionality in the form of libraries that can be integrated into applications. This functionality is tied to a specific language and runtime environment. For example, Google App Engine supports applications written in Java, Python, and Go. Microsoft Azure supports applications written in .NET, and more recently applications written in Java, PHP, and Node.js.

The incentives for PaaS adoption are primarily rapid development and deployment and the potential for these applications to serve a greater number of clients. Buying into a PaaS provider means buying into a platform in the same way that organizations traditionally have, which is based on added value, skills, cost, and any other criteria.

Providers can make applications more interoperable by selecting platforms that support more standardized tools and languages, such as those based on the Java language or standard dataaccess interfaces, including Java Database Connectivity (JDBC), Open Database Connectivity (ODBC), and SQL. However, even among providers that support the same programming language, the interfaces to basic services such as authentication, files, queues, hash tables, and tasks may not be compatible [Badger 2011]. In addition, native options may be more powerful (i.e., have greater benefit that can motivate an adoption decision) than standardized options. For example, the default data store in Google App Engine is the High Replication data store that offers automatic replication of data across data centers. A user can access the data store with a standard API or a low-level API. The tradeoff is that the standard API makes an application more portable but offers less control and less provider-specific value-added features than the low-level API, resulting in a lowest common denominator for features [Badger 2011].

5.3 2BSoftware as a Service (SaaS)

SaaS is a somewhat different model than IaaS and PaaS because it is a licensing agreement to third-party software instead of a different deployment model for existing resources that range

from data storage to applications.
Benefits of standardization for SaaS are even more limited than for PaaS. For SaaS offerings such as Salesforce.com CRM, the user is an end user. However, there are other SaaS offerings such as Google Maps or Yahoo Social in which the user can be a developer who is integrating functionality from these services into other applications [Google 2012c, Yahoo! 2012]. In the latter case, standardized APIs are useful because they facilitate the development process [Linthicum 2010a]. However, unless the APIs are identical from a functional perspective, this standardization helps little with migration.
Migration for the case when the SaaS user is an end user would occur in the same way as with any software migration because each SaaS provider has its own processing logic; it is simply a different way to license software [Harding 2010]. In this case, the only area where SaaS would benefit from standardization is data storage because the most important concern for SaaS consumers, especially
for enterprise software SaaS such as CRM or human resources, is how to extract their data. In one widely publicized incident, an online storage service shut down and a SaaS provider lost access to 45% of its customer data [Armbrust 2009]. In this case, the consumer would have to extract its data from the SaaS provider, write logic to perform data transformations, and then upload data to a new SaaS provider. The standardized APIs could potentially make this task easier.

5.4 23B**Do Standards Make Sense Beyond IaaS?**

The answer to this question is that they probably do not. A decision to adopt IaaS extends an organization's
IT department mainly by adding resources (primarily computation and storage) that exist outside of the organization and for which there is a pay-per-use fee as opposed to acquisition, maintenance, and obsolescence costs. In this case, the advantage of standards is that an organization can move these basic resources if another provider offers better prices or the organization experiences problems with its provider. Also, there is not much differentiation among IaaS providers other than price and SLAs.
A decision to adopt a PaaS or SaaS provider goes beyond extending basic IT resources. The service model usually involves value-added features in the form of libraries and platforms in the case of PaaS and application software in the case of SaaS. An organization selects a PaaS or SaaS provider based on these value-added features, and the choice involves a commitment similar to the commitment to traditional development platforms, deployment platforms, and software packages. PaaS and SaaS providers' focus on offering precisely the best set of value-added features creates many differences among them. Expecting PaaS and SaaS providers to standardize feature sets is equivalent to asking ERP software vendors to standardize feature sets. This is not likely to happen because it is not in their best interest.

5.5 24B**Can Existing Standards Support Cloud Interoperability Instead of Portability, or Do Clouds Require New Standards?**

Interoperability refers to the ability of a collection of communicating entities to share specific information
and operate on it according to agreed-on operational semantics [Brownsword 2009]. As mentioned earlier, even though the community desires standards for cloud interoperability, the reality is that existing standards efforts are so far focusing mainly on portability, which is the ability to migrate workloads and data from one provider to another.
Cloud interoperability, based on Brownsword's definition, refers to the ability of resources on one cloud provider to communicate with resources on another cloud provider. With this definition in mind, I examine whether each of the three types of service models would benefit from existing standards that promote interoperability, such as those that support service-oriented systems, or whether they would require new standards specific to the type of service model a cloud provider uses.
There are two basic use cases (UCs) for IaaS that exercise this service model's potential for interoperability:
UC1: Workload W1 on Cloud C1 can communicate with Workload W2 on Cloud C2.
UC2: Workload W1 on Cloud C1 can access Data Store DS in Cloud C2.
To support UC1, the following conditions must be true:
1. Workload W2 is accessible over the network and has a known address, uniform resource identifier (URI), or other unique identifier.
2. Workload W1 is authorized to communicate with Workload W2.

3. Workload W2 exposes an interface that Workload W1 can use.
This is a common interoperability scenario between two systems that does not require standards built specifically for the cloud. Standards such as SOAP and REST as well as existing userauthentication
standards could support this scenario if the cloud meets the conditions listed above.
Once workloads are running in a cloud instance, they behave like any other server.
Similarly to supporting UC1, to support UC2 the following conditions must be true:
1. DS is accessible over the network and has a known address, URI, or other unique identifier.
2. Workload W1 is authorized to access DS.
3. DS exposes an interface that Workload W1 can use.
This use case does benefit from standards for cloud data access such as CDMI and the Amazon S3 API.
The basic use case that exercises the PaaS service model's potential for interoperability is similar to UC1 for IaaS: Application A1 deployed on Cloud C1 can communicate with Application A2 on Cloud C2. Also similarly to supporting UC1, to support this use case the following must be true:
1. Application A2 is accessible over the network and has a known address, URI, or other unique identifier.
2. Application A1 is authorized to interact with Application A2.
3. Application A2 exposes an interface that Application A1 can use.
This is also a common interoperability scenario that does not require standards built specifically for the cloud.
The basic use case that exercises the SaaS service model's potential for interoperability is the same as for PaaS, except that it refers to interoperability between SaaS products instead of between applications. Interoperability between PaaS-deployed applications and IaaS workloads/data stores and SaaS products could also be supported the same way, if the cloud meets the conditions listed above.
The bottom line is that existing standards such as those that support service-oriented systems can support real cloud interoperability. However, there are different levels of system interoperability, as shown in Figure 1. Technical interoperability is about exchanging data, semantic interoperability is about exchanging meaningful data, and organizational interoperability is about participating in multi-organizational business processes.
Standards such as SOAP and REST enable technical (or syntactic) interoperability but do not guarantee semantic or organizational interoperability. Systems or data deployed inside cloud providers
will have to rely on documentation or formal/informal agreements to provide meaning to the interaction, just as in any use case that required systems to interoperate.

UNIT-3
Scalability and Fault Tolerance

What is scalability?

Scalability refers to the idea of a system in which every application or piece of infrastructure can be expanded to handle increased load. For example, suppose your web application gets featured on a popular website like ProductHunt. Suddenly, thousands of visitors are using your app – can your infrastructure handle the traffic? Having a scalable web application ensures that it can scale up to handle the load and not crash. Crashing (or even just slow) pages leave your users unhappy and your app with a bad reputation. Systems have four general areas that scalability can apply to:

1. Disk I/O
2. Memory
3. Network I/O
4. CPU

When talking about scalability in cloud computing, you will often hear about two main ways of scaling – horizontal or vertical. Let's look deeper into these terms.

Vertical scaling

Vertical is often thought of as the "easier" of the two methods. When scaling a system vertically, you add more power to an existing instance. This can mean more memory (RAM), faster storage such as Solid State Drives (SSDs), or more powerful processors (CPUs).The reason this is thought to be the easier option is that hardware is often trivial to upgrade on cloud platforms like AWS, where

servers are already virtualized. There is also very little (if any) additional configuration you are required to do at the software level.

Horizontal scaling

Horizontal scaling is slightly more complex. When scaling your systems horizontally, you generally add more servers to spread the load across multiple machines.With this, however, comes added complexity to your system. You now have multiple servers that require the general administration tasks such as updates, security and monitoring but you must also now sync your application, data and backups across many instances.

So which is better?

Horizontal scaling is often considered a long term advantage, whereas vertical scaling is usually considered a short term advantage.The reason for this is that you can typically add as many servers as you need to your infrastructure, but at some point, hardware upgrades are just not plausible.

Performance

One of the primary reasons for scaling your system is to increase performance. This is only one aspect of performance though – scaling ties in with many other concepts such as elasticity and fault tolerance.

Response time

Performance of a system is measured by many different metrics – one of the main ones is response time. Interestingly, scaling your system may increase response times. If you move away from the type of system architecture that has all of the components (database, application code, caching) on one server to a type of system architecture that separates these components onto their own serversthen the response time will naturally increase as you now have network latency and other considerations. Let's look at two popular system architecture types below.

Monolith

A monolith system architecture is the idea of having many of your components in one place. When talking about an application then it may mean that you have all of your services coupled together

such as your data layer, caching layer, file layer and business logic. When talking about hardware and servers it can mean that you run all of your processes in one place such as your database, web server and file system.

Microservices

A microservices system architecture is the process of splitting up core services into their own ecosystems. A key part of your application may be an image processing service that can save, delete, cache and manipulate images. This service could be set up as its own infrastructure which means that it would be separated from the other application services. You'll often hear the term *separationof concerns* when referring to microservices. Although each core service having its own infrastructure can make scalability easier, it can still add a lot of complexity to your application. You'll now have to manage multiple servers but also change your application code to handle these changes.

Scalability and databases

Each application is different but the key is to identify key services that may be a bottleneck and the first ones to cripple under increased load pressure. One of the most common bottlenecks can be the database.The database is used to store data in an application. You may use a traditional relational database such as MySQL or a NoSQL database such as MongoDB. In simple terms, the database is used to write data (save it) and read it (view it). The database can often be one of the first components to fall down under high load pressure in an application environment.

Sharding

To shard a database for scalability is to split your data up into separate database servers. Instead of having all of your data on one database server you would split the data into "shards". This can help with performance in a few ways:

- The data requests are shared across multiple servers instead of a the same database server each time
- Less data on each shard reduces index sizes which can improve data seek time
- Less data on each shard means there are less rows of data, this can allow queries to run quicker since there is less data to traverse or calculate

Partitioning

Database partitioning is similar to database sharding, but not exactly the same. Database partitioning separates the data into distinct parts. Certain partitioning methods include:

- Splitting data by range (alphabetically or numerically)
- Row wise (horizontal partitioning)
- Column wise (vertical partitioning)

Application code database optimizations

You can also perform application-level database optimizations, such as:

- Using database indexes
- Table partitioning
- Caching database queries
- De-normalization
- Running large queries/batch queries offline

Scalability benefits

The main benefit of scalable architecture is performance and the ability to handle bursts of traffic or heavy loads with little or no notice. A scalable system can help keep your application or online business running during peak times and not end up losing you money or damaging your reputation. Having your system set up into services such as the microservices system architecture can make monitoring, feature updates, debugging and scaling easier.

Scalability caveats

Scalability does have its caveats; it is certainly not a silver bullet. Creating a fully scalable system and infrastructure can be a large task that requires planning, testing and more testing. If you already have an application in place, splitting up that system can be a tedious process that may require code changes, software updates and more monitoring.

Scalability on AWS

Amazon Web Services as a platform has scalability built in. They offer many services that can help set up your application scale up or down depending on the resource requirements. One AWSproduct, the Elastic Load Balancer scales automatically on demand with the traffic it receives for your application. It also integrates with the Auto Scaling on your back end services (such as EC2 instances) to offer a full end to end scaling layer to handle different levels of traffic.

Review

Companies that use the cloud don't stay the same forever. They're using the cloud to help them expand their business. Scalability is one of the core concepts that aspiring Solutions Architects need to understand in order to be as effective as possible.That's a wrap! In this post you've learned all about scalability, how it affects systems and applications, its benefits and caveats, optimizing your database for scalability and how scalability is used with Amazon Web Services.

Explicating Fault Tolerance in Cloud Computing

Fault tolerance in cloud computing is about designing a blueprint for continuing the ongoing work whenever a few parts are down or unavailable. This helps the enterprises to evaluate their infrastructure needs and requirements, and provide services when the associated devices are unavailable due to some cause. It doesn't mean that the alternate arrangement can provide 100% of the full service, but this concept keeps the system in running mode at a useable, and most importantly, at a reasonable level. This is important if the enterprises are to keep growing in a continuous mode and increase their productivity levels.

Main Concepts behind Fault Tolerance in Cloud Computing System

- **Replication:** The fault-tolerant system works on the concept of running several other replicates for each and every service. Thus, if one part of the system goes wrong, it has other instances that can be placed instead of it to keep it running. Take, for example, a database cluster that has 3 servers with the same information on each of them. All the actions like data insertion, updates, and deletion get written on each of them. The servers, which are redundant, would be in inactive mode unless and until any fault tolerance system doesn't demand the availability of them.
- **Redundancy:** When any system part fails or moves towards a downstate, then it is important to have backup type systems. For example, a website program that has MS SQL as its database may failin between due to some hardware fault. Then a new database has to be availed in the redundancy concept when the original is in offline mode. The server operates with the emergency database which comprises several redundant services within.

Techniques for Fault Tolerance in Cloud Computing

- All the services have to be given priority when designing a fault tolerance system. The database has to be given special preference because it powers several other units.

- After deciding the priorities, the enterprise has to work on the mock test. Take, for example, the enterprise has a forum website that enables users to log in and posts comments. When the authentication services fail due to some problem, the users will not be able to log in. Then, the forum becomes a read-only one and does not serve the purpose. But with the fault tolerant systems, remediation will be ensured and the user can search for information with minimal impact.

Major Attributes of Fault Tolerance in Cloud Computing

- **None Point Failure:** The concepts of redundancy and replication defines that fault tolerance can be had but with some minor impacts. If there isn't even a single point failure then the system is not a fault tolerant one.
- **Accept the Fault Isolation Concept:** The fault occurrence has to be handled separately from other systems. This helps the enterprise to isolate it from the existing system failure.

Existence of Fault Tolerance in Cloud Computing

- **System Failure:** This may be either software or hardware issue. The software failure results in a system crash or hanging situation that may be due to stack overflow or other reasons. Any improper maintenance of the physical hardware machines will result in hardware system failure.
- **Security Breach Occurrences:** There are several reasons why fault tolerance occurs due to security failures. The hacking of the server negatively impacts the server and results in a data breach. Other reasons for the necessity of fault tolerance in the form of security breaches include ransomware, phishing, virus attack, etc.

Cloud Scalability

In cloud computing, cloud scalability refers to the ability to increase or reduce IT resources as required to meet evolving demands. One of the hallmarks of the cloud and the key factor of its burgeoning popularity with companies is scalability.

Using existing cloud computing technology, data storage space, processing power and networking can all be escalated. Better still, scaling, usually with little or no interruption or downtime, can be achieved rapidly and easily. Third-party cloud providers now have the entire infrastructure in place; in the past, the process could take weeks or months to scale with on-site physical infrastructure and entail enormous costs.

How to achieve cloud scalability?

To set up a personalized, scalable cloud solution via a public cloud, private cloud, or hybrid cloud, businesses have several options.

In cloud computing, two specific forms of scalability exist vertical and horizontal scaling.

We can add or subtract power to an existing cloud server memory upgrade, storage, or computing power with vertical scaling, also known as "scaling up" or "scaling down". This generally indicates that scaling has an upper limit based on the scaling capability of the server or machine; scaling above that also includes downtime.

We can add more resources like servers to our system using horizontal scalability to spread the workload across computers, which in turn improves efficiency and storage space. For companies with high-availability services that need limited downtime, horizontal scaling is essential.

Cloud Fault Tolerance

In cloud computing, fault tolerance is conceptually the same as in private or hosted environments. In other words, it means the infrastructure's ability to continue to provide service/services to underlying applications even when one or more component fails. To continue to work through failure or repair, we do not need to configure certain facilities for our infrastructure to use.

Objectives of Fault Tolerance in Cloud Computing

The fault-tolerant system uses backup components that take the place of failed components automatically, ensuring no service loss. They include:

- **Hardware systems**
  Hardware systems can be backed up using identical or equivalent systems. For instance, using an identical server running in parallel, with all operations mirrored to the backup server, a server can be made fault-tolerant.
- **Software systems**
  Software systems can be backed up using software instances. For example, it is possible to continuously replicate a database with customer information on another machine and operations can be mechanically redirected to another database in case a primary database goes down.
- **Power sources**
  Power sources use alternative sources using fault-tolerant. In many instances, organizations have power generators that can be used in case the electricity fails. Similarly, using redundancy, any system or component that is a single point of failure can be made fault-tolerant.
- **Security Breach Occurrences**
  Owing to security failures, there are many explanations about why fault tolerance exists. The server's hacking adversely affects the server and results in a leak of data. Ransomware, phishing, virus attack, etc. are other explanations for the need for fault tolerance in the form of security violations.

Key principles behind Cloud Computing Device Fault Tolerance

- **Replication**
  For every operation, the fault-tolerant system operates on the principle of running many other replicates. Therefore, if one aspect of the device goes wrong, it has other instances that can be put to keep it going instead. For example, a database of clusters that has 3 servers with the same information on each of them. All the acts are written on each of them, such as adding data, upgrading, and deleting. The redundant servers will be in inactive mode unless and until the availability of them is requested by any fault tolerance scheme.
- **Redundancy**
  If any part of the system fails or moves to a downstate, then it is necessary to have backup systems. For example, due to some hardware faults, a website programmer that has MS SQL as its database can fail in between. In the redundancy principle, a server works with an emergency database comprising many backup resources.

Techniques for Fault Tolerance in Cloud Computing

When developing a fault tolerance scheme, all the facilities have to be given priority. Special priority needs to be given to the database since it drives many other units.

The enterprise has to work on the test after deciding the objectives. Take the company's forum website, for example, which allows users to log in and make comments. If any problem causes the authentication services to malfunction, users may not be able to log in. The platform then becomes a read-only one and does not fulfill the objective. But remediation can be assured with fault-tolerant systems, and the user will search for details with minimal effect.

CLOUD ECOSYSTEM

A cloud ecosystem is a complex system of interdependent components that all work together to enable cloud services. In nature, an ecosystem is composed of living and nonliving things that are connected and work together. In cloud computing, the ecosystem consists of hardware and software as well as cloud customers, cloud engineers, consultants, integrators and partners.

Werner Vogels, CTO at Amazon, first compared the cloud to an ecosystem in a keynote address at the Cloud Connect 2011 conference. At the time, enterprise cloud computing was usually thought ofin terms of three broad service areas -- infrastructure-as-a-service (IaaS), platform-as-a-service (PaaS) and software-as-a-service (SaaS). Vogels proposed that the cloud was really more complex and its description also needed to include the array of service providers that companies rely on to operate in the cloud.

How a cloud ecosystem works

The center of a cloud ecosystem is a public cloud provider. It might be an IaaS provider such as Amazon Web Services (AWS) or a SaaS vendor such as Salesforce. Radiating out from the center of the cloud are software companies that use the provider’s anchor platform, as well as consultants and companies that have formed strategic alliances with the anchor provider. There is no vendor lock-in because these companies overlap, making the ecosystem more complex. For example, AWSis the center of its own ecosystem, but it's also a part of the Salesforce ecosystem. Salesforce runs a number of its services on AWS's infrastructure, and Salesforce customers can gain access, through devices called connectors, to pieces of AWS, such as its Simple Storage Service (S3).

A robust ecosystem provides a cloud provider's customers with an easy way to find and purchase business applications and respond to changing business needs. When the apps are sold through a provider’s app store such as AWS Marketplace, Microsoft Azure Marketplace (for cloud software) or Microsoft AppSource (for business applications), the customer essentially has access to a catalog of different vendors' software and services that have already been vetted and reviewed for security, risk and cost.

The benefits of a cloud ecosystem

Companies can use a cloud ecosystem to build new business models. It becomes relatively easy for a medical device manufacturer, for example, to launch a heart-monitoring service on its cloud service provider's cloud infrastructure and then sell the service alongside its main business of manufacturing heart monitors for hospitals.

In a cloud ecosystem, it is also easier to aggregate data and analyze how each part of the system affects the other parts. For example, if an ecosystem consists of patient records, smart device logs and healthcare provider records, it becomes possible to analyze patterns across an entire patient population.

CLOUD BUSINESS PROCESS MANAGEMENT

1. Business process management (BPM) is a mature business discipline that has spawned a number of technologies to support it.
2. Today it is the agile who survive those organizations who are able to adapt to change, to innovate as well as continuously improve, and to continuously monitor and analyze the results of these adaptations.
3. In the current web enabled business environment, processes in many cases depend on the discovery and recognition of components that exist as web services.
4. The current trend is towards increased emphasis on mobility and collaboration as essential elements to support the agility and currency of business processes.
5. This means that BPM vendors are increasingly seeking to augment their BPM packages by incorporating greater Web 2.0 type functionality.
6. Cloud based BPM is one response to these new demands.
7. BPM governs organizations cross functional, customer focused end to end core business process.

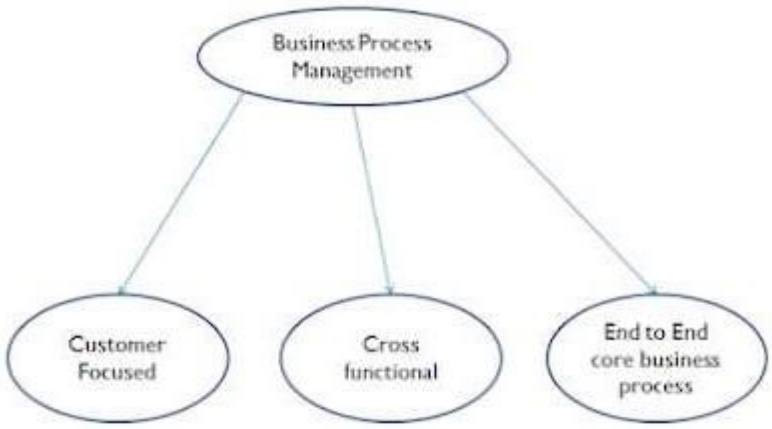

PORTABILITY AND INTEROPERABILITY

Cloud interoperability refers to the ability of customers to use the same management tools, server images and other software with a variety of cloud computing providers and platforms.

Standards are important in cloud computing for a variety of reasons. Standards for interoperability and data and application portability can ensure an open competitive market in cloud computing because customers are not locked-in to cloud providers and can easily transfer data or applications between cloud providers.

Why cloud interoperability and standards?

Vendor lock-in can prevent a customer from switching to another competitor's solution.
If switching is possible, it happens at considerable conversion cost and requires significant amounts of time.
Switching heppen because may be customer wants to find a more suitable solution for customer needs.
Or vendor may not be able to provide the service required.

So, the presence of standards that are actually implemented and adopted in the cloud computing community gives power for interoperability and then lessen the risks resulting from vendor lock-in.

CLOUD SERVICE MANAGEMENT

Service management:

1. A system integral of supply chain management that contents actual company sales and the customer.
2. The goal of service management is to maximize service supply chains.
3. The purpose of service management are to reduce high costs by integrating products and services and keep inventory levels smaller.

Cloud Service Management:

1. Cloud monitoring and cloud service management tools allow cloud providers to ensure optimal performance, continuity and efficiency in virtualized, on-demand environments.
2. The delivery of dynamic, cloud-based infrastructure, platform and application services doesn't occur in a vacuum.
3. In addition to best practices for effective administration of all the elements associated with cloud service delivery, cloud service management and cloud monitoring tools enable providers to keep up with the continually shifting capacity demands of a highly-elastic environment.

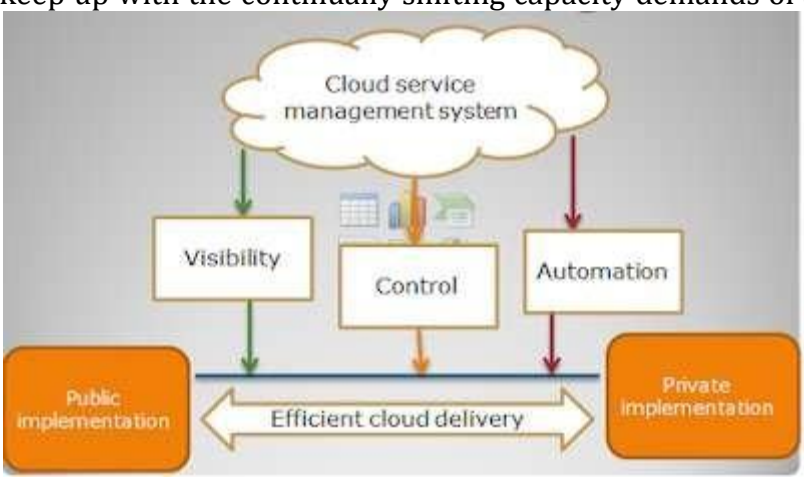

4. The Above fig illustrates that service management provides the visibility, control and automation needed for efficient cloud delivery in both public and private implementations.

Simplify user interaction with it:

1. The user friendly self-service accelerates time to value.
2. Service catalogue enables standards which drives consistent service delivery.

Enable policies to lower cost with provisioning:

1. Automatic allocating and de-allocating of resources will make delivery of services fast.
2. Provisioning policies allow release and reuse of assets.

Increase system admin productivity:

1. Providing the benefits to the broker will probably become a critical success factor in cloud computing.
2. Due to the growth of service brokerage business will increase the ability of cloud consumers to use services in a trustworthy manner.
3. These cloud mediators will help companies to choose the right platform, deploy the apps across multiple clouds.

Following are the opportunities for cloud brokers:

1. Cloud service intermediation : The broker must need to manage the additional securities or management capabilities over the cloud.
2. Cloud aggregation: It includes the deployment of services over multiple cloud platforms.
3. The ability to group an application across multiple clouds will become important i.e. if one service goes down the another can be started.

CLOUD OFFERINGS

Patterns of this category cover different functionality found in clouds regarding the functionality they provide to customers and the behavior they display.

Cloud Environments

Patterns of this category descibe the hosting environments of cloud in detail and refer to other offerings composed to form these environments.

- Elastic Infrastructure
- Elastic Platform
- Node-based Availability
- Environment-based Availability

Processing Offerings

Patterns of this category describe how computation can be performed in the cloud.

- Hypervisor
- Execution Environment
- Map Reduce

Storage Offerings

Patterns of this category describe how data can be stored in the cloud.

- Block Storage
- Blob Storage
- Relational Database
- Key-Value Storage
- Strict Consistency
- Eventual Consistency

Communication Offerings

Patterns of this category describe how data can be exchanged in the cloud.

- Virtual Networking

Message-oriented Middleware

Exactly-once Delivery

At-least-once Delivery

Transaction-based Delivery

Timeout-based Delivery

TESTING UNDER CONTROL

Cloud testing typically involves monitoring and reporting on real-world user traffic conditions as well as load balance and stress testing for a range of simulated usage conditions.

Load and performance testing conducted on the applications and services provided via cloud computing particularly the capability to access these services in order to ensure optimal performance and scalability under a wide variety of conditions.

Consumers can access the IT resources in the test environment.

Testing under the cloud gives very good sign by decreasing the manual intervention and reducing the processes in the typical testing environment.

After enabling of resources as and when they are required ,it reduces the investment on capital as well as enables the business to handle the ups and downs of the testing requirements.

The fig clearly shows that on the basis of these six parameters a cloud testing process can be performed.

Advantages of Cloud Testing:

1. Reduces capital investment and operational costs and not effect goal critical production application.
2. Offers new and attractive services to the clients and present an opportunity to speed cycles of innovations and improve the solution quality.

CLOUD SERVICE CONTROLS

VIRTUAL DESKTOP INFRASTRUCTURE

Virtual desktop infrastructure (VDI) is defined as the hosting of desktop environments on a central server. It is a form of desktop virtualization, as the specific desktop images run within virtual machines (VMs) and are delivered to end clients over a network.

Virtual Desktop Infrastructure (VDI) is a concept in which a solution based on a server based computing model that is not so different from the traditional terminal server centralized computing model used to deliver applications to remote users.

Virtual Desktop Infrastructure or VDI is the name given to a collection of technologies and processes that dramatically extend the concept of a remote desktop.

It centers on the idea that companies can virtualized their desktop operating systems like Windows XP or Vista using platforms like VMware ESX or Xenserver. and run said desktops from within the secured datacenter.

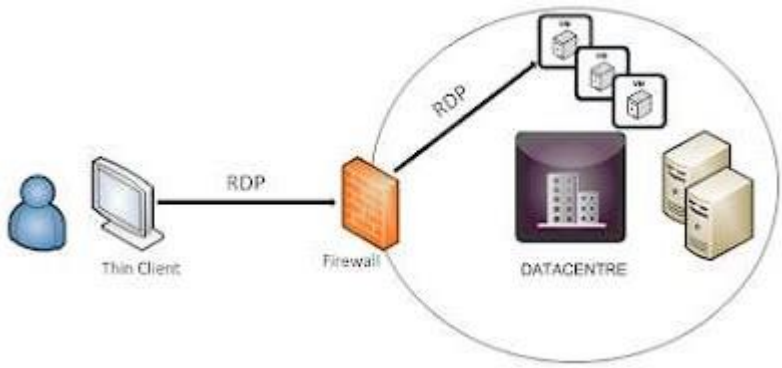

VDI Benefits:

1. Data Security
2. Reduced Hardware Expenditures
3. Easier Management: Perform centralized patching and application installation or streaming without loading the network
4. Mobile Workforce: User desktops are portable – users can reconnect from any location with a variety of devices.
5. Resource Pooling.

UNIT-4

Cloud Management and Virtualisation Technology

4.1. Create a virtualised Architecture
4.2. Data Centre
4.3. Resilience
4.4. Agility
4.5. Cisco Data Centre Network architecture
4.6. Storage
4.7. Provisioning
4.8. Asset Management
4.9. Concept of Map Reduce
4.10. Cloud Goverance
4.11. Load Balancing
4.12. High Availability
4.13. Disaster Recovery

Create a virtualised Architecture

Create and deploy end-to-end virtualizations that help;

- Reduce cost
- Provision new application quickly
- Maintain a high level of application performance

Whether you want to start with server virtualizations, extent virtualization across the data centre or implement virtual desktop infrastructure, network solution and cisco provides a comprehensive architecture approach that helps reduce costs, protect application performance and secure the virtual infrastructure.

Data Centre

A data center houses servers and/or data storage for an organization. This includes the hardware itself, the space in which it is housed, the power systems and backup systems, environmental controls, and anything else needed to keep those servers running. A data center can be a single server or complex with hundreds of servers on racks. Companies (like Amazon or Microsoft) that offer public cloud computing services have data centers that they then make available to other organizations. So why do people often contrast data centers with the cloud?

While cloud companies have their own data centers, organizations often have their own data centers as well, which are referred to as on-premises or on-prem for short. When most people talk about their data centers, the implication is that they are talking about on-premises data centers. So, on-prem data centers are owned and managed by the organization in question, for their own internal uses. With the cloud, data are stored and applications run off-premises and accessed remotely through the internet.

Resilience

Cloud **Resiliency:**

1. Cloud Resiliency is the capacity to rapidly adapt and respond to risks, as well as opportunities. In simple words resiliency refers to improve our business for handle risks.
2. This also maintains the continuous business operations that support growth.

3. The assessment process examines business-driven, data-driven, and event-driven risks. The goal is to understand the risks to the company and the business process in one building.

4. Risks in one geography are different from other locations. So we will be looking across different parts of the company, we have to find out common risks by focusing on one specific area first.

5. By using resilience framework to look at different parts of the company, we are trying to understand whether we have a risk that we can accept or whether we have risk that we want to avoid.

6. In other words either we may choose to do nothing about a risk, or we may improve our infrastructure handle the risks if they occur.

7. The resiliency blueprint includes different layers- facilities, technology, applications and data, processes, organization, strategy and vision.

8. The resiliency framework enables us to examine the business, understand what areas of vulnerability might come across business-driven, data-driven and event driven risks.

Resiliency capabilities: The strategy combines multiple parts to mitigate risks(that means to reduce the effect of risks) and improve business resilience.

1. From a facilities perspective, we may want to implement power protection.

2. From a security perspective, to protect our data and applications we may want to implement remote backup, identity management, email filtering, or email archiving.

3. From a process perspective, we may implement identification and documentation of most critical business processes.

4. From a organizational perspective, we may want to implement a virtual workstation environment.

5. From a strategy and vision perspective, we may want to look at the kind of crisis management process.

Resiliency tiers can be defined as a common set of infrastructure services that are delivered to meet or to provide a corresponding set of business availability expectations.

Agility

The highly competitive business world requires a lot of **business acumen and agility** on their part to survive. It requires agility not only in their strategic management process, but also their cloud services. The next question, is how does agility boosts your business?

Agility mean the ability of a business to adapt rapidly and cost efficiently in response to the changes in the business environment. Wikipedia defines Business Agility as "the ability of an organization to rapidly adapt to market and environmental changes in productive and cost-effectiveways." In that framework, cloud applications are fast in several ways.

Breaking down "Agility on Cloud"

Using the Cloud means agility and adding business value. In the cloud computing context, agility often refers to the ability to rapidly develop, test and launch software applications that drivebusiness growth. Cloud Agility allows them to focus on other issues such as security, monitoring and analysis, instead of provisioning and maintaining the resources.

You can achieve Agility in the cloud in a number of ways:

- **Quicker Time-to-market:**Cloud computing allows companies to significantly decrease the time it takes to provision and de-provision IT infrastructure, speeding delivery of IT projects that are critical to revenue growth or cost reduction. While a physical server could take days or weeks to procure and provision, a cloud server takes minutes. Faster time to market means faster time to revenue.

- **Automated allocation of resources**: Cloud computing simplifies provisioning, de-provisioning and re-deploying resources through automation and easy-to-use web consoles and APIs. The time for an IT systems administrator spent on managing and supporting cloud infrastructure is reduced greatly compared to that seen in a physical environment.

- **Flexibility and Scalability:**Cloud computing allows the flexibility for businesses to scale up or down their resources to meet the on-demands or sudden burst in demand or website traffic to meet unpredictable application development or production needs. The pay-per-use flexibility of the cloud, allows end-users to scale fast or "fail fast" based on the demands of the business. Common workloads that require on-demand scalability: testing and development, load testing, seasonal spikes in traffic, a new application etc.

- **Adaptive Auto-Scaling:**Cloud computing uses API's, software etc. for accessibility of cloud platforms and services. It is easier to automate IT management and provisioning in a cloud environment. You can integrate business intelligence and analytics platforms, IT monitoring tools with the cloud, allowing the systems to be more adaptive. Ex. new servers can be automatically provisioned (or de-provisioned) when load balancing thresholds are met.
- **Faster Innovation:**Cloud computing allows companies to support an increased pace of product development and marketing programs that better align IT infrastructure and management costs with the goals and objectives of the business.

Cisco Data Centre Network architecture

What Is the Cisco Data Center Network Architecture?
A comprehensive architecture that enables IT executives to:

- Consolidate and virtualize computing, storage and network resources
- Deliver secure and optimized employee, partner and customer access to information and applications
- Protect and rapidly recover IT resources and applications

Built with:

- **Networked Infrastructure:** Gigabit/10Gigabit, InfiniBand and storage switching and optical transport
- **Interactive Services:** Storage Fabric Services, computer services, security services, and application optimization services
- **Management:** Fabric manager (element and network management) and Cisco VFrame (server andservice provisioning)

Based on:
Cisco Service-Oriented Network Architecture (SONA), the enterprise implementation of the Intelligent Information Network (IIN) technology vision. Cisco SONA emphasizes the value of the interactive services provided in the networked infrastructure, such as application optimization, security, and server and storage fabric switching, to enhance business applications.
Cisco Data Center Network Architecture in Support of SONA

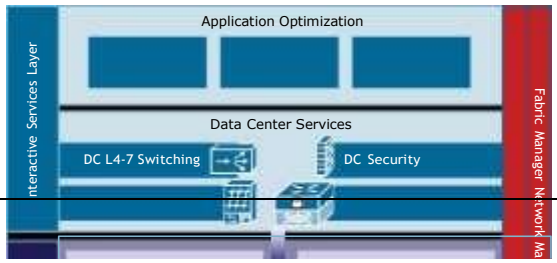

Benefits

- Lower-priced server and storage infrastructure
- Increased business agility and adaptability
- Ability to meet regulatory compliance standards with integrated network security and support for business continuance
- Tested and verified design and extensive service offerings for lower implementation costs and reduced risk
- Investment protection for core data center platforms offering multiyear deployment lifecycles
- Rapid application development and time to market of business-critical services

Why Cisco?

Cisco is the only vendor that delivers a complete architecture with advanced services, support, and industry-leading products. Cisco can help design the optimal end-state data center architecture and meet each tactical deployment phase of network evolution with the best products and services to achieve it.

What Is the Evolution of the Data Center?

- **Consolidation** of the front-end data network and back-end storage network infrastructures achieves greater administrative efficiency and increases utilization, increasing return on investment and lowering total cost of ownership.
- **Virtualization** increases productivity and business agility decoupling the application environment from the constraints of particular hardware. This way, computing, network, and storage resources can be allocated to an application in a way that best meets the needs of the organization.
- **Automation** manages the data center as a cohesive system by facilitating easier provisioning of resources while providing faster troubleshooting and easier recovery from security threats.

Cisco Data Center Network Architecture Overview and Products

Cisco Data Center Network Architecture can be grouped into four key areas:

1. **Server Fabric:** Used to interconnect servers to create high- performance cluster computing. Enterprises are starting to use low-latency interconnects to support parallel
and tightly coupled applications that provide finan- cial modeling, fluid dynamics and data mining.
By using InfiniBand and Remote Direct Memory Access (RDMA) technology, enterprises can reduce complex computing jobs to minutes and hours instead of days and weeks. **Cisco Products:**
Cisco SFS 3000 Series Multifabric Server Switches, Cisco SFS 7000 Series InfiniBand Server Fabric Switches. Cisco Infiniband Host Channel Adapters and blade switches integrated into IBM and
Dell blade servers.
2. **Storage Area/Fabric:** Used to consolidate and virtualize storage resources, so that they can be shared more effectively, virtual storage area networks (VSANs), and multiprotocol storage access through Fibre Channel, Small Computer System Interface over IP (iSCSI), and IBM Fiber Connection (FICON) enable large, heterogeneous storage networks. Support for advanced storage services like

virtualization, serverless backup, data replication, and continuous data protection allow for enhanced business continuance and data migration. **Cisco Products:** Cisco MDS 9000 Multilayer Directors and Fabric Switches

3. **Data Center Interconnect:** Connects the primary data center to a backup or secondary data center over optical or traditional WAN circuits. Data

replication and business continuance best practices mandate the need for high-speed, low-latency connections between data center locations. An optical network's inherent features—low latency, high bandwidth, and high density—are ideal for interconnecting storage area networks (SANs), cluster nodes, and server farms between multiple data centers. When optical networks are not feasible, data center protocols including Fibre Channel can
be transported over IP across traditional WANs. **Cisco Products:**

Cisco ONS 15302 Multiservice Customer Access Platform,
Cisco MDS 9000 and IPS Modules, Cisco Catalyst 6500 Switches

4. **Access Network:** Provides secure access to employees, customers, or partners connected remotely over the intranet, Internet, or extranet. The majority of users are not located close to the
data center, so robust, secure connectivity to the data center is mandatory. **Cisco Products:** Cisco Catalyst 6500 Series/7600 Series Supervisor Engine 720-3BXL, Cisco VPN 3000 Series Concentrators.

Business Resilience
The Data Center Network Architecture offers companies the ability to minimize theimpact of disaster scenarios through an architecture that helps mitigate risks and also provides tools and technologies that expedite recovery. The Data Center Network Architecture can also be a key part of an organizations' strategy for regulatory compliance and protection and management of company and customer data.

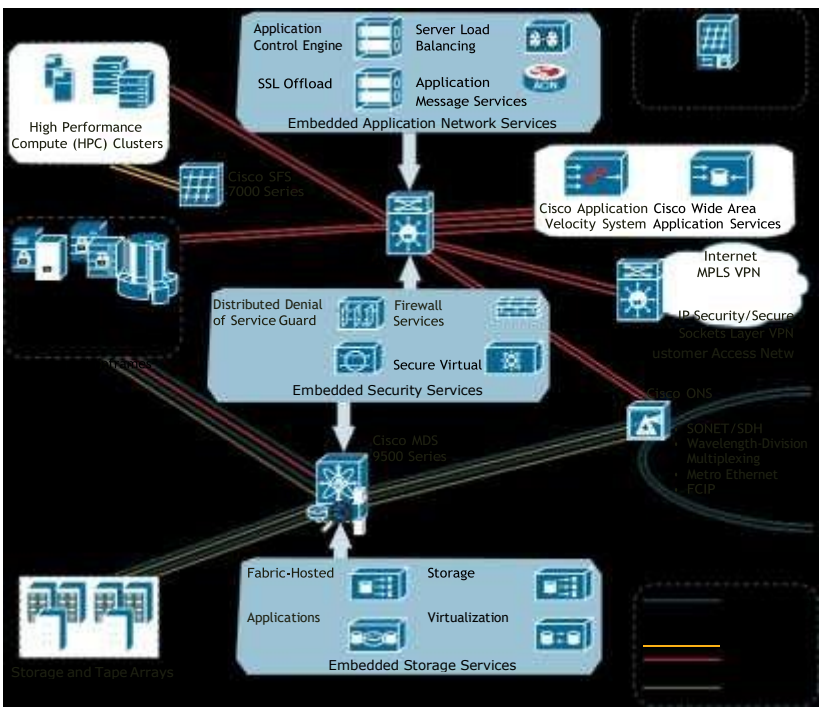

Components of the Data Center

#1 Optimization of Web Applications

Business Challenges	• Poor performance of HTTP-based Enterprise Applications. Examples: Portals, Siebel, SAP, Oracle, OWA, inotes
Decision Maker	Anyone Responsible forAn Application Service Level Agreement (SLA) • Director of Systems/Applications • Director of Operations/Networks
Business Benefits	• Ability to run delay sensitive applications over lower bandwidth links • Increased user satisfaction and adoption due to increase web application speed and performance
Cisco Solutions	Application Velocity System (AVS) 3100 (FineGround) • Dynamically caches/transforms/compresses content, secures web w/full proxy functionality • Reduces latency and improves responsiveness by handling all redirects

#2 Business Continuity and Disaster Recovery

Business Challenges	• Recovering business functions after disruptions and preventing data loss from any failure/attack
Decision Maker	• Storage Manager • Business Continuance Planning Manager • Risk/Compliance Manager
Business Benefits	• Ensure compliance with industry and other regulatory requirements • Overall improve business agility by

	creating a scalable and resilient solution • Improve customer and partner trust with a resilient design for applications and data
Cisco Solutions	• MDS 9500: Synchronous Mirroring and Asynchronous replication • Catalyst 6500: High-performance xWDM and 10GB Ethernet • ONS 15454/15540/15530: Supports high density, low-latency and high-bandwidth SAN extension solutions between Data Centers, providing native layer 2 extension for server clusters • Global Site Selector: Continuous Access with Automatic Site Selection
Note about Cisco on Cisco:	• Cisco synchronously replicates between data centers on Cisco San Jose campus, and asynchronously between San Jose and RTP sites for true fault-tolerant disaster recovery

#3 Consolidation of Branch File Servers

Business Challenges	• High operational cost of managing branch file and print servers
Decision Maker	• Director of Systems/Computing/Server Operations
Business Benefits	• Reduce maintenance delay of having to patch and maintain remote servers • Improve user experience with reduction of delay in accessing files over the WAN • Reduce cost by consolidating remote file and print servers
Cisco Solutions	• Cisco Wide Area File Services (WAFS) Core and Edge File Engines (formerly Actona and FineGround)

#4 Consolidate and Simplify Storage Management

Business Challenges	• Data backup on primary network exceeds backup window • High operational cost of managing underutilized and costly storage devices • Inability to dynamically scale servers and storage as business requirements change
Decision Maker	• Storage Manager • Data center Manager

Cisco Lifecycle Services for the Data Center Network Cisco Customer Advocacy (CA) Data Center Networking Services can bring together depth and breadth of expertise across the data center networking technologies to assist customers throughout the prepare, plan, design, implement, operate and optimize (PDIOO) network lifecycle. Cisco CA also advises customers on aligning their data center strategy with their business objectives and operational processes to industry standards and best practices.

Cisco services for data center networking complement those of our partners to form an end-to-end solution.

Programs to Help Implement Cisco Data Center Cisco Storage Networking CCIE Certification

CCIE certification in Storage Networking indicates expert level knowledge of intelligent storage solutions over extended network infrastructure using multiple transport options such as Fibre Channel, iSCSI, FCIP and FICON.

Storage

Provisioning

Cloud Provisioning:

The cloud provisioning is the allocation of a cloud provider's resources to a customer.

When a cloud provider accepts a request from a customer, it must create the appropriate number of virtual machines (VMs) and allocate resources to support them. The process is conducted in several different ways: advance provisioning, dynamic provisioning and user self-provisioning. The term provisioning simply means "to provide."

Cloud provisioning primarily defines how, what and when an organization will provision cloud services. These services can be internal, public or hybrid cloud products and solutions.

Cloud providers deliver cloud solutions through on-demand, pay-as-you-go systems as a service to customers and end users. Cloud provider customers access cloud resources through

Internet and programmatic access and are only billed for resources and services used according to a subscribed billing method.

Depending on the business model, a cloud provider may provide various solutions, such as:

1. IaaS
2. PaaS
3. SaaS

Types Of Provisioning:

1. Advanced Provisioning
2. Dynamic Provisioning
3. Self Provisioning

Asset Management

Cloud Asset management:

Cloud Asset management is a dedicated application which is used to record and track an asset throughout its life cycle, from occupying to disposal. It provides an organization with information like where certain assets are located, who is using them, how they are being utilized and details about the asset.

CAM is primarily about managing the challenges of cloud applications, platforms and infrastructure (SaaS, PaaS and IaaS). For instance:

1. Inability to track and manage the growing use of SaaS applications and providers
2. Lack of a centralized view of Cloud resources and consumption
3. Limited access to SaaS subscription data
4. Limited access to actual SaaS, IaaS and PaaS usage data

Benefits of Cloud Asset Management (CAM):

1. Accurate tracking of key applications delivered in the Cloud
2. Overcome the limitations of Cloud portals, by providing access to a single centralized view
3. Expanded access to data and improved analysis and reporting
4. Granular insight into SaaS, IaaS and PaaS usage across your organization
5. Combine Cloud and on-premise deployment data for a complete end-to-end view of your IT ecosystem
6. Accurate, complete view of investments and their usage across the whole IT estate enables better cost control
7. Access all the information needed to ensure a successful migration to the Cloud.

Concept of Map Reduce

MapReduce is a programming model or pattern within the Hadoop framework that is used to access big data stored in the Hadoop File System (HDFS). It is a core component, integral to the functioning of the Hadoop framework.

MapReduce facilitates concurrent processing by splitting petabytes of data into smaller chunks, and processing them in parallel on Hadoop commodity servers. In the end, it aggregates all the data from multiple servers to return a consolidated output back to the application.

For example, a Hadoop cluster with 20,000 inexpensive commodity servers and 256MB block of data in each, can process around 5TB of data at the same time. This reduces the processing time as compared to sequential processing of such a large data set.

With MapReduce, rather than sending data to where the application or logic resides, the logic is executed on the server where the data already resides, to expedite processing. Data access and storage is disk-based—the input is usually stored as files containing structured, semi-structured, or unstructured data, and the output is also stored in files.

MapReduce was once the only method through which the data stored in the HDFS could be retrieved, but that is no longer the case. Today, there are other query-based systems such as Hive and Pig that are used to retrieve data from the HDFS using SQL-like statements. However, these usually run along with jobs that are written using the MapReduce model. That's because MapReduce has unique advantages.

How MapReduce Works

At the crux of MapReduce are two functions: Map and Reduce. They are sequenced one after the other.

- The Map function takes input from the disk as <key,value> pairs, processes them, and produces another set of intermediate <key,value> pairs as output.
- The Reduce function also takes inputs as <key,value> pairs, and produces <key,value> pairs as output.

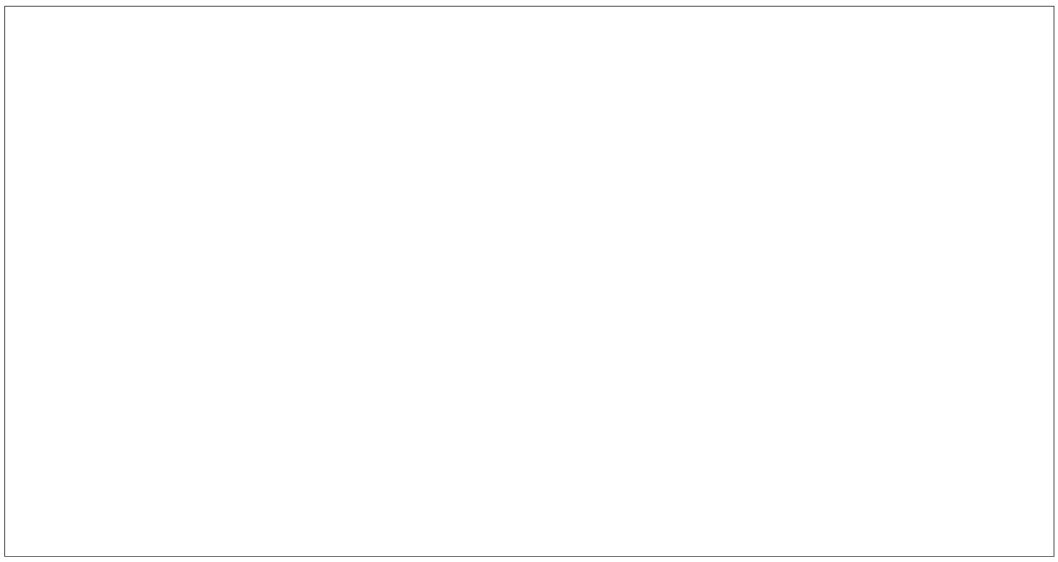

The types of keys and values differ based on the use case. All inputs and outputs are stored in the HDFS. While the map is a mandatory step to filter and sort the initial data, the reduce function is optional.

<k1, v1> -> Map() -> list(<k2, v2>)
<k2, list(v2)> -> Reduce() -> list(<k3, v3>)

Mappers and Reducers are the Hadoop servers that run the Map and Reduce functions respectively. It doesn't matter if these are the same or different servers.

Map

The input data is first split into smaller blocks. Each block is then assigned to a mapper for processing.

For example, if a file has 100 records to be processed, 100 mappers can run together to process one record each. Or maybe 50 mappers can run together to process two records each. The Hadoop framework decides how many mappers to use, based on the size of the data to be processed and the memory block available on each mapper server.

Reduce

After all the mappers complete processing, the framework shuffles and sorts the results before passing them on to the reducers. A reducer cannot start while a mapper is still in progress. All the map output values that have the same key are assigned to a single reducer, which then aggregates the values for that key.

Combine and Partition

There are two intermediate steps between Map and Reduce.

Combine is an optional process. The combiner is a reducer that runs individually on each mapper server. It reduces the data on each mapper further to a simplified form before passing it downstream.

This makes shuffling and sorting easier as there is less data to work with. Often, the combiner class is set to the reducer class itself, due to the cumulative and associative functions in the reduce function. However, if needed, the combiner can be a separate class as well.

Partition is the process that translates the <key, value> pairs resulting from mappers to another set of <key, value> pairs to feed into the reducer. It decides how the data has to be presented to the reducer and also assigns it to a particular reducer.

The default partitioner determines the hash value for the key, resulting from the mapper, and assigns a partition based on this hash value. There are as many partitions as there are reducers. So, once the partitioning is complete, the data from each partition is sent to a specific reducer.

A MapReduce Example

Consider an ecommerce system that receives a million requests every day to process payments. There may be several exceptions thrown during these requests such as "payment declined by a payment gateway," "out of inventory," and "invalid address." A developer wants to analyze last four days' logs to understand which exception is thrown how many times.

Example Use Case

The objective is to isolate use cases that are most prone to errors, and to take appropriate action. For example, if the same payment gateway is frequently throwing an exception, is it because of an unreliable service or a badly written interface? If the "out of inventory" exception is thrown often, does it mean the inventory calculation service has to be improved, or does the inventory stocks need to be increased for certain products?

The developer can ask relevant questions and determine the right course of action. To perform this analysis on logs that are bulky, with millions of records, MapReduce is an apt programming model. Multiple mappers can process these logs simultaneously: one mapper could process a day's log or a subset of it based on the log size and the memory block available for processing in the mapper server.

Map

For simplification, let's assume that the Hadoop framework runs just four mappers. Mapper 1, Mapper 2, Mapper 3, and Mapper 4.

The value input to the mapper is one record of the log file. The key could be a text string such as "file name + line number." The mapper, then, processes each record of the log file to produce key value pairs. Here, we will just use a filler for the value as '1.' The output from the mappers look like this:

Mapper 1 -> <Exception A, 1>, <Exception B, 1>, <Exception A, 1>, <Exception C, 1>, <Exception A, 1>
Mapper 2 -> <Exception B, 1>, <Exception B, 1>, <Exception A, 1>, <Exception A, 1> Mapper 3 -> <Exception A, 1>, <Exception C, 1>, <Exception A, 1>, <Exception B, 1>, <Exception A, 1>
Mapper 4 -> <Exception B, 1>, <Exception C, 1>, <Exception C, 1>, <Exception A, 1>

Assuming that there is a combiner running on each mapper—Combiner 1 … Combiner 4—that calculates the count of each exception (which is the same function as the reducer), the input to Combiner 1 will be:

<Exception A, 1>, <Exception B, 1>, <Exception A, 1>, <Exception C, 1>, <Exception A, 1>

Combine

The output of Combiner 1 will be:

<Exception A, 3>, <Exception B, 1>, <Exception C, 1>

The output from the other combiners will be:

Combiner 2: <Exception A, 2> <Exception B, 2>
Combiner 3: <Exception A, 3> <Exception B, 1> <Exception C, 1>
Combiner 4: <Exception A, 1> <Exception B, 1> <Exception C, 2>

Partition

After this, the partitioner allocates the data from the combiners to the reducers. The data is also sorted for the reducer.

The input to the reducers will be as below:

Reducer 1: <Exception A> {3,2,3,1}
Reducer 2: <Exception B> {1,2,1,1}
Reducer 3: <Exception C> {1,1,2}

If there were no combiners involved, the input to the reducers will be as below:

Reducer 1: <Exception A> {1,1,1,1,1,1,1,1,1}
Reducer 2: <Exception B> {1,1,1,1,1}
Reducer 3: <Exception C> {1,1,1,1}

Here, the example is a simple one, but when there are terabytes of data involved, the combiner process' improvement to the bandwidth is significant.

Reduce

Now, each reducer just calculates the total count of the exceptions as:

Reducer 1: <Exception A, 9>
Reducer 2: <Exception B, 5>
Reducer 3: <Exception C, 4>

The data shows that Exception A is thrown more often than others and requires more attention. When there are more than a few weeks' or months' of data to be processed together, the potential of the MapReduce program can be truly exploited.

Cloud Goverance

Cloud governance is a general term for applying specific policies or principles to the use of cloud computing services.

In other terms we can say that cloud governance refers to the decision making processes, criteria and policies involved in the planning, architecture, acquisition, deployment, operation and

management of a cloud computing capability.

The goal of cloud governance is to secure applications and data when they are located remotely.

There are five reasons of cloud governance:

1. Enable "business at cloud speed" and establish a cloud centric IT operating model based on the speed, agility and cost of cloud computing.
2. Enable appropriate cloud decision making without friction.
3. Integrated with existing Enterprise IT Governance processes, policies, boards and tools.
4. Balanced appropriate coverage for key decisions, Investments and risks while achieving the benefits of clouds.
5. Proactive to anticipate and prevent shadow clouds and unauthorized cloud activities that expose organizational risks.

We can define cloud governance as the framework to:

1. Convert rules, decisions and rights for the usage of IT resources into policies.
2. Ensure that cloud resource accessibility, provisioning, security, and operating procedures are executed in accordance with policies.
3. Provide automatic altering mechanism and remediation responses if policies are violated.
4. Provide capability to track policy changes and generate audit trails.

Effective governance tools are necessary to avoid careless or unauthorized use of cloud based IT resources, which includes the practice known as "shadow IT".

The governance is applied in cloud for:

1. Setting company policies in cloud computing.
2. Risk based decision which cloud provider, if any, to engage.
3. Assigning responsibilities for enforcing and monitoring of the policy compliance.
4. Set corrective action for non-compliance.

Cloud Governance model example: Microsoft's Cloud Governance Model.

Load Balancing

High Availability and Disaster Recovery

HIGH AVAILABILITY:

1. In simple words we can say that high availability refers to the availability of resources in a computer system.

2. In terms of cloud computing it refers to the availability of cloud services.

3. High availability is the heart of the cloud.

4. It provides the idea of anywhere, anytime access to service of cloud environment.

5. Availability is also related to reliability.

6. Availability is a technology issue as well as business issue.

7. High Availability can be simply defined by the simple equation:

$$HA = \frac{MTBF}{MTBF * MTTR}$$

Where ,
MTBF – mean time between failures,
MTTR- means time to repair and
HA- high availability.

8. There is two way improve the availability:-

1. Increase MTBF to very large values.
2. Reduce MTTR to very low values.

DISASTER RECOVERY:

1. Disaster recovery (DR) is the process, policies and procedures that are related to preparing for recovery or continuation of technology infrastructure which are an organization after a natural or human-induced disaster.

2. A disaster recovery is the process by which an organization can recover and access their software, data, and hardware.

3. It is necessary for faster disasters recovery to have an infrastructure supporting high availability.

4. The failure of disaster recovery plan mainly due to lack of high availability preparation, planning and maintenance to occurrence of the disaster.

Strategies of Disaster Recovery:

1. RTO (Recovery Time Objective): RTO is the period of time within which system, application, or functions must be discovered after an outage. RTOs are often used as the basis for the development of recovery strategies and as determinant as to whether or not to implement the recovery strategies during a disaster situation.
2. RPO (Recovery point objective): RPO is the point to time to which systems and data must be recovered after an outage. RPOs are often used as the basis for the development of backup strategies, and as a determinant of the amount of data that may need to be recreated after the systems or function have been recovered.

Unit-5
Virtualisation

5.1. Virtualisation
5.2. Betwork Virtualisation

5.3. Desktop and Application Virtualisation
5.4. Desktop as a service
5.5. Local desktop Virtualisation
5.6. Virtualisation benefits
5.7. Server Virtualisation
5.8. Block and File level Storage Virtualisation
5.9. Virtual Machine Monitor
5.10. Infrastructure Requirements
5.11. VLAN and VSAN

Virtualisation

Over an existing operating system & hardware, we generally create a virtual machine that and above it, we run other operating systems or applications. This is called Hardware Virtualization. The virtual machine provides a separate environment that is logically distinct from its underlying hardware. Here, the system or the machine is the host & the virtual machine is the guest machine. This virtual environment is managed by firmware, which is termed as a hypervisor.

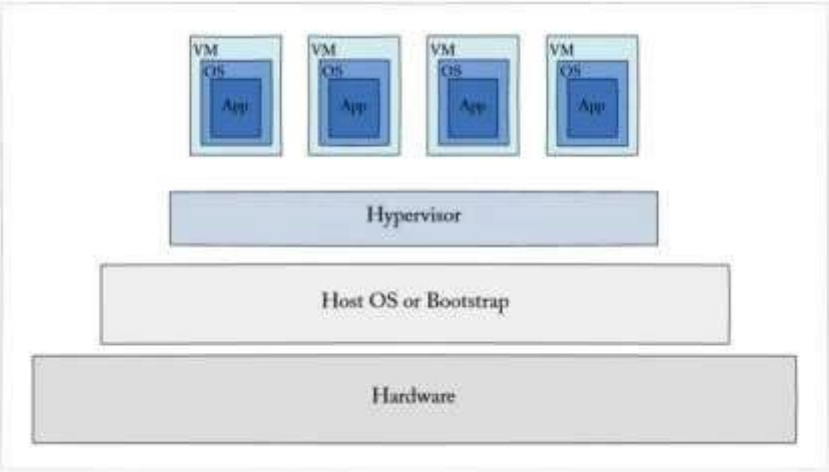

There are several approaches or ways to virtualizes cloud servers.

These are:

- **Grid Approach**: where the processing workloads are distributed among different physical servers, and their results are then collected as one.
- **OS - Level Virtualization**: Here, multiple instances of an application can run in an isolated form on a single OS
- **Hypervisor-based Virtualization**: which is currently the most widely used technique

The virtualization of the cloud has been categorized into four different types based on their characteristics. These are:

1. Hardware Virtualization
1. Full Virtualization
2. Emulation Virtualization
3. Para-virtualization
2. Software Virtualization

3. OS Virtualization
4. Server Virtualization
5. Storage Virtualization

Types of Virtualization:

1.Application Virtualization.

2.Network Virtualization.

3.Desktop Virtualization.

4.Storage Virtualization.

5.Server Virtualization.

6.Data virtualization.

1. Application Virtualization:

Application virtualization helps a user to have remote access of an application from a server. The server stores all personal information and other characteristics of the application but can still run on a local workstation through the internet. Example of this would be a user who needs to runtwo different versions of the same software. Technologies that use application virtualization are hosted applications and packaged applications.

2. Network Virtualization:

The ability to run multiple virtual networks with each has a separate control and data plan. It co-exists together on top of one physical network. It can be managed by individual parties that potentially confidential to each other.
Network virtualization provides a facility to create and provision virtual networks—logical switches, routers, firewalls, load balancer, Virtual Private Network (VPN), and workload security within days or even in weeks.

3. Desktop Virtualization:

Desktop virtualization allows the users' OS to be remotely stored on a server in the data centre. It allows the user to access their desktop virtually, from any location by a different machine. Users who want specific operating systems other than Windows Server will need to have a virtual desktop. Main benefits of desktop virtualization are user mobility, portability, easy management of software installation, updates, and patches.

4. Storage Virtualization:

Storage virtualization is an array of servers that are managed by a virtual storage system. The servers aren't aware of exactly where their data is stored, and instead function more like worker bees in a hive. It makes managing storage from multiple sources to be managed and utilized as a single repository. storage virtualization software maintains smooth operations, consistent performance and a continuous suite of advanced functions despite changes, break down and differences in the underlying equipment.

5. Server Virtualization:

This is a kind of virtualization in which masking of server resources takes place. Here, the central-server(physical server) is divided into multiple different virtual servers by changing the identity number, processors. So, each system can operate its own operating systems in isolate manner. Where each sub-server knows the identity of the central server. It causes an increase in the performance and reduces the operating cost by the deployment of main server resources into a sub-server resource. It's beneficial in virtual migration, reduce energy consumption, reduce infrastructural cost, etc.

6. Data virtualization:

This is the kind of virtualization in which the data is collected from various sources and managed that at a single place without knowing more about the technical information like how data is collected, stored & formatted then arranged that data logically so that its virtual view can be accessed by its interested people and stakeholders, and users through the various cloud services

remotely. Many big giant companies are providing their services like Oracle, IBM, At scale, Cdata, etc.

It can be used to performing various kind of tasks such as:

- Data-integration
- Business-integration
- Service-oriented architecture data-services
- Searching organizational data

How Virtualization Works in Cloud

Virtualization plays a significant role in cloud technology and its working mechanism. Usually, what happens in the cloud - the users not only share the data that are located in the cloud-like application but also share their infrastructures with the help of virtualization. Virtualization is used mainly to provide applications with standard versions for cloud customers. With the release of the latest version of an application, the providers can efficiently provide that application to the cloud and its users, and it is possible using virtualization only. By using this virtualization concept, all servers & software other cloud providers require those are maintained by a third-party, and the cloud provider pays them on a monthly or yearly basis.

In reality, most of today's hypervisors use a combination of different types of hardware virtualization. Mainly virtualization means running multiple systems on a single machine but sharing all resources (hardware) & it helps to share IT resources to get benefits in the business field.

Difference Between Virtualization and Cloud

1. Essentially there is a gap between these two terms, though cloud technology requires the concept of virtualization. Virtualization is a technology - it can also be treated as software that can manipulate hardware. At the same time, cloud computing is a service that is the result of manipulation.
2. Virtualization is the foundation element of cloud computing, whereas Cloud technology is the delivery of shared resources as a service-on-demand via the internet.
3. Cloud is essentially made-up of the concept of virtualization.

Advantages of Virtualization

- The number of servers gets reduced by the use of the virtualization concept.
- Improve the ability of technology.
- The business continuity was also raised due to the use of virtualization.
- It creates a mixed virtual environment.
- Increase efficiency for the development and test environment.
- Lowers Total Cost of Ownership (TCO).

Features of Virtualization

1. **Partitioning**: Multiple virtual servers can run on a physical server at the same time.

2. **Encapsulation of data**: All data on the virtual server, including boot disks, is encapsulated in a file format.
3. **Isolation**: The Virtual server running on the physical server is safely separated and don't affect each other.
4. **Hardware Independence**: When the virtual server runs, it can migrate to a different hardware platform.

NETWORK VIRTUALIZATION

What is network virtualization?

Network Virtualization (NV) refers to abstracting network resources that were traditionally delivered in hardware to software. NV can combine multiple physical networks to one virtual, software-based network, or it can divide one physical network into separate, independent virtual networks.

Network virtualization software allows network administrators to move virtual machines across different domains without reconfiguring the network. The software creates a network overlay that can run separate virtual network layers on top of the same physical network fabric.

Why network virtualization?

Network virtualization is rewriting the rules for the way services are delivered, from the software-defined data center (SDDC), to the cloud, to the edge. This approach moves networks from static, inflexible, and inefficient to dynamic, agile, and optimized. Modern networks must keep up with the demands for cloud-hosted, distributed apps, and the increasing threats of cybercriminals while delivering the speed and agility you need for faster time to market for your applications. With

network virtualization, you can forget about spending days or weeks provisioning the infrastructure to support a new application. Apps can be deployed or updated in minutes for rapid time to value.
How does network virtualization work?

Network virtualization decouples network services from the underlying hardware and allows virtual provisioning of an entire network. It makes it possible to programmatically create, provision, and manage networks all in software, while continuing to leverage the underlying physical network as the packet-forwarding backplane. Physical network resources, such as switching, routing, firewalling, load balancing, virtual private networks (VPNs), and more, are pooled, delivered in software, and require only Internet Protocol (IP) packet forwarding from the underlying physical network.

Network and security services in software are distributed to a virtual layer (hypervisors, in the data center) and "attached" to individual workloads, such as your virtual machines (VMs) or containers, in accordance with networking and security policies defined for each connected application. When a

workload is moved to another host, network services and security policies move with it. And when new workloads are created to scale an application, necessary policies are dynamically applied to these new workloads, providing greater policy consistency and network agility.
Benefits of network virtualization

Network virtualization helps organizations achieve major advances in speed, agility, and security by automating and simplifying many of the processes that go into running a data center network and managing networking and security in the cloud. Here are some of the key benefits of network virtualization:

- Reduce network provisioning time from weeks to minutes
- Achieve greater operational efficiency by automating manual processes
- Place and move workloads independently of physical topology
- Improve network security within the data center
Network Virtualization Example

One example of network virtualization is virtual LAN (VLAN). A VLAN is a subsection of a local area network (LAN) created with software that combines network devices into one group, regardless of physical location. VLANs can improve the speed and performance of busy networks and simplify changes or additions to the network.

Another example is network overlays. There are various overlay technologies. One industry-standard technology is called virtual extensible local area network (VXLAN). VXLAN provides a framework for overlaying virtualized layer 2 networks over layer 3 networks, defining both an encapsulation mechanism and a control plane. Another is generic network virtualization encapsulation (GENEVE), which takes the same concepts but makes them more extensible by being flexible to multiple control plane mechanisms.

VMware NSX Data Center – Network Virtualization Platform

VMware NSX Data Center is a network virtualization platform that delivers networking and security components like firewalling, switching, and routing that are defined and consumed in software. NSX takes an architectural approach built on scale-out network virtualization that delivers consistent, pervasive connectivity and security for apps and data wherever they reside, independent of underlying physical infrastructure.

Desktop and Application Virtualisation

Desktop Virtualization

To have a robust desktop environment management system, desktop virtualization can be used in combination with application virtualization and user profile management systems, called desktop virtualization. All the components of the desktop are virtualized in this mode, which enables a highly scalable and much more reliable model of desktop delivery. Besides, a more robust desktop recovery plan is enabled by this approach as all components are effectively stored in the data center and backed up by conventional redundant maintenance systems. If the device or hardwareof a user is missing, the repair is quick and simple, since the components from another device would be available at login. Moreover, since no data is stored on the computer of the user, if the device is lost, there is far less risk that any sensitive data can be recovered and compromised.

Advantages of Desktop Virtualization

- **Security**: Desktop virtualization software provides IT administrators with centralized security control over which users can access which data and which applications. If a user's permissions change because they leave the company, desktop virtualization makes it simple for IT to quickly revoke that user's access to their persistent virtual desktop and all its data instead of having to manually uninstall all of that user's computers. And a lost or stolen device does not face the same data danger because, rather than on each machine, all company data lives inside the data centre. If anyone uses desktop virtualization to steal a laptop, there is no real machine company details and thus less chance of a breach.
- **Resource Management**: Desktop virtualization lets IT departments get the most from their hardware investments by consolidating most of their computing in a data centre. Desktop virtualization then helps companies to issue lower-cost computers and tablets to end users where none of the intensive computing work in the data centre takes place. IT departments can save money by purchasing less costly machines by minimizing how much computation is needed for end-users on the endpoint computers.
- **Remote work**: Desktop virtualization enables IT, administrators, to support remote employees by giving central IT power over desktop virtual implementation across an organization's devices. Desktop virtualization allows IT, rather than manually setting up a new desktop for each user, to simply install a ready-to-go virtual desktop to that user's laptop. The user can now communicate with the operating system and applications on that desktop from every place, and the experience of the employee would be the same as if they were working locally. Once the user has finished using this virtual desktop, they can log off and return the desktop image to the shared pool.

Application Virtualization

Virtualization of applications is a technology that encapsulates an application on which it is executed from the underlying operating system. It allows an application to be accessed without needing to install it on the target device.

The application works and interacts with its device from the user's perspective. As well as using familiar keyboard and mouse operations, the user can resize, move, or minimize the application window. At times, there may be slight variations, but the consumer has a smooth experience for the most part.

Advantages of Application Virtualization

There are several advantages to application virtualization, some of which are described below:

- **Better Portability**: Virtualized apps can be used on any endpoint, whether Windows, Android, or IOS, anywhere and anywhere. Furthermore, this added portability enables sensitive data to remain on the server, and if an endpoint is compromised or stolen, there is no need to worry.
- **Simplified Support**: If there are problems with the operation of the virtual apps, then helpdesk workers can quickly see it from a central location and centrally address the problems.
- **Independence from the Operating System**: Virtualized apps are independent of the operating system used, making them available for every Microsoft, IOS, or Android endpoint.

- **Simple to get rid of applications**: To get rid of them, we can easily uninstall virtual applications. Uninstallation is therefore not necessary on each device.
- **Reduced conflicts between applications**: Installing one application may sometimes lead to problems and the other may crash. Because virtualized apps are virtual versions running on endpoints, this problem is largely reduced.
- **Simple Installation**: We can install an application once on the server and simply virtualized the application to as many endpoints as we want. The need to install the application on each endpoint is therefore reduced.
- **Easy deployment**: It's also easy to deploy applications for clients or partners. We can simply send them the already configured executable file, and it becomes easier to deploy these apps.
- **Easier Rollback**: If an application is not working as expected, it can be centrally reverted or rolled-back to its most stable state easily.
- **Improved Security**: Because virtualized applications, as well as the operating system, are isolated from each other, malware that appears in one cannot impact the other.
- **Easier updates**: The virtualized apps can be updated once, from a central location, and these updates do not have to be done on all desktops individually.

Differences between Desktop Virtualization & Application Virtualization

Desktop Virtualization	**Application Virtualization**
Provides greater virtual infrastructure versatility.	The smaller degree of versatility by contrast.
Richer and clear experience with desktops.	Desktop experience varies between software and applications.
It is difficult to manage applications because simple adjustments enable the "golden image" to be redeployed in all VDI instances.	Enables faster programmer maintenance, allowing improvements to take place without even knowing that changes have taken place by the user.
Depending on the use situation, the cost may be a concern.	Cost-effective cure.
In the underlying OS, applications are still related to.	Isolates the programmer entirely from the underlying OS.
Provides users with the experience of a full desktop.	Provides consumers with the individualized experience to render it application-specific.

5.4. Desktop as a service

Desktop as a Service (DaaS) is a cloud computing offering where a service provider delivers virtual desktops to end users over the Internet, licensed with a per-user subscription.

The provider takes care of backend management for small businesses that find creating theirown virtual desktop infrastructure to be too expensive or resource-consuming. This management typically includes maintenance, back-up, updates, and data storage. Cloud service providers may also handle security and applications for the desktop, or users may manage these service aspects individually. There are two kinds of desktops are available in DaaS—persistent and non-persistent.

1. **Persistent desktop**: Users have the ability to customize and save a desktop so it will look the same way each time a particular user logs on. Persistent desktops require more storage than non-persistent desktops, which can make them more expensive.
2. **Non-persistent desktop**: Desktops are wiped each time the user logs out—they are merely a way to access shared cloud services.

Cloud providers may allow customers to choose from both, allowing workers with specific needs to access a persistent desktop and providing access to temporary or occasional workers via a non-persistent desktop.

Advantages of Desktop as a Service (DaaS)

Desktop as a Service (DaaS) offers some clear advantages over a traditional desktop model. Deploying or decommissioning active end users with DaaS is much faster and less expensive.

- **Faster deployment and decommissioning of active end users**: The desktop is already configured, it just needs to be connected to a new device. For seasonal businesses that consistently experience spikes and drops in demand or employees, DaaS can save a lot of time and money.
- **Reduced downtime for IT support**: Desktop as a Service also allows companies to provide remote IT support to their employees, reducing downtime.
- **Cost savings**: Because the devices that run DaaS require much less computing power than a traditional desktop machine or laptop, they are less expensive and use less power.
- **Increased device flexibility**: DaaS runs on a variety of operating systems and device types, which supports the trend of users bringing their own devices into the office and shifts the burden of supporting the desktop on all of those devices to the cloud service provider.
- **Enhanced security**: Because the data is stored in the data center with DaaS, security risks are considerably lower. If a laptop or mobile device is stolen, it can simply be disconnected from the service. Since none of the data lives on that stolen device, the risk of a thief accessing sensitive data is minimal. Security patches and updates are also easier to install in a DaaS environment because all of the desktops can be updated simultaneously from a remote location.

How does Desktop as a Service (DaaS) work?

With Desktop as a Service (DaaS), the cloud services provider hosts the infrastructure, network resources, and storage in the cloud and streams a virtual desktop to the user's device, where the user can access the desktop's data and applications through a web browser or other software. Organizations may purchase as many virtual desktops as they need through a subscription model.

Because desktop applications stream over the Internet from a centralized server, graphics- intensive applications have historically been hard to use with DaaS. New technology has changed this, and even applications such as computer-aided design (CAD) that require an immense amount of computer power to display quickly can now run easily on DaaS. When the workload on oneserver gets too high, IT administrators can migrate a running virtual machine from one physical server to another in just a few seconds, allowing graphics accelerated or GPU-accelerated applications to run uninterrupted. GPU-accelerated Desktop as a Service (GPU-DaaS) has implications for any industry that requires 3D modeling, high-end graphics, simulations, or video

production. The engineering and design, broadcasting, and architecture industries can all benefit from this technology.

What is the difference between VDI and DaaS?

A virtual desktop infrastructure (VDI) allows organizations to remotely host desktop operating systems on endpoint devices from a centralized server. All of the data lives in the data center server—the endpoint is merely a way for users to access that data over the Internet. VDI requires a costly investment in network, storage, and compute infrastructure in the data center, in addition to an IT Team that is skilled in setting up and managing virtual infrastructures. With the DaaS model, cloud service providers bear the infrastructure set-up cost and the management cost, which can make DaaS much more affordable than setting up a new virtual desktop infrastructure in house, depending on the number of end users served and the price of a subscription.

A company with a large number of users can save money with either VDI or DaaS because the endpoint devices don't need much computing power (most of the processing is happening in the data center). However, serving a large number of users requires a large IT staff that can handle any issues that come up. DaaS lets a company function with a leaner IT staff because the DaaS vendor will deal with deployment, connectivity problems, and other issues that come up for end users.

However, VDI gives a company's IT staff more control over the desktop offering and more control over security than DaaS. A business with specific or stringent security or application requirements may not be able to find a DaaS provider that meets all of its needs in a cost-effective way.

Why Desktop as a Service?

The benefits of Desktop as a Service (DaaS) include simplified management, increased flexibility, and lower cost of ownership compared to traditional models. Businesses that aim to offer remote work options and personal device flexibility can use DaaS to quickly and easily create a digital workspace. Users may log in to their virtual desktop from anywhere, via many different kinds of devices, and their desktop will look exactly the same as when they last visited from a different location. All they need is an internet connection. Since the data lives in a centralized, remote location, it can be constantly backed up – no need for users tot manage back-ups on their own or worry about data existing on a computer at the office but not at home.

5.5. Local desktop Virtualisation

5.6. Virtualisation benefits

BENEFITS OF VIRTUALIZATION

1. More flexible and efficient allocation of resources.
2. Enhance development productivity.
3. It lowers the cost of IT infrastructure.
4. Remote access and rapid scalability.
5. High availability and disaster recovery.
6. Pay peruse of the IT infrastructure on demand.
7. Enables running multiple operating systems.

5.7. Server Virtualisation

It is the division of physical server into several virtual servers and this division is mainly done to improvise the utility of server resource. In other word it is the masking of resources that are located in server which includes the number & identity of processors, physical servers & the operating system.

This division of one physical server into multiple isolated virtual servers is done by server administrator using software. The virtual environment is sometimes called the virtual private-servers.

Usage of Server Virtualization

This technique is mainly used in web-servers which reduces the cost of web-hosting services.Instead of having separate system for each web-server, multiple virtual servers can run on the same system/computer.

- To centralize the server administration
- Improve the availability of server
- Helps in disaster recovery
- Ease in development & testing
- Make efficient use of server resources.

Approaches To Virtualization:

For Server Virtualization, there are three popular approaches.
These are:

- Virtual Machine model
- Para-virtual Machine model
- Operating System (OS) layer Virtualization

1. **Virtual Machine model:** are based on host-guest paradigm, where each guest runs on a virtual replica of hardware layer. This technique of virtualization provide guest OS to run without modification. However it requires real computing resources from the host and for this a hypervisor or VM is required to coordinate instructions to CPU.
2. **Para-Virtual Machine model:** is also based on host-guest paradigm & uses virtual machine monitor too. In this model the VMM modifies the guest operating system's code which is called 'porting'. Like that of virtual machine, similarly the Para-virtual machine is also capable of executing multiple operating systems. The Para-virtual model is used by both Xen & UML.
3. **Operating System Layer Virtualization:** Virtualization at OS level functions in a different way and is not based on host-guest paradigm. In this model the host runs a single operating system kernel as its main/core and transfers its functionality to each of the guests. The guest must use the same operating system as the host. This distributed nature of architecture eliminated system calls between layers and hence reduces overhead of CPU usage. It is also a must that each partition remains strictly isolated from its neighbors because any failure or security breach of one partition won't be able to affect the other partitions.

Advantages of Server Virtualization

- **Cost Reduction:** Server virtualization reduces cost because less hardware is required.
- **Independent Restart:** Each server can be rebooted independently and that reboot won't affect the working of other virtual servers.

5.8. Block and File level Storage Virtualisation

Virtualisation on block level means that storage capacity is made available to the operating system or the applications in the form of virtual disks

In virtualisation on block level the task of file system management is the responsibility of the operating system or the applications

The task of the virtualisation entity is to map these virtual blocks to the physical blocks of the real storage devices

Virtualisation on file level means that the virtualisation entity provides virtual storage to the operating systems or applications in the form of files and directories

The applications work with files instead of blocks and the conversion of the files to virtual blocks is performed by the virtualisation entity itself(This means, the task of file system management is performed by the virtualisation entity, unlike in block level which is done by OS or application)

The physical blocks are presented in the form of a virtual file system and not in the form of virtual blocks.

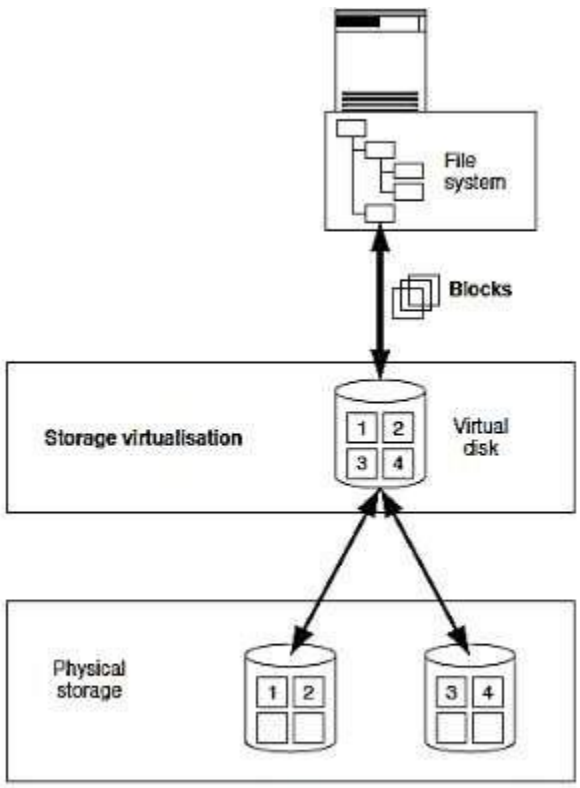

Figure 5.12 In virtualisation on block level the virtualisation entity provides the virtual storage to the servers in the form of a virtual disk.

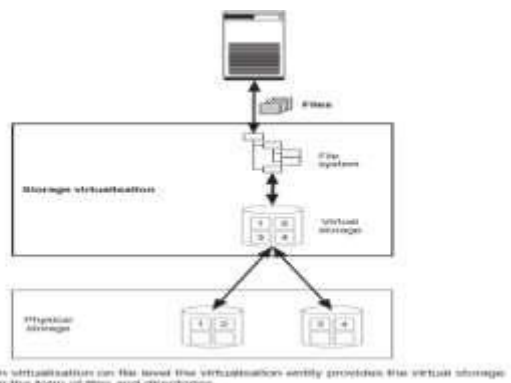

5.9. Virtual Machine Monitor

VMM stands for Virtual Machine Monitor. It is a software program which allows management, governance, and creation of VM (Virtual Machines) and also manages the virtualized environment's operation.

It is also called as **Hypervisor** and Virtual Machine Manager. Though, provided services and architectural implementation differs by product of seller.

Behind implementations and virtualization environments, it is the major software. When VMM is installed on host machine, it facilitates VMs creation and each with different apps and Operating Systems. In addition, it manages backend operations via allocating required **storage**, memory, input/output resources, and computing.

It also offers centralized interface to manage the status, availability, and operation of VMs installed on a host or spread across interconnected hosts.

Getting Virtualized

Since a visitor OS isn't in hardware control, the VMM works as go-between. It blocks calls tomemory tables and peripheral devices from every visitor and intervenes on the behalf of them. At the point of time when peripheral device makes interfere, for instance, when a disk is completed, VMM infuses that disturbance in the suitable visitor OS.

VMM programming and as per the stage, a visitor OS might run in integral in VM environment with no modification or the source code of VMM might need to be adjusted.

Types of VMM

There are three VMM architectures that shows the relationship in between device drivers, VMM, and guest OS. In paravirtualized, in all three architectures, guest OS changes could be made or not a single change can be made in guest OS in the case of fully virtualized. 'Hypervisor' term refers to the component of VMM near the hardware.

Here are Three Types of VMM –

1. **Host OS** – This method helps in enabling VMM in order to get installed on guest OS and running computer with no modification.

2. **Hypervisor** – It offers best performance, flexibility, and most control in **VM (Virtual Machine)** environment.

3. **Service OS** – It combines hypervisor's robustness with hosted model's flexibility that uses existing OS.

5.10. Infrastructure Requirements

Beyond data centers, cloud computing has been a revolutionary technology trend for businesses of all sizes across virtually every industry, and it's become a core component of a modern ecosystem and application integration strategy. Instead of investing in costly hardware while having to manage and maintain a data center in-house, companies are turning to cloud providers like AmazonWeb Services, Google Cloud, and Microsoft Azure for flexible cloud infrastructure to provide modernized computing, networking, and storage resources.

But what does an organization need to consider when selecting and implementing a cloud computing infrastructure that's the best fit for its business ecosystem while meeting workflow requirements and providing ideal results?

This blog will cover the answers to these important questions and a whole lot more you need to know about cloud computing.

What is Cloud Infrastructure?

Defining exactly what a cloud infrastructure is can be broad and complex. But when it comes down to it, a cloud-based infrastructure has several key components, including, but not limited to a combination of:

- Servers
- Software
- Network devices, and
- Other storage resources

It is these components, all of which are necessary to create applications that are then accessed via the cloud. These apps can be retrieved remotely over the internet, telecom services, WANs (wide area networks), and other network means.

For example an EDI provider might offer their services using a cloud EDI software model, allowing clients to access the platform without having to maintain the required physical infrastructure on premise.

How is Cloud Infrastructure Categorized?

Cloud infrastructure generally is categorized into three parts that all collaborate to create a cloud service:

1. Computing: The computing portion of the infrastructure is delivered by server racks in order to deliver cloud services for various services and partners.

2. Networking: To transfer data externally as well as between computer and storage systems, this part of the infrastructure relies on routers and switches.

3. Storage: A cloud infrastructure will likely need considerable storage often using a combination of hard disks and flash storage.

Cloud Infrastructure vs. SaaS, PaaS, and IaaS

There are generally three models when it comes to cloud services: SaaS (Software as aService), PaaS (Platform as a Service), and IaaS (Infrastructure as a Service). Each service hasvarying levels of benefits and differences.

SaaS: Software as a Service (aka cloud application services), is the most widely used type of cloud service. Popular for business as SaaS companies deliver strong customer experience through information exchange and services, SaaS takes a large part of the IT burden off the hands of a business. SaaS employs the internet to provide distributed applications and services, eliminating the need for clients to download any software. With SaaS, a cloud vendor fully manages the entire offering: applications, data, runtime, middleware, operating systems, services, storage, networking, and virtualization.

PaaS: Platform as a Service is somewhat comparable to SaaS, but instead offers a platform to create software. The PaaS method is delivered via the internet, giving IT teams the ability to design software without bothering with other aspects. With PaaS, the cloud vendor shoulders a majority ofthe service, including runtime, middleware, operating systems, servers, storage, networking, and virtualization. The company therefore only needs to worry about managing its applications and data.

IaaS: Infrastructure as a Service offers the most in-house control, allowing access and direct maintenance to most cloud resources. IaaS is extensively automated and scalable, as clients are able to buy resources as-needed without relying on in-house hardware. With IaaS, the cloud services are mostly managed by a company, including applications, data, runtime, middleware, and operating systems. But the cloud vendor is responsible for the services, storage, networking, and virtualization.

Cloud Infrastructure as a Service

A cloud infrastructure as a service (IaaS), just like with all cloud technology, is accessed via the internet through a cloud vendor's data center, which is responsible for maintaining and managing traditional on-premise hardware like servers and other storage devices as well as networking and visualization. That means the customer has the freedom and control to manage application, data, middleware, and other operating systems.

With a cloud IaaS solution, there are important infrastructure services such as network monitoring, security, billing, disaster recovery, and load balancing. There is also advanced automation and orchestration to simplify application performance and management as well as make it easier to install operating systems, deploy middleware, launch virtual machines, and create workload storage and backups.

Cloud Infrastructure Management

The main purpose of cloud infrastructure management is to provide business scalability while consolidating IT resources and enabling a variety of users to share the same infrastructure without compromising each other's data. In the long run, this lowers operating costs.

A cloud infrastructure management interface (CIMI) is an open standard API requirement for handling cloud infrastructure and allows users to manage it in a simplified manner with homogeneous communication between cloud ecosystems. This attains interoperable management among cloud vendors, developers, and customers.

Requirements for Building a Cloud Infrastructure

When building out a cloud strategy, there are several in-depth steps that must be taken to ensure a robust infrastructure.

Requirement 1: Service and Resource Management

A cloud infrastructure virtualizes all components of a data center. Service management is a measured package of applications and services that end users can easily deploy and manage via a public and/or private cloud vendor. And a simplified tool to outline and gauge services is vital for cloud administrators to market functionality. Service management needs to contain resource maintenance, resource guarantees, billing cycles, and measured regulations. Once deployed, management services should help create policies for data and workflows to make sure it's fully efficient and processes are delivered to systems in the cloud.

Requirement 2: Data Center Management Tools Integration

Most data centers utilize a variety of IT tools for systems management, security, provisioning, customer care, billing, and directories, among others. And these work with cloud management services and open APIs to integrate existing operation, administration, maintenance, and provisioning (OAM&P) systems. A modern cloud service should support a data center's existing

infrastructure as well as leveraging modern software, hardware, and virtualization, and other technology.

Requirement 3: Reporting, Visibility, Reliability, a Security

Data centers need high levels of real-time reporting and visibility capabilities in cloud environments to guarantee compliance, SLAs, security, billing, and chargebacks. Without robust reporting and visibility, managing system performance, customer service, and other processes are nearly impossible. And to be wholly reliable, cloud infrastructures must operate regardless of oneor more failing components. For to safeguard the cloud, services must ensure data and apps are secure while providing access to those who are authorized.

Requirement 4: Interfaces for Users, Admins, and Developers

Automated deployment and self-service interfaces ease complex cloud services for end users, helping lower operating costs and deliver adoption. Self-service interfaces offer customers theability to effectively launch a cloud service by managing their own data centers virtually, designing and driving templates, maintaining virtual storage, networking resources, and utilizing libraries. Administrator interfaces present better visibility to all resources, virtual machines, templates, service offers, and various cloud users. And all of these structures integrate by way of APIs for developers.

Advantages of Using Cloud Infrastructure

The arguments in favor of using the cloud are only getting stronger as the technology continues to improve. So, there are some obvious key benefits to migrating to a cloud infrastructure that helps companies streamline business processes.

Cost: First and foremost, the cloud removes or greatly reduces the operating expense of a company setting up and managing its own data center. Taking on this process begins to add up with all the various hardware, software, servers, energy bills, IT experts, and the updates that come along with this multi-faceted set-up. With cloud infrastructure, a company simply pays for it all to be managed while paying only for as-needed services.

Agility and flexibility: Most cloud service infrastructures are offered as self-managed, where service changes can be made within minutes. This improves the uptime and efficiency of business systems while allowing off-site coworkers and partners to access shared data on mobile devices whenever and wherever. And with a cloud infrastructure managing processes, a company becomes more business-focused than IT-focused.

Security: There's a common misconception that cloud services are generally not secure and that data can easily be compromised. There is some truth in that, however, the risks are often blown out of proportion at least in terms of enterprise-level cloud infrastructure and services. Cloud infrastructure technologies and providers are always improving protection against hackers, viruses, and other data breaches with stronger firewalls, advanced encryption keys, and a hybrid approach that stores sensitive data in a private cloud and other data, even apps, in a public cloud.

Disadvantages of Using Cloud Infrastructure

That being said, not all cloud infrastructures are perfect. And while there are far more advantages, there are still some drawbacks.

Vendor overturn: The cloud is still an evolving, albeit improving, technology that rapidly fluctuates. Meaning, some cloud services companies get it right and some don't. If a company goes out of business or sees a massive overhaul, that could be destructive to a business that relies on justone infrastructure for its entire database.

Connection reliance: A cloud infrastructure is only as good as its network connection. Therefore, the cloud can't stay afloat without a dependable connection. Any glitches in an internet or intranet connection due to a technical outage or storm mean the cloud goes down along with all the data, software, and/or applications in it. A reliable network means business promises and SLAs are delivered.

Control: Since a company's cloud infrastructure is generally controlled by its service provider, there are times organizations have limited access to data. And business customers have even less control than they might want, with limited access to applications, data, and tools stored on a server.

5.11. VLAN and VSAN

The main differences between VLAN and VSAN are given below:

S.No.	VLAN(Virtual Local Area Network)	VSAN(Virtual Storage Area Network)
1	VLAN is a network technology used to logically separate large broadcast domains using layer 2 devices.	VSAN is a logical partition in a storage area network.
2	It divides the network into different virtual sub-networks reduces unnecessary traffic and improve performance.	VSANs allow traffic to be isolated within specific portions of a storage area network.
3	VLANs are implemented to achieve scalability, security and ease of network management.	The use of multiple VSAN's can make a system easier to configure and scale out.
4	VLAN's can quickly adapt to change in network requirements and relocation of workstations and server nodes.	In this subscribers can be added or relocated without the need for changing the physical layout.
5	The purpose of implementing a VLAN is to improve the performance of a network or apply appropriate security features.	The VSANs minimizes the total system's vulnerability, security is improved. VSANs also offer the possibility of data redundancy, minimizing the risk of catastrophic data loss.

UNIT-6
Cloud Security

6.1. Cloud Security Fundamentals

There are many reasons why enterprises are using cloud platforms such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) to facilitate web applications. Outsourcing to the cloud adds scalability, efficiency, and reliability, while also reducing workloads for IT teams.

These are positives, for sure – but while enterprise IT leaders celebrate the benefits that the cloud brings to their businesses, they may be missing a big negative. Organizations that move critical infrastructure to cloud platforms often mistakenly assume that their cloud providers also lock down security. In fact, this is often not the case; and when enterprises let critical cloud security slip, their security teams may also fail to configure critical controls or adopt the necessary secure architecture practices, leaving gaps that attackers can compromise.

For example, many of the risks that security teams have to deal with when they're working tokeep attackers from breaching on-premises architecture – like improper segmentation, overly permissive firewall rules, or weak passwords – also exist in the cloud. And there are always new risks that can affect cloud platform security, such as exposure of API keys in source repositories or open web directories. None of these risks should be left unmonitored – yet the default configurations for AWS, Azure, and GCP often don't include turning on event logging, encryption, data retention, multifactor authentication, and other preventative controls.

The steps below detail how to configure and monitor your cloud platforms for improved visibility, which cloud-native tools are needed to secure cloud platforms, and why integration can help secure a multi-cloud environment.

Step 1: Determine where sensitive data lives, and prioritize integrations that increase visibility

Because cloud deployments are relatively easy, it's also easy to move data around from cloud to cloud. For this reason, security teams need to understand where data is stored and how it's used. Without this knowledge, and without controls that manage visibility into sensitive data, it could be painfully easy to transfer customer data from a private server to a public storage repository.

Typically, the flow of data in the cloud should be traced from the point where the application is accessed, and back to where a company's developers eventually access the systems on which data is stored. Security teams need knowledge of how data moves through the environment – ifnot, they'll waste time and money securing potentially lower-priority infrastructure devices.

Fortunately, the major cloud platforms have out-of-the-box security tools and open APIs that make log ingestion easy. Once security teams know what they're looking for, they can use proper tuning and integrations to make alerting and visibility into attacks as simple to managing on-premises tools.

Step 2: Configure cloud platforms to maximize the security of their architecture

It's worth spending time to figure out which features a specific cloud platform already provides for visibility and automation. Once these basic tools are enabled, security experts can start the process of fine-tuning and tightening controls. For example, teams can set up alerts for unusual calls from accounts, repeated denies, policy changes, and other actions and content that help pinpoint attacker activity. For more information on platform-specific features on AWS, Azure,and GCP, check out the Tactical Guide to Securing Data on Cloud Platforms.

At this stage, review the areas of vulnerability that can allow bad actors to pivot their attacks, including Identity Access Management (IAM), cloud infrastructure, server infrastructure, and application security. Application security and IAM should be the primary areas of concern. IAM can be secured through the combination of access logs and regular auditing and tightening of permissions; application security requires security by design, via developers who care about the security of their application.

Step 3: Monitor the cloud through integration

Most cloud environments generate a large number of logs that can quickly take up all of a security team's time. To constantly parse, tune, and respond to alerts, teams need automation and integration. To efficiently and effectively respond to alerts, teams need a centralized view of their data across cloud and on-premises environments, such as through an Open XDR solution. And, to see what's happening in the cloud environment and export logs into an alerting engine such as a SIEM, enterprise organizations need to plan a cloud-logging strategy by answering these questions:

- **Where will log aggregation and filtering tools reside?** It's efficient to have collection tools doing the filtering and aggregating as close to the source as possible.
- **How big are the Internet connections between your cloud environment and local data centers?** It may be more cost-effective to keep raw data at both places, and send actionable events used in alerting rules to the central SIEM or monitoring tool.
- **How will you collect and parse cloud infrastructure logs?** In addition to standard operating system or application logs from servers, many essential cloud infrastructure logs should be gathered and monitored for unauthorized or malicious activities.

6.2. Cloud security services

6.3. Design Principles

6.4. Secure Cloud software requirements

6.5. Policy Implementation

6.6. Cloud Computing Security Challenges

UNIT-9
Hadoop

Traditional Approach

In this approach, an enterprise will have a computer to store and process big data. For storage purpose, the programmers will take the help of their choice of database vendors such as Oracle, IBM, etc. In this approach, the user interacts with the application, which in turn handles the part of data storage and analysis.

Limitation

This approach works fine with those applications that process less voluminous data that can be accommodated by standard database servers, or up to the limit of the processor that is processing the data. But when it comes to dealing with huge amounts of scalable data, it is a hectic task to process such data through a single database bottleneck.

Google's Solution

Google solved this problem using an algorithm called MapReduce. This algorithm divides the task into small parts and assigns them to many computers, and collects the results from them which when integrated, form the result dataset.

Hadoop

Using the solution provided by Google, **Doug Cutting** and his team developed an Open Source Project called **HADOOP**.

Hadoop runs applications using the MapReduce algorithm, where the data is processed in parallel with others. In short, Hadoop is used to develop applications that could perform complete statistical analysis on huge amounts of data.

Case Study of Haddop

Hadoop is an open-source java-based software framework sponsored by the Apache Software Foundation for distributed storage and distributed processing of very large data sets on computer clusters built from commodity hardware.

It provides storage for big data at reasonable cost. Hadoop process big data in a single place as in a storage cluster doubling as a compute cluster.

Hadoop Architecture and Components:

Apache Hadoop consist of two major parts:

1. Hadoop Distributed File System (HDFS)
2. MapReduce

1. Hadoop Distributed File System:

HDFS is a file system or storage layer of Hadoop. It can store data and can handle very large amount of data.

When capacity of file is large then it is necessary to partition it. And the file systems manage the storage across a network of machine are called distributed file systems.

An HDFS cluster has two types of node operating in a master-worker pattern- Name Node and No. of Data Nodes.

Hadoop keep data safe by duplicating data across nodes.

2. MapReduce:

MapReduce is a programming framework. It organize multiple computers in a cluster in order to perform the calculations. It takes care of distributing the work between computers and putting results together.

Apache Hadoop is an open source framework that is used to efficiently store and process large datasets ranging in size from gigabytes to petabytes of data. Instead of using one large computer to store and process the data, Hadoop allows clustering multiple computers to analyze massive datasets in parallel more quickly.

Hadoop consists of four main modules:

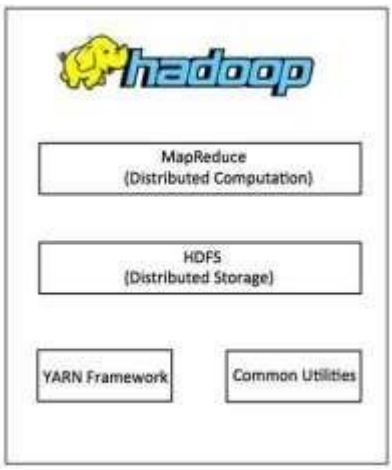

- Hadoop Distributed File System (HDFS) – A distributed file system that runs on standard or low- end hardware. HDFS provides better data throughput than traditional file systems, in addition to high fault tolerance and native support of large datasets.
- Yet Another Resource Negotiator (YARN) – Manages and monitors cluster nodes and resource usage. It schedules jobs and tasks.
- MapReduce – A framework that helps programs do the parallel computation on data. The map task takes input data and converts it into a dataset that can be computed in key value pairs. The output of the map task is consumed by reduce tasks to aggregate output and provide the desired result.
- Hadoop Common – Provides common Java libraries that can be used across all modules.

Hadoop works in a Master-Worker / Master-slave fashion:-

1. Master:
Master contains Name node and Job tracker components.

1. **Name node:** It holds information about all the other nodes in the Hadoop Cluster, files in the cluster, blocks of files, their locations etc.
2. **Job tracker:** It keeps track of the individual tasks assigned to each of the nodes and coordinates the exchange of information and result.

2. Worker:
Worker contains Task tracker and Data node components.

1. **Task Tracker:** It is responsible for running the task assigned to it.
2. **Data node:** It is responsible for holding the data.

Other components of Hadoop architecture are : -Chukwa, Hive, HBase, Mahoutetc.

Characteristics of Hadoop:

1. Hadoop provides a reliable shared storage(HDFS) and analysis system (Map Reduce).
2. Hadoop is highly scalable. It can contain thousands of servers.
3. Hadoop works on the principles of write once and read multiple times.
4. Hadoop is highly flexible, can process both structured as well as unstructured data.

What is the Hadoop ecosystem?

The term **Hadoop** is a general term that may refer to any of the following:

- **The overall Hadoop ecosystem,** which encompasses both the core modules and related sub-modules.
- **The core Hadoop modules,** including Hadoop Distributed File System (HDFS™), Yet Another Resource Negotiator (YARN), MapReduce, and Hadoop Common (discussed below). These are the basic building blocks of a typical Hadoop deployment.
- **Hadoop-related sub-modules,** including: Apache **Hive**™, Apache **Impala**™, Apache **Pig**™, and Apache **Zookeeper**™, among others. These related pieces of software can be used to customize, improve upon, or extend the functionality of core Hadoop.

DATA SOURCES - HADOOP

We live in the data age.It's not easy to measure the total volume of data stored electronically,but an IDC estimate put the size of the "digital universe" at 0.18 zettabytes in 2006, and is forecasting a tenfold growth by 2011 to 1.8 zettabytes.* A zettabyte is 1021 bytes, or equivalently one thousand exabytes, one million petabytes, or one billion terabytes. That's roughly the same order of magnitude as one disk drive for every personin the world.

This flood of data is coming from many sources. Consider the following:

- The New York Stock Exchange generates about one terabyte of new trade data perday.
- Facebook hosts approximately 10 billion photos, taking up one petabyte of storage.
- Ancestry.com, the genealogy site, stores around 2.5 petabytes of data.
- The Internet Archive stores around 2 petabytes of data, and is growing at a rate of 20 terabytes per month.
- The Large Hadron Collider near Geneva, Switzerland, will produce about 15 petabytes of data per year.

So there's a lot of data out there. But you are probably wondering how it affects you. Most of the data is locked up in the largest web properties (like search engines), orscientific or financial institutions, isn't it? Does the advent of "Big Data," as it is being called, affect smaller organizations or individuals.

The trend is for every individual's data footprint to grow, but perhaps more important,the amount of data generated by machines will be even greater than that generated by people. Machine logs, RFID readers, sensor networks, vehicle GPS traces, retail transactions all of these contribute to the growing mountain of data.

The volume of data being made publicly available increases every year, too. Organizations no longer have to merely manage their own data: success in the future will bedictated to a large extent by their ability to extract value from other organizations' data.

Initiatives such as Public Data Sets on Amazon Web Services, Infochimps.org, and the info.org exist to foster the "information commons," where data can be freely (or inthe case of AWS, for a modest price) shared for anyone to download and analyze. Mashups between different information sources make for unexpected and hitherto unimaginable applications.

Take, for example, the Astrometry.net project, which watches the Astrometry groupon Flickr for new photos of the night sky. It analyzes each image and identifies which part of the sky it is from, aswell as any interesting celestial bodies, such as stars orgalaxies. This project shows the kind of things that are possible when data (in this case,tagged photographic images) is made available and used for something (image analysis)that was not anticipated by the creator.

It has been said that "More data usually beats better algorithms," which is to say that for some problems (such as recommending movies or music based on past preferences),however fiendish your algorithms are, they can often be beaten simply by having more data (and a less sophisticated algorithm).

The good news is that Big Data is here. The bad news is that we are struggling to store and analyze it.

DATA STORAGE AND ANALYSIS

The problem is simple: while the storage capacities of hard drives have increased massively over the years, access speeds the rate at which data can be read from drives have not kept up.One typicaldrive from 1990 could store 1,370 MB of data and hada transfer speed of 4.4 MB/s, so you could read all the data from a full drive in around five minutes. Over 20 years later, one terabyte drivesare the norm, but the transfer speed is around 100 MB/s, so it takes more than two and a half hoursto read all the data off the disk.

This is a long time to read all data on a single drive and writing is even slower.The obvious way to reduce the time is to read from multiple disks at once.Imagine if we had 100 drives, each holding one hundredth of the data.Working in parallel,we could read the data in under two minutes.

Only using one hundredth of a disk may seem wasteful.But we can store one hundred data sets, each of which is one terabyte, and provide shared ould be likely to be spread over time, so they wouldn't interfere with each other too much.

There's more to being able to read and write data in parallel to or from multiple disks,though.

The first problem to solve is hardware failure: as soon as you start using many pieces of hardware,the chance that one will fail is fairly high. A common way of avoiding data loss is through replication: redundant copies of the data are kept by the system so that in the event of failure, there is another copy available. This is how RAID works, for instance, although Hadoop's file system, the Hadoop Distributed Filesystem (HDFS),takes a slightly different approach, as you shall see later.

The second problem is that most analysis tasks need to be able to combine the data in some way; data read from one disk may need to be combined with the data from any of the other 99 disks. Various distributed systems allow data to be combined from multiple sources, but doing this correctly is notoriously challenging. Map Reduce provides a programming model that abstracts the problem from disk reads and writes,transforming it into a computation over sets of keys andvalues. We will look at the details of this model in later chapters, but the important point for the present discussion is that there are two parts to the computation, the map and the reduce, and it's theinterface between the two where the "mixing" occurs. Like HDFS, MapReduce has built-in reliability.

This, in a nut shell, is what Hadoop provides: a reliable shared storage and analysis system. The storage is provided by HDFS and analysis by MapReduce. There are other parts to Hadoop, but these capabilities are its kernel.

COMPARISON WITH OTHER SYSTEMS - HADOOP

The approach taken by MapReduce may seem like a brute-force approach. The premise is that the entire dataset or at least a good portion of it is processed for each query. But this is its power. MapReduce is a batch query processor, and the ability to run an ad hoc query against your whole dataset and get the results in a reasonable time is transformative.It changes the way you think about data, and unlocks data that was previously archived on tape or disk.It gives people the opportunity to innovate with data.Questions that took too long to get answered before can now be answered, which in turn leads to new questions and new insights.

For example,Mailtrust, Rackspace's mail division, used Hadoop for processing email logs. One ad hoc query they wrote was to find the geographic distribution of their users. In their words:

This data was so useful that we've scheduled the MapReduce job to run monthly and we will be using this data to help us decide which Rackspace data centers to place new mail servers in as we grow.

By bringing several hundred gigabytes of data together and having the tools to analyze it, the Rackspace engineers were able to gain an understanding of the data that they otherwise would never have had, and, further more, they were able to use what they had learned to improve the service for their customers.

RDBMS

Why can't we use databases with lots of disks to do large-scale batch analysis? Why is MapReduce needed?

The answer to these questions comes from another trend in disk drives: seek time is improving more slowly than transfer rate. Seeking is the process of moving the disk's head to a particular place on the disk to read or write data. It characterizes the latency of a disk operation, where as the transfer rate corresponds to a disk's bandwidth.

If the data access pattern is dominated by seeks, it will take longer to read or write large portions of the dataset than streaming through it,which operates at the transfer rate. On the other hand, for updating a small proportion of records in a database,a traditional B-Tree (the data structure used in relational databases, which is limited by the rate it can perform seeks) works well.For updating the majority of a database, a B-Tree is less efficient than MapReduce, which uses Sort / Merge to rebuild the database.

In many ways, MapReduce can be seen as a complement to an RDBMS.(The differences between the two systems are shown in below) MapReduce is a good fit for problems that need to analyze the whole dataset, in a batch fashion, particularly for ad hoc analysis. An RDBMS is good for point queries or updates, where the data set has been indexed to deliver low-latency retrieval and updatetimes of a relatively small amount of data. MapReduce suits applications where the data is written once, and read many times, whereas a relational database is good for datasets that are continually updated.

Table RDBMS compared to MapReduce

	Traditional RDBMS	MapReduce
Data size	Gigabytes	Petabytes
Access	Interactive and batch	Batch
Updates	Read and write many times	Write once, read many times
Structure	Static schema	Dynamic schema
Integrity	High	Low
Scaling	Nonlinear	Linear

Another difference between MapReduce and an RDBMS is the amount of structure in the datasets that they operate on. Structured data is data that is organized into entities that have a defined format, such as XML documents or database tables that conform to a particular predefined schema. This is the realm of the RDBMS. Semi-structured data, on the other hand, is looser, and though theremay be a schema, it is often ignored, so it may be used only as a guide to the structure of the data: for example, a spread sheet, in which the structure is the grid of cells, although the cells themselves may hold any form of data. Unstructured data does not have any particular internal structure: for example, plain text or image data. MapReduce works well on unstructured or semistructured data, since it is designed to interpret the data at processing time. In other words, the input keys and values for MapReduce are not an intrinsic property of the data, but they are chosen by the person analyzing the data.

Relational data is often normalized to retain its integrity and remove redundancy. Normalization poses problems for MapReduce, since it makes reading a record a non local operation, and one of the central assumptions that MapReduce makes is that it is possible to perform (high-speed) streaming reads and writes.

A web server log is a good example of a set of records that is not normalized (for example, the client host names are specified in full each time, even though the same client may appear many times), and this is one reason that logfiles of all kinds are particularly well-suited to analysis with MapReduce.

MapReduce is a linearly scalable programming model. The programmer writes two functions a map function and a reduce function each of which defines a mapping from one set of key-value pairs to another. These functions are oblivious to the size of the data or the cluster that they are operating

on, so they can be used unchanged for a small dataset and for a massive one. More important, if you double the size of the input data, a job will run twice as slow. But if you also double the size of the cluster, a job will run as fast as the original one. This is not generally true of SQL queries.

Over time, however, the differences between relational databases and MapReduce systems arelikely to blur both as relational databases start incorporating some of the ideas from MapReduce (such as Aster Data's and Greenplum's databases) and, from the other direction, as higher-level query languages built on MapReduce (such as Pig and Hive) make MapReduce systems more approachable to traditional database programmers.

Grid Computing

The High Performance Computing (HPC) and Grid Computing communities have been doing large-scale data processing for years, using such APIs as Message Passing Interface (MPI). Broadly, the approach in HPC is to distribute the work across a cluster of machines, which access a shared filesystem, hosted by a SAN. This works well for predominantly compute-intensive jobs, but becomes a problem when nodes need to access larger data volumes (hundreds of gigabytes, the point at which MapReduce really starts to shine), since the network bandwidth is the bottleneck and compute nodes become idle.

MapReduce tries to allocate the data with the compute node, so data access is fast since it is local. This feature, known as data locality, is at the heart of MapReduce and is the reason for its good performance. Recognizing that network bandwidth is the most precious resource in a data center environment (it is easy to saturate network links by copying data around), MapReduce implementations go to great lengths to conserve it by explicitly modelling network topology. Notice that this arrangement does not pre include high-CPU analysis in MapReduce.

MPI gives great control to the programmer, but requires that he or she explicitly handle the mechanics of the data flow, exposed via low-level C routines and constructs, such as sockets, as well as the higher-level algorithm for the analysis. MapReduce operates only at the higher level: the programmer thinks in terms of functions of key and value pairs, and the data flow is implicit.

Coordinating the processes in a large-scale distributed computation is a challenge. The hardest aspect is gracefully handling partial failure when you don't know if a remote process has failed or not and still making progress with the overall computation. MapReduce spares the programmer from having to think about failure, since the implementation detects failed map or reduce tasks and reschedules replacements on machines that are healthy. MapReduce is able to do this since it is a shared-nothing architecture, meaning that tasks have no dependence on one other. (This is a slight over simplification, since the output from mappers is fed to the reducers, but this is under the control of the MapReduce system; in this case, it needs to take more care rerunning a failed reducer than rerunning a failed map, since it has to make sure it can retrieve the necessary map outputs,and if not, regenerate them by running the relevant maps again.) So from the programmer's point ofview, the order in which the tasks run doesn't matter. By contrast, MPI programs have to explicitly manage their own check pointing and recovery, which gives more control to the programmer, but makes them more difficult to write.

MapReduce might sound like quite a restrictive programming model, and in a sense it is: you are limited to key and value types that are related in specified ways, and mappers and reducers run with very limited coordination between one another (the mappers pass keys and values to reducers). The answer is yes. MapReduce was invented by engineers at Google as a system for

building production search indexes because they found themselves solving the same problem over and over again (and MapReduce was inspired by older ideas from the functional programming, distributed computing, and database communities), but it has since been used for many other applications in many other industries. It is pleasantly surprising to see the range of algorithms that can be expressed in MapReduce, from image analysis, to graph-based problems, to machine learning algorithms.It can't solve every problem, of course, but it is a general data-processing tool.

Difference Between RDBMS and Hadoop

S.No.	RDBMS	Hadoop
1.	Traditional row-column based databases, basically used for data storage, manipulation and retrieval.	An open-source software used for storing data and running applications or processes concurrently.
2.	In this structured data is mostly processed.	In this both structured and unstructured data is processed.
3.	It is best suited for OLTP environment.	It is best suited for BIG data.
4.	It is less scalable than Hadoop.	It is highly scalable.
5.	Data normalization is required in RDBMS.	Data normalization is not required in Hadoop.
6.	It stores transformed and aggregated data.	It stores huge volume of data.
7.	It has no latency in response.	It has some latency in response.
8.	The data schema of RDBMS is static type.	The data schema of Hadoop is dynamic type.
9.	High data integrity available.	Low data integrity available than RDBMS.
10.	Cost is applicable for licensed software.	Free of cost, as it is an open source software.