SATHYABAMA
INSTITUTE OF SCIENCE AND TECHNOLOGY
(DEEMED TO BE UNIVERSITY)
Accredited "A" Grade by NAAC | 12B Status by UGC | Approved by AICTE
www.sathyabama.ac.in

SCHOOL OF ELECTRICAL AND ELECRONICS ENGINEERING

DEPARTMENT OF ELECTRONICS AND COMMUNICATION ENGINEERING

UNIT – I – Cloud Computing – SECA5303

UNIT 1

INTRODUCTION TO CLOUD COMPUTING

Introduction to Cloud Computing-Definition, Characteristics, Components,

1.0 Introduction

Cloud computing is the on-demand availability of computer system resources, especially data storage (cloud storage) and computing power, without direct active management by the user. The term is generally used to describe data centers available to many users over the Internet. Large clouds, predominant today, often have functions distributed over multiple locations from central servers. If the connection to the user is relatively close, it may be designated an edge server.

1.0.1 Difference between Cloud Computing and Web 2.0

Cloud Computing	**Web 2.0**
It is more specific and definite	Programming and business models
It is a way of searching through data.	It is sharing entire pieces of data between different websites.
Cloud computing is about computers.	Web 2.0 is about people.
The internet as a computing platform	Attempt to explore and explain the business rules of that platform
Google apps are considered in Cloud computing.	A web-based application is considered in Web 2.0.
It is a business model for hosting these services.	It is a technology which allows web pages to act as more responsive applications

1.1.Key challenges in cloud computing

- Bandwidth cost
- Cloud computing saves the hardware acquisition costs but their expenditure on bandwidth rises considerably.
- Sufficient bandwidth is required to deliver intensive and complex data over the network.
- Continuous monitoring and supervision
- It is important to monitor the cloud service continuously as well as to supervise its

performance, business dependency and robustness.

- Security concerns.

To prevent cloud infrastructure damages, some of the measures include tracking unusual behaviour across servers, buying security hardware and using security applications.

1.2 Data access and integration

In cloud where the data stored, how to access it, who is the owner and how to control it.

Companies are often concerned about data ownership and loss of data control while moving to cloud.

The integration of existing applications in the cloud for smooth running is another challenge.

1.2.1 Proper usability

Enterprises need to have a good and clear view of how to use the technology to add value to their unique businesses

1.2.2 Migration issues

Migrating data from system to the cloud can pose major risks, if it not handled properly.

It need to develop migration strategy that integrates well with the current IT infrastructure.

1.2.3 Cost assessment

Scalable and on-demand nature of cloud services makes the assessment of cost difficult.

Heavy use of a service for a few days may consume the budget of several months.

Current Cloud Computing Challenges

1.3 Requirement of Cloud

In January 2016, RightScale conducted its fifth annual State of the Cloud Survey on the latest cloud computing trends. They questioned 1,060 technical professionals across a broad cross-section of organizations about their adoption of cloud infrastructure.

1.3.1 Lack of resources/expertise

The organizations are increasingly placing more workloads in the cloud while cloud technologies continue to rapidly advance.

Due to these factors organizations are having a hard time keeping up with the tools. Also, the need for expertise continues to grow.

These challenges can be minimized through additional training of IT and development staff.

1.3.2. Security issues

The start of cloud computing technology: you are unable to see the exact location where your data is stored or being processed.

1.3.3. Compliance

The organization needs to be able to comply with regulations and standards, no matter where your data is stored.

Speaking of storage, also ensure the provider has strict data recovery policies in place.

1.3.4. Cost management and containment

The most part cloud computing can save businesses money.

To avoid purchase of new hardware

Use pay-as-you go models

1.3.5. Governance / Control

Proper IT governance should ensure that IT assets are implemented and used according to agreed upon policies and procedures.

In today's cloud-based world, IT does not always have full control over the provisioning, de-provisioning and operations of infrastructure.

This has increased the difficulty for IT to provide the governance, compliance and risk management required.

1.3.6. Performance

When a business moves to the cloud it becomes dependent on the service providers.

When your provider is down, you are also down.

Make sure your SaaS provider has real time monitoring policies in place to help mitigate these issues.

So, a robust cloud adoption strategy in place when they started to move to the cloud.

The Top 5 Cloud-Computing Vendors:

#1 Microsoft,

#2 Amazon,

#3 IBM,

#4Salesforce,

#5 SAP

#1 Microsoft remains an absolute lock at the top due to four factors: its deep involvement at all three layers of the cloud (IaaS, PaaS and SaaS); its unmatched commitment to developing and helping customers deploy AI, ML and Blockchain in innovative production environments; its market-leading cloud revenue, which I estimate at about $16.7 billion for the trailing 12 months (not to be confused with the forward-projected $20.4 billion annualized run rate the company released on Oct. 26); and the extraordinary vision and leadership of CEO Satya Nadella.

#2 Amazon might not have the end-to-end software chops of the others in the Top 5 but it was and continues to be the poster-child for the cloud-computing movement: the first-moving paradigm-buster and category creator. I believe Amazon will make some big moves to bolster its position in software, and no matter how you slice it, the $16 billion in trailing-12-month cloud revenue from AWS is awfully impressive.

#3 IBM has leapfrogged both Salesforce.com (formerly tied with Amazon for #2 and now in the #4 spot) and SAP (formerly #4) on the strength of its un-trendy but highly successful emphasis on transforming its vast array of software expertise and technology from the on-premises world to the cloud. In so doing, IBM has quietly created a $15.8-billion cloud business (again on trailing-12-month basis) that includes revenue of $7 billion from helping big global corporations convert legacy systems to cloud or cloud-enabled environments. And like #1 Microsoft, IBM plays in all three layers of the cloud—IaaS, PaaS and SaaS—which is hugely important for the elite cloud vendors because it allows them to give customers more choices, more seamless integration, better cybersecurity, and more reasons for third-party developers to rally to the IBM Cloud. Plus, its relentless pairing of "cloud and cognitive" is an excellent approach toward weaving AI and ML deeply into customer-facing solutions.

#4 Salesforce.com falls a couple of spots from its long-time tie with Amazon at #2 but—and this will be the case as long as founder Marc Benioff is CEO—remains a powerful source of digital innovation and disruptive strategy. However, to remain in the rarified air near the top of the Cloud Wars Top 10, Benioff and Salesforce must find a way to extend their market impact beyond their enormously successful SaaS business and become more of a high-impact player in the platform or PaaS space. At this stage, it's simply not possible for Salesforce to become a player in IaaS, so Benioff needs to crank up the genius machine and hammer his way into top contention as a platform powerhouse.

#5 SAP has what all of the other cloud vendors would kill for: unmatched incumbency within all of the world's leading corporations as the supplier of mission-critical business applications that run those companies. It's also fashioned, under CEO Bill McDermott, powerful new partnerships with Amazon and Google to complement its long-standing relationships with IBM and Microsoft, all of which give customers a heightened sense of confidence that SAP will be willing and able to play nice in heterogeneous environments. Plus, SAP's HANA technology is now in full deployment across thousands of businesses, and as it takes root and SAP continues to rationalize its massive product portfolio around HANA in the cloud, SAP has a very bright future ahead of it in the cloud.

So the Cloud Wars are clearly intensifying—particularly among the five most-

powerful players at the top—as business customers are fully embracing the cloud as the best and most- capable approach toward digital transformation.And in such an environment, the long-term winners in the Cloud Wars will be those tech companies that choose to view the world by and create their strategies around what business customers want and need, rather than by the distorted and distorting tech-centric perspectives of the Silicon Valley bubble.

1.4 Cloud Deployment Models

The four types of Cloud Deployment Models identified by NIST.

- Private cloud
- Community cloud
- Public cloud
- Hybrid cloud

1.4.1 Private Cloud

Cloud environments are defined based on hardware location and owner. Private clouds are accessible only to a respective customer residing either on-site or be outsourced by a third party.

The cloud infrastructure is operated solely for an organization.

Contrary to popular belief, private cloud may exist off premises and can be managed by a third party.

Thus, two private cloud scenarios exist, as follows:

- On-site Private Cloud
- Applies to private clouds implemented at a customer’s premises.
- Outsourced Private Cloud
- Applies to private clouds where the server side is outsourced to a hosting company.
- Examples of Private Cloud:
- Eucalyptus
- Ubuntu Enterprise Cloud - UEC (powered by Eucalyptus)
- Amazon VPC (Virtual Private Cloud)
- VMware Cloud Infrastructure Suite
- Microsoft ECI data center.

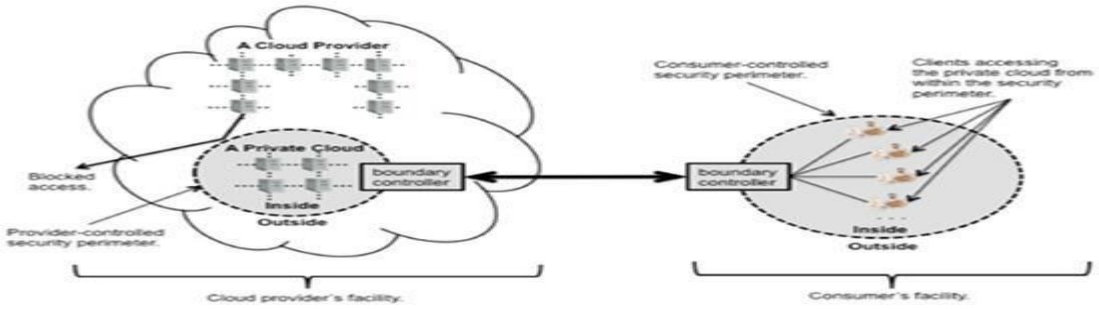

Figure 1.1 Schematic Sketch of Private Cloud

1.4.2 Community Cloud

The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). Government departments, universities, central banks etc. often find this type of cloud useful.

Community cloud also has two possible scenarios:

On-site Community Cloud Scenario

Applies to community clouds implemented on the premises of the customers composing a community cloud

Outsourced Community Cloud

- Applies to community clouds where the server side is outsourced to a hosting company.
- Examples of Community Cloud:
- Google Apps for Government
- Microsoft Government Community Cloud

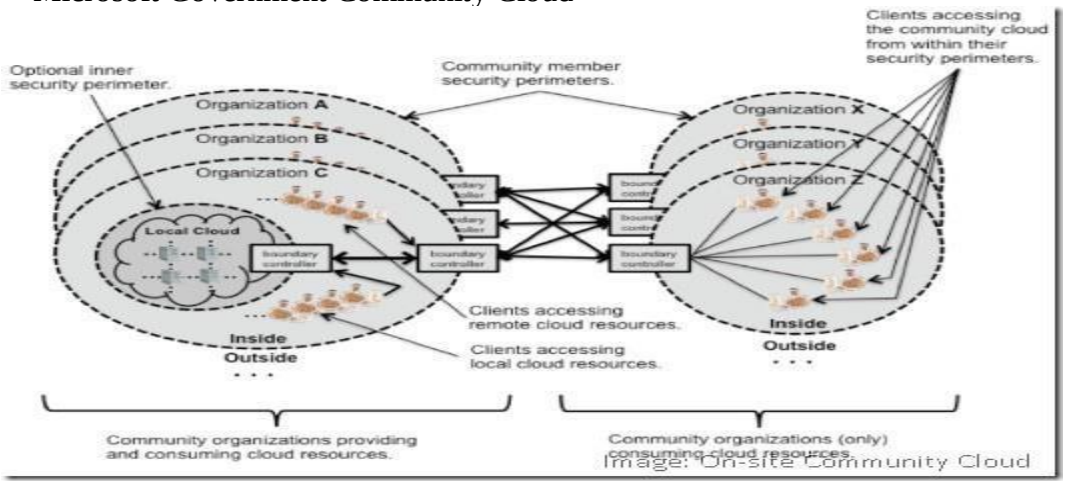

Image: On-site Community Cloud

Figure 1.2. Schematic Sketch of Public Cloud

1.4.3 Public Cloud

The most ubiquitous, and almost a synonym for, cloud computing. The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

Examples of Public Cloud:

- Google App Engine
- Microsoft Windows Azure
- IBM Smart Cloud
- Amazon EC2

1.4.4 Hybrid Cloud

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

Examples of Hybrid Cloud:

- Windows Azure (capable of Hybrid Cloud)
- VMware vCloud (Hybrid Cloud Services)

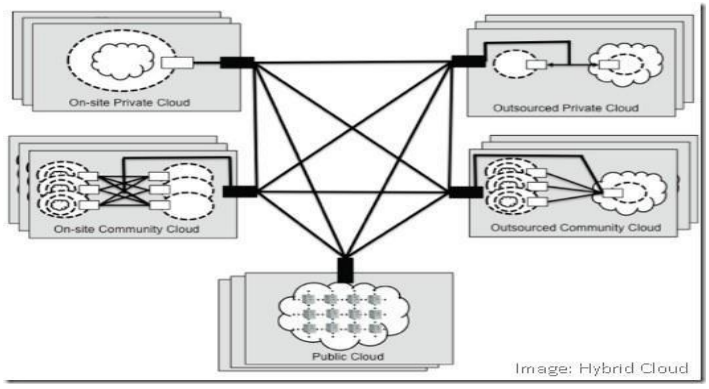

Figure 1.3. Schematic Sketch of Hybrid Cloud

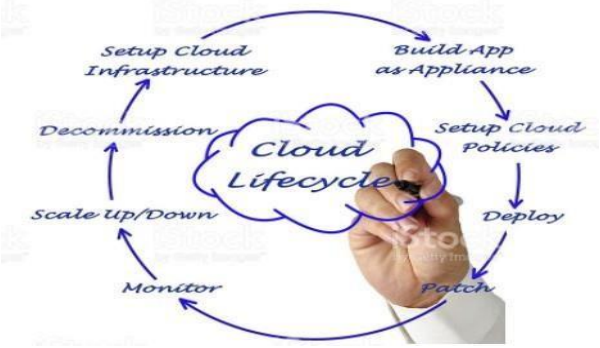

Figure 1.4. Cloud Computing Life Cycle

1.5 Cloud Components

It has three components

- Client computers
- Distributed Servers
- Datacenters

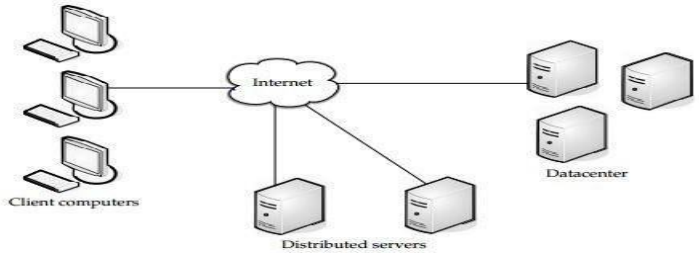

Figure 1.4. Schematic Sketch of Interconnection of Cloud Components

1.Clients

Clients are the device that the end user interacts with cloud.

Three types of clients: 1.) Mobile

2. Thick

3.Thin (Most Popular)

Datacenter

It is collection of servers where application is placed and is accessed via internet.

Distributed servers

Often servers are in geographically different places, but server acts as if they are working next to each other

Who Benefits from Cloud Computing

- Collaborators
- Road Warriors
- Cost-Conscious Users
- Cost-Conscious IT Departments
- Users with Increasing Needs

Who Should not be using Cloud Computing

- The internet impaired
- Offline workers

- The security conscious
- Anyone married to existing applications

Dark Clouds: Barriers to use web-based applications

- Technical issues
- Business model issues
- Internet issues
- Security issues
- Compatibility issues
- Social issues

Cloud Computing Applications

- Clients would be able to access their applications and data from anywhere at any time.
- It could bring hardware costs down.
- The right software available in place to achieve goals.
- Servers and digital storage devices take up space.
- Corporations might save money on IT support.
- The client takes advantage of the entire network's processing power.

Virtualization in cloud computing

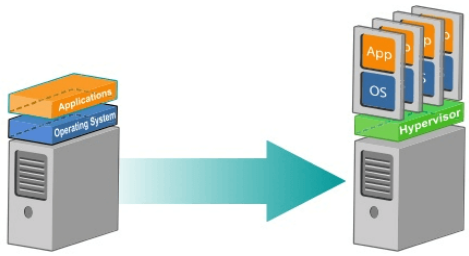

Figure 1.5. Concept of Virtualization

- Virtualization is Nothing but Creating Virtual Version of Something, in Computer Terms it can be OS, Storage Device, Network Resources.
- Virtualization is the ability to run multiple operating systems on a single physical system and share the underlying hardware resources.
- It is the process by which one computer hosts the appearance of many computers.
- Virtualization is used to improve IT throughput and costs by using physical resources as a pool from which virtual resources can be allocated.
- With VMware virtualization solutions you can reduce IT costs while increasing the

efficiency, utilization and flexibility of their existing computer hardware.

- You don't need to own the hardware

Resources are rented as needed from a cloud

You get billed only for what you used

Virtualization Architecture

A Virtual machine (VM) is an isolated runtime environment (guest OS and applications)

Multiple virtual systems (VMs) can run on a single physical system

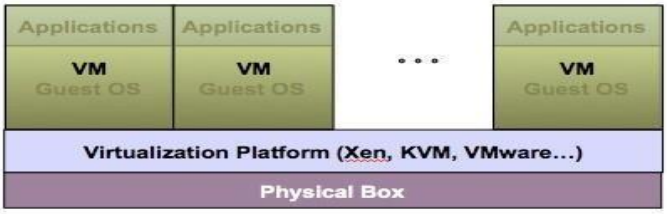

Figure 1.6. Concept of Virtualization

Before Virtualization

- Single OS image per machine
- Software and hardware tightly coupled
- Running multiple applications on same machine often creates conflict
- Inflexible and costly infrastructure

After virtualization

- Hardware-independence of operating system and applications
- Virtual machines can be provisioned to any system
- Can manage OS and application as a single unit by encapsulating them into virtual Machines
- Benefits of Virtualization
- Sharing of resources helps cost reduction.

Isolation: Virtual machines are isolated from each other as if they are physically separated.

Encapsulation: Virtual machines encapsulate a complete computing environment.

Hardware Independence: Virtual machines run independently of underlying hardware.

Portability: Virtual machines can be migrated between different hosts.

What makes virtualization possible?

There is a software that makes virtualization possible. This software is known as a Hypervisor, also known as a virtualization manager.

It sits between the hardware and the operating system and assigns the amount of access that the applications and operating systems have with the processor and other hardware resources.

1.6 Hypervisor

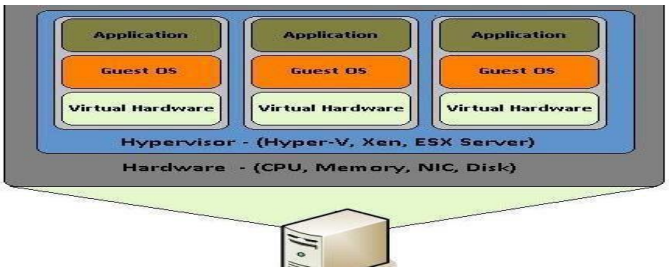

Figure 1.7. Schematic Sketch of Hypervisor

The term hypervisor was first coined in 1956 by IBM

Hypervisor acts as a link between the hardware and the virtual environment and distributes the hardware resources such as CPU usage, memory allotment between the different virtual environments.

A hypervisor or virtual machine monitor (VMM) is computer software, firmware or hardware that creates and runs virtual machines. A computer on which a hypervisor runs one or more virtual machines is called a host machine, and each virtual machine is called a guest machine.

A hypervisor is a hardware virtualization technique that allows multiple guest operating systems (OS) to run on a single host system at the same time. The guest OS shares the hardware of the host computer, such that each OS appears to have its own processor, memory and other hardware resources.A hypervisor is also known as a virtual machine manager (VMM).

1.6.1Types of virtualization

The 7 Types of Virtualization

- Hardware Virtualization.
- Software Virtualization.
- Network Virtualization.
- Storage Virtualization.
- Memory Virtualization.
- Data Virtualization.
- Desktop Virtualization.

1.6.1.1 Hardware Virtualization

Hardware or platform virtualization means creation of virtual machine that act like real computer.

Ex. Computer running Microsoft Windows 7 may host the virtual machine look like a Ubundu

Hardware virtualization also knows as hardware-assisted virtualization or server virtualization.

The basic idea of the technology is to combine many small physical servers into one large physical server, so that the processor can be used more effectively and efficiently.

Each small server can host a virtual machine, but the entire cluster of servers is treated as a single device by any process requesting the hardware.

The hardware resource allotment is done by the hypervisor.

The advantages are increased processing power as a result of maximized hardware utilization and application uptime.

Hardware virtualization is further subdivided into the following types

1.6.1.2Full Virtualization – Guest software does not require any modifications since the underlying hardware is fully simulated

Para Virtualization – The hardware is not simulated and the guest software run their own isolated domains.

Partial Virtualization – The virtual machine simulates the hardware and becomes independent of it. The guest operating system may require modifications.

1.6.1.3 Software Virtualization

- The ability to computer to run and create one or more virtual environments.
- It is used to enable a computer system in order to allow a guest OS to run.
- Ex. Linux to run as a guest that is natively running a Microsoft Windows OS

Subtypes:

Operating System Virtualization – Hosting multiple OS on the native OS

Application Virtualization – Hosting individual applications in a virtual environment separate from the native OS

Service Virtualization – Hosting specific processes and services related to a particular application

1.6.2 Network Virtualization

- It refers to the management and monitoring of a computer network as a single managerial entity from a single software-based administrator's console.
- Multiple sub-networks can be created on the same physical network, which may or may not is authorized to communicate with each other.
- It allows network optimization of data transfer rates, scalability, reliability, flexibility, and security
- Subtypes:
- Internal network: Enables a single system to function like a network.

1.6.3 Storage Virtualization

- Multiple physical storage devices are grouped together, which look like a single storage device.
- Ex. Partitioning your hard drive into multiple partitions
- Advantages
- Improved storage management in a heterogeneous IT environment
- Easy updates, better availability
- Reduced downtime
- Better storage utilization
- Automated management
- Two types
- Block- Multiple storage devices are consolidated into one
- File- Storage system grants access to files that are stored over multiple hosts

1.6.4 Memory Virtualization

- The way to decouple memory from the server to provide a shared, distributed or networked function.
- It enhances performance by providing greater memory capacity without any addition to the main memory.
- Implementations
- Application-level integration – Applications access the memory pool directly
- Operating System Level Integration – Access to the memory pool is provided through an operating system.

1.6.5 Data Virtualization

Without any technical details, you can easily manipulate data and know how it is formatted or where it is physically located.

It decreases the data errors and workload

The data is presented as an abstract layer completely independent of data structure and database systems

1.6.6 Desktop Virtualization

The user's desktop is stored on a remote server, allowing the user to access his/her desktop from any device or location.

It provides the work convenience and security

It provides a lot of flexibility for employees to work from home or on the go

Since the data transfer takes place over secure protocols, any risk of data theft is minimized

Which Technology to use?

Virtualization is possible through a wide range of Technologies which are available to use and are also Open Source.

They are,

- XEN
- KVM
- OpenVZ

1.7 Parallelization in cloud computing

- Parallel computing is a type of computing architecture in which several processors execute or process an application or computation simultaneously.
- Parallel computing helps in performing large computations by dividing the workload between more than one processor, all of which work through the computation at the same time.

- Most supercomputers implemented parallel computing principles to operate.
- Parallel computing is also known as parallel processing.
- Parallel processing is generally implemented in operational environments/scenarios that require massive computation or processing power.
- The primary objective of parallel computing is to increase the available computation power for faster application processing.
- Typically, parallel computing infrastructure is housed within a single facility where many processors are installed in a server rack or separate servers are connected together.
- The application server sends a processing request that is distributed in small components, which are concurrently executed on each processor/server.
- Parallel computation can be classified as bit-level, instructional level, data and task parallelism.

Cloud Resource Management

- Critical function of any man-made system.
- It affects the three basic criteria for the evaluation of a system:
- Functionality.
- Performance.
- Cost.
- Scheduling in a computing system deciding how to allocate resources of a system, such as CPU cycles, memory, secondary storage space, I/O and network bandwidth, between users and tasks.
- Policies and mechanisms for resource allocation.
- Policy: principles guiding decisions.
- Mechanisms: the means to implement policies
- Cloud resources
- Requires complex policies and decisions for multi-objective optimization.
- It is challenging - the complexity of the system makes it impossible to have accurate global state information.
- Affected by unpredictable interactions with the environment, e.g.,system failures, attacks.
- Cloud service providers are faced with large fluctuating loads which challenge the claim of cloud elasticity

- The strategies for resource management for IaaS, PaaS, and SaaS are different.

Cloud resource management (CRM) policies

- Admission control: prevent the system from accepting workload in violation of high-level system policies.
- Capacity allocation: allocate resources for individual activations of a service.
- Load balancing: distribute the workload evenly among the servers.
- Energy optimization: minimization of energy consumption
- Quality of service (QoS) guarantees: ability to satisfy timing or other conditions specified by a Service Level Agreement

1.7 Dynamic resource allocation

- Cloud Computing environment can supply of computing resources on the basis of demand and when needed
- Managing the customer demand creates the challenges of on-demand resource allocation.
- Effective and dynamic utilization of the resources in cloud can help to balance the load and avoid situations like slow run of systems.
- Cloud computing allows business outcomes to scale up and down their resources based on needs.
- Virtual Machines are allocated to the user based on their job in order to reduce the number of physical servers in the cloud environment
- If the VM is available then job is allowed to run on the VM.
- If the VM is not available then the algorithm finds a low priority job taking into account the job's lease type.
- The low priority job is paused its execution by pre-empting its resource.
- The high priority job is allowed to run on the resources pre-empted from the low priority.
- When any other job running on VMs are completed, the job which was paused early can be resumed if the lease type of the job is suspendable.
- If not, the suspended job has to wait for the completion of high priority job running in its resources, so that it can be resumed.
- There are three types
- Cancellable: These requests can be scheduled at any time after their arrival time
- Suspendable: Suspendable leases are flexible in start time and can be scheduled at any time after their ready time

- Non-Preemptable: The leases associated with such requests cannot be pre-empted at all.

1.8 Optimal allocation of cloud models

The optimal allocation of computing resources is a core part for implementing cloud computing.

High heterogeneity, high dynamism, and virtualization make the optimal allocation problem more complex than the traditional scheduling problems in grid system or cloud computing system.

SATHYABAMA
INSTITUTE OF SCIENCE AND TECHNOLOGY
(DEEMED TO BE UNIVERSITY)
Accredited "A" Grade by NAAC | 12B Status by UGC | Approved by AICTE
www.sathyabama.ac.in

SCHOOL OF ELECTRICAL AND ELECRONICS ENGINEERING

DEPARTMENT OF ELECTRONICS AND COMMUNICATION ENGINEERING

UNIT – II – Cloud computing – SECA5303

UNIT 2

CLOUD COMPUTING SERVICES

Cloud provider, SAAS, PAAS, IAAS and other Organizational scenarios of clouds.

2.1 Types of Services provided by Cloud

- Software as a Service (SaaS)
- Infrastructure as a Service (IaaS)
- Platform as a Service (PaaS)

Service Oriented Architecture

- Elastic Computing
- On Demand Computing

2.2 Cloud Services

Software as a Service

Introduction

SaaS (software-as-a-service). WAN-enabled application services (e.g., Google Apps, Salesforce.com, WebEx).This is a public cloud service model where the application is 100% managed by the cloud provider. SaaS removes the need for organizations to install and run applications on their own computers or in their own data centers.This eliminates the expense of hardware acquisition, provisioning and maintenance, as well as software licensing, installation and support.Software-as-a-Service (SaaS) has evolved from limited on-line software delivery in 1990s to a fully matured “direct-sourcing” business model for enterprise applications.

SaaS is one of the fastest growing concepts: more than 10 million companies will be using SaaS in the next 5 - 10 years; more than 50% of all Fortune 500 companies are already using SaaS.According to influential IT institutes, SaaS is the leading business model of choice for 2008/2009Virtually all big software/service vendors (IBM, Microsoft, Oracle, Cisco) are investing heavily in SaaSWith the continuously increasing bandwidth and reliability of the internet, using web services over the (public) internet has become a viable option.Microsoft Office 365 is available with the Azure cloud platform.The architecture uses an application instance instead of server instances.There is no actual migration of company servers to the cloud.

The SaaS model provides single-tenant and multiple tenant services.The single-tenant dedicates the application instance to the assigned tenant.The multiple tenant application is shared by multiple tenants.The company can manage the security and storage with the single-

tenant model.The SaaS application is well suited to internet connectivity.The employees along with their partners and customers can access the application with a variety of network access devices.The SaaS billing model is based on either per usage or monthly subscription.The security compliance requirements for some applications prevented deployment to the SaaS cloud.

Some SaaS providers offer Virtual Private SaaS (VPS) services for additional application security.It is a hybrid deployment model that allows peering with an enterprise or VPC database server.The peering is for storage purposes only and used for security compliance. Salesforce.com is a leading SaaS provider with a CRM application to customers.

2.2.1Benefits of SaaS

- Flexible payments
- Scalable usage
- Automatic updates
- Accessibility and persistence
- On demand computing
- Opportunities of SaaS

Software provided as a service by a software vendor to multiple customers with the following main characteristics:

- Standardization of software
- Service including maintenance, support and upgrades
- Web based – usage over the (public) internet
- SaaS offers potential for lowering the Total Cost of Ownership
- Lower operational costs
- No large scale, costly, high risk implementations of applications
- Need few operational resources for application management
- No platform and hardware (maintenance) costs for application servers
- Reduced operational complexity: software delivered as a transparent service through the web
- Minimized software development costs – No lengthy software development and testing cycles
- Lower costs for software use
- No software license and annual maintenance fees
- No expensive software upgrades
- Lower application consultancy and support costs

- SaaS allows corporations to focus on core business activities and responsibilities
- Transparent overview and usage of electronic data and information
- Automation of iterative, manual tasks
- Faster Time to Market – easy to scale software
- More flexibility in changing and modifying application services for business needs – Full-scale integration of business processes
- Control over IT
- Minimized IT Service Management efforts mainly focused on availability –
- Well-defined SLAs between the corporation and the IT vendor
- More predictable cash flow – easier licensing based on access/usage of software
- Increased productivity and improved user satisfaction
- Automatic software upgrades with minimal outage

2.2.2 Limitations

- Businesses must rely on outside vendors to provide the software, keep that software up and running, track and report accurate billing and facilitate a secure environment for the business' data.

2.3Platform as a Service

The Platform as a Service (PaaS) is a way to rent hardware, operating systems, storage and network capacity over the Internet.

PaaS services are,

- Data services
- Application runtime
- Messaging & queueing
- Application management.
- The PaaS is a computing platform that abstracts the infrastructure, OS, and middle-ware to drive developer productivity.
- The PaaS is foundational elements to develop new applications
- E.g., Google Application Engine, Microsoft Azure, Coghead.
- Microsoft Azure
- Pay per role instance
- Add and remove instances based on demand
- Elastic computing!
- Load balancing is part of the Azure fabric and automatically allocated.
- The PaaS is the delivery of a computing platform and solution stack as a service

- The Solution stack is integrated set of software that provides everything a developer needs to build an application for both software development and runtime.

2.3.1 PaaS offers the following

Facilities for application design

- Application development
- Application testing, deployment
- Application services are,
- Operating system
- Server-side scripting environment
- Database management system
- Server Software
- Support
- Storage
- Network access
- Tools for design and development
- Hosting

All these services may be provisioned as an integrated solution over the web

2.3.2 Properties and characteristics of PaaS

- Scalability
- Availability
- Manageability
- Performance
- Accessibility

2.3.3 PaaS Features

- It delivers the computing platform as service
- The capacities to abstract and control all the underlying resources
- It helps to providers any smallest unit of resources
- To provide a reliable environment for running applications and services
- Act as a bridge between consumer and hardware
- Do not need to care about how to build, configure, manage and maintain the backend environment
- It provides a development and testing platform for running developed applications
- Reduce the responsibility of managing the development and runtime environment

2.3.4 Advantages of PaaS

It helps to provide deployment of application without the cost and complexity of buying and managing the hardware and software

It provides all the required to support the complete life cycle of building and delivering web applications and services entirely available from the internet

2.3.5Disadvantages of PaaS

- Less flexible than IaaS
- Dependency on provider
- Adoption of software / system architecture required
- Evolving from different standard.
- Evolving “upwards” from IaaS
- Amazon (Mail, Notification, Events, Databases, Workflow, etc.)
- Evolving “downwards” from SaaS
- Force.com – a place to host additional per-tenant logic.
- Google App Engine
- Evolving “sideways” from middleware platforms
- WSO2, Tibco, vmWare, Oracle, IBM
- Generic PaaS Model
- Infrastructure as a Service
- This service offers the computing architecture and infrastructure i.e. all types of computing resources

All resources are offered in a virtual environment, so that multiple users can access it.

The resources are including,

Data storage

- Virtualization
- Servers
- Networking
- The vendors are responsible for managing all the computing resources which they provided.
- It allows existing applications to be run on a supplier’s hardware.
- User Task in IaaS Cloud
- Multiple user can access Virtual instances

The user responsible for handling other resources such as,

2.3.6 Applications

- Data

- Runtime
- Middleware

Example IaaS service providers

- AWS EC2 / S3 / RD
- GoGrid
- RackSpace

2.3.7 Pros

- The cloud provides the infrastructure with enhanced scalability i.e. dynamic workloads are supported
- It is flexible
- Cons
- Security issues
- Network and service delay

2.3.8 Comparison of cloud services

- Blue indicates the levels owned and operated by the organization / Customer
- White levels are run and operated by the service provider / Operator
- Cloud Computing Services Pros
- Lower computer costs
- Improved performance:
- Reduced software costs
- Instant software updates
- Improved document format compatibility
- Unlimited storage capacity
- Increased data reliability
- Universal document access
- Latest version availability
- Easier group collaboration
- Device independence

2.3.9Cons

- Requires a constant Internet connection
- Does not work well with low-speed connections
- Features might be limited
- Can be slow
- Stored data can be lost

- Stored data might not be secure
- Service Oriented Architecture

2.4Service

A service is a program you interact with via message exchanges A system is a set of deployed services cooperating in a given task.

2.5Architecture

It serves as the blueprint for the system

Team structure

- Documentation organization
- Work breakdown structure
- Scheduling, planning, budgeting
- Unit testing, integration
- Architecture establishes the communication and coordination mechanisms among components

2.5.1 Software Architecture

- It is collection of the fundamental decisions about a software product/solution designed to meet the project's quality attributes (i.e. requirements).
- The architecture includes the main components, their main attributes, and their collaboration (i.e. interactions and behavior) to meet the quality attributes.
- Architecture can and usually should be expressed in several levels of abstraction (depending on the project's size).
- Architecture is communicated from multiple viewpoints

2.5.2 Service Oriented Architecture (SOA)

- It is a design pattern or software architecture which provides application functionality as a service to other applications.
- The basic principles of service-oriented architecture are independent of vendors, products and technologies.
- The services are provided to the other components through a communication protocol over a network.
- Every service has its own business logic
- Consumer interface layer – this layer is used by the customer
- Business process layer – it provides the business process flow
- Service layer – this layer comprises of all the services in the enterprises
- Component layer – this layer has the actual service to be provided

- Operational system layer – this layer contains the data model

SOA – Architecture in details

2.5.3 Principles of SOA

- Service loose coupling – service does not have high dependency
- implementation from outside world
- Service reusability – services can be used again and again instead of rewriting them
- Service statelessness – they usually do not maintain the state to reduce the resource consumption
- Service discoverability – services are registered in registry, so that the client can discover them in the service registry.
- Applications
- Manufacturing – E.g. Inventory management
- Insurance – Take up the insurance of the employees in companies

Companies using SOA

- Banking Sector
- ICICI Bank
- HDFC Bank
- UTI Bank etc..
- Manufacturing Sector
- Apollo Tyres
- Maruthi
- Hyundai

2.5.4 Advantages

- Interoperability
- Programs to run different vendors / locations
- To interact with different networks
- Different operating systems
- Solution: XML
- Scalability
- To extend the processing power of the servers
- Reusability
- If any new systems are introduced, no need to create a new service for every time.
- Parallel application development
- Modular approach

- Easy maintenance
- Greater Reliability
- Improved Software Quality
- Platform Independence
- Increased Productivity

2.5.5 Disadvantages

- Stand alone, non-distributed applications
- Homogenous application environments
- GUI based applications
- Short lived applications
- Real time applications
- One-way asynchronous communication applications

SATHYABAMA
INSTITUTE OF SCIENCE AND TECHNOLOGY
(DEEMED TO BE UNIVERSITY)
Accredited "A" Grade by NAAC | 12B Status by UGC | Approved by AICTE
www.sathyabama.ac.in

SCHOOL OF ELECTRICAL AND ELECRONICS ENGINEERING

DEPARTMENT OF ELECTRONICS AND COMMUNICATION ENGINEERING

UNIT – III – Cloud Computing– SECA5303

UNIT 3

CLOUD ADMINISTRATION AND MANAGEMENT

Administering & Monitoring cloud services, benefits and limitations,

3.1Difference between a data center and cloud computing

The main difference between a cloud and a data center is that a cloud is an off-premise form of computing that stores data on the Internet, whereas a data center refers to on-premise hardware that stores data within an organization's local network.

3.1.1Location of Data Storage in Cloud

Cloud storage is a model of data storage in which the digital data is stored in logical pools, the physical storage spans multiple servers (and often locations), and the physical environment is typically owned and managed by a hosting company.

3.1.2Host in Data Centre

Data center hosting is the process of deploying and hosting a data center on a third-party or external service provider's infrastructure. It enables the use of the same services, features and capabilities of a data center but from a hosted platform external to the on-premises data center or IT infrastructure.

3.1.3Point of a data Centre

A data center is a facility that centralizes an organization's IT operations and equipment, as well as where it stores, manages, and disseminates its data. Data centers house a network's most critical systems and are vital to the continuity of daily operations.

The term ―data center‖ can be interpreted in a few different ways. First, an organization can run an in-house data center maintained by trained IT employees whose job it is to keep the system up and running. Second, it can refer to an offsite storage center that consists of servers and other equipment needed to keep the stored data accessible both virtually and physically.

3.1.4Pros: Data centers come with a number of pros. Organizations able to have an in-house data storage center are far less reliant on maintaining an Internet connection. Data will be accessible as long as the local network remains stable. Remote storage has its advantages as well. If the organization’s location is compromised via fire, break-in, flooding, etc., the data will remain untouched and unharmed at its remote location.

3.1.5Cons: Having all or most of your data stored in one location makes it more easily accessible to those you don’t want having access, both virtually and physically. Depending on your organization’s budget, it could prove too expensive to maintain an organization-owned and operated data center.

3.2The Basics

The main difference between a cloud and a data center is that a cloud is an off-premise form of computing that stores data on the Internet, whereas a data center refers to on-premise hardware that stores data within an organization's local network. While cloud services are outsourced to third-party cloud providers who perform all updates and ongoing maintenance, data centers are typically run by

an in-house IT department.

Although both types of computing systems can store data, as a physical unit, only a data center can store servers and other equipment. As such, cloud service providers use data centers to house cloud services and cloud-based resources. For cloud-hosting purposes, vendors also often own multiple data centers in several geographic locations to safeguard data availability during outages and other data center failures.

For companies considering whether or not to use cloud computing versus staying with or building their own data center, there are three primary factors affecting their decision: their business needs, data security and system costs.

Does your business need a cloud or a data center?

A data center is ideal for companies that need a customized, dedicated system that gives them full control over their data and equipment. Since only the company will be using the infrastructure's power, a data center is also more suitable for organizations that run many different types of applications and complex workloads. A data center, however, has limited capacity -- once you build a data center, you will not be able to change the amount of storage and workload it can withstand without purchasing and installing more equipment.

On the other hand, a cloud system is scalable to your business needs. It has potentially unlimited capacity,based on your vendor's offerings and service plans. One disadvantage of the cloud is that you will not have as much control as you would a data center, since a third party is managing the system. Furthermore, unless you have a private cloud within the company network, you will be sharing resources with other cloud users in your provider's public cloud.

3.2.1Cloud security vs. data center security

Because the cloud is an external form of computing, it may be less secure or take more work to secure than a data center. Unlike data centers, where you are responsible for your own security, you will be entrusting your data to a third-party provider that may or may not have the most up-to-date security certifications. If your cloud resides on several data centers in different locations, each location will also need the proper security measures.

A data center is also physically connected to a local network, which makes it easier to ensure that only those with company-approved credentials and equipment can access stored apps and information. The cloud, however, is accessible by anyone with the proper credentials anywhere that there is an Internet connection. This opens a wide array of entry and exit points, all of which need to be protected to make sure that data transmitted to and from these points are secure.

3.2.2Cloud vs. data center costs

For most small businesses, the cloud is a more cost-effective option than a data center. Because you will be building an infrastructure from the ground up and will be responsible for your own maintenance and administration, a data center takes much longer to get started and can cost businesses

$10 million to $25 million per year to operate.

Unlike a data center, cloud computing does not require time or capital to get up and running. Instead, most cloud providers offer a range of affordable subscription plans to meet your budget and scale the service to your performance needs. Whereas data centers take time to build, depending on your provider, cloud services are available for use almost immediately after registration.

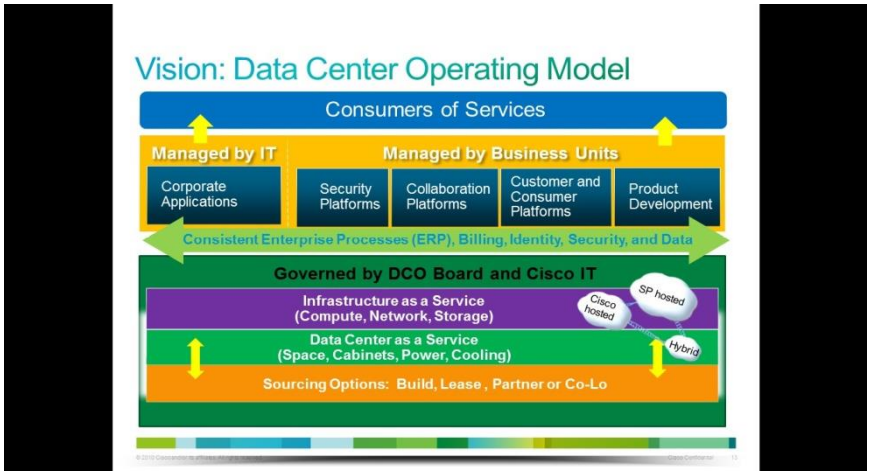

Figure 3.1 Data Centre Operating Model

Data Center Components (Preview)

Basic Data Center facility systems:

- Physical space
- Raised flooring
- In-room electrical
- Standby power
- Data cabling
- Cooling
- Fire suppression

3.2.3 Physical Space

Physical space refers to the footprint that Data Center-related items occupy. This generally applies to the overall area of the Data Center and its associated spaces, such as electrical rooms or storage areas. On a smaller scale this might refer to key dimensions within the Data Center, such as the external measurements of a server cabinet or aisle clearances.

3.2.4Raised Flooring

Raised flooring is an elevated grid system that is frequently installed in large Data Centers. Cooled air, electrical whips, and data cabling are routed through the space under the raised floor, promoting better air flow and enabling easier management of power and cable runs. Water pipes, fire suppressant cylinders, moisture detectors, and smoke detectors may be located here as well.Raised flooring can

vary in height from a few inches to several feet, or a few centimeters to several meters. In extreme cases they are as tall as the story of a building, enabling workers to walk upright under the plenum. Regardless of their height, the floors are typically composed of standard 2 foot (60 centimeter) square floor tiles. The tiles can vary in weight, strength, and finish depending upon their use. Tiles featuring either small perforations or large cut-out sections are placed in key locations to enable pass-through of air and cabling between the areas above and below the floor.

3.2.5 In-Room Electrical

In-room electrical refers to all power-related facilities within the Data Center. This normally includes electrical panels, conduits, and several types of receptacles. Power to this system usually comes from an outside commercial power source, namely your local utility company, and is likely conditioned at the company site. Voltage varies from one country to another.

3.2.6Standby Power

Standby power includes all backup power systems responsible for support of the Data Center's electrical load in the event that normal utility power fails for any reason. This system traditionally includes large batteries, known as an uninterruptible power source or uninterruptible power supply, and one or more generators.

3.2.7 Cabling

The cabling system is all structured cabling within the Data Center. Copper and fiber cabling are the typical media and are terminated via several types of connectors. Common components include fiber housings, patch panels, multimedia boxes, and data faceplates. Cabinets, raceways, and other items used to route structured cabling are also considered part of the cabling system. Users plug servers in to the Data Center's structured cabling system with pre-terminated patch cords.

3.2.8 Cooling

The cooling system refers to the chillers and air handlers used to regulate ambient temperature and control humidity within the Data Center. This system might incorporate the air conditioning system used to cool regular office space within the same building, known as house air, or might be independent of it. Individual server cabinets can also possess their own cooling measures, such as fans or water-cooling.

Fire Suppression

Fire suppression includes all devices associated with detecting or extinguishing a fire in the Data Center. The most obvious components are water-based sprinklers, gaseous fire suppression systems, and hand-held fire extinguishers. Others can include devices that detect smoke or measure air quality.

3.2.9 Other Infrastructure Components

There are also some infrastructure items that do not strictly fall under the prior categories but are commonly found in server environments. These include leak detection devices, seismic

mitigation, and physical security controls such as card readers and security cameras.

3.3 Data Center Design Criteria

How many layers of infrastructure should your Data Center possess?

Will it be the only server environment for your company or one of several?

Will the room house production servers and be a business-critical site or contain a minimum of equipment for disaster recovery purposes and serve as a failover location?

How long is its initial construction expected to meet your company's needs?

What is it all going to cost? Availability

3.4Availability:

The degree to which Data Center devices function continuously is known as the room's availability or its uptime.

Availability is represented as a percentage of time. How many days, hours, and minutes is the Data Center's electrical infrastructure operational and supplying power over a given time period

Infrastructure Tiers

The higher the availability you want your Data Center to achieve, the more layers of infrastructure it must have.

N capacity is the amount of infrastructure required to support all servers or networking devices in the Data Center, assuming that the space is filled to maximum capacity and all devices are functioning.

N most commonly used when discussing standby power, cooling, and the room's network.

N+1 infrastructure can support the Data Center at full server capacity and includes an additional component

Alternately called a 2N or system-plus-system design, it involves fully doubling the required number of infrastructure components

Even higher tiers exist or can be created: 3N, 4N, and so on.

One Room or Several?

One large Data Center is simpler to manage than several smaller ones. Having only one server environment puts all of your eggs in one basket.

3.4.1Life Span

How long it is expected to support your company's needs without having to be expanded or retrofitted, or otherwise undergo major changes.

The most effective strategy is to design a Data Center with a projected life span of a few years.

3.4.2 Budget Decisions

It is no good to spend millions of dollars on a server environment to protect your company's assets if that cost drives your business into bankruptcy.

The most obvious costs for a Data Center are labor and materials associated with its initial construction, which, even for a room smaller than 1000 square feet or 100 square meters, normally runs into hundreds of thousands of dollars. This includes:

Initial construction Consulting fees Real estate

Ongoing operational expenses

3.4.3Mobile Cloud Computing Service Models:

Mobile computing is using a computer while on the move.

Mobile devices allow users to run powerful applications that take advantage of the growing availability of built-in sensing and better data exchange capabilities of mobile devices.

As a result, mobile applications seamlessly integrate with real-time data streams and Web 2.0 applications, such as mashups, open collaboration, social networking and mobile commerce.

The mobile execution platform is being used for more and more tasks, e.g., for playing games; capturing, editing, annotating and uploading video; handling finances; managing personal health, micro payments, ticket purchase, interacting with ubiquitous computing infrastructures.

Even mobile device hardware and mobile networks continue to evolve and to improve, mobile devices will always be resource-poor, less secure, with unstable connectivity, and with less energy since they are powered by battery.

Resource poverty is major obstacle for many applications. Therefore, computation on mobile devices will always involve a compromise.

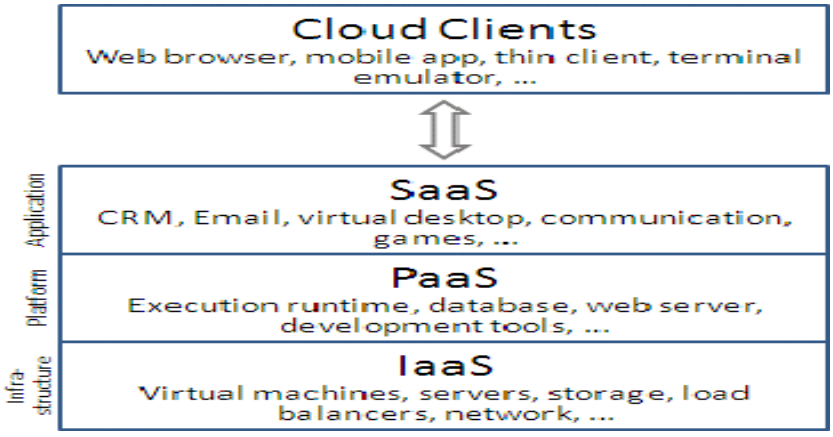

Figure 3.2 Types of Cloud Services

3.4.5 Mobile Cloud Computing Service Models

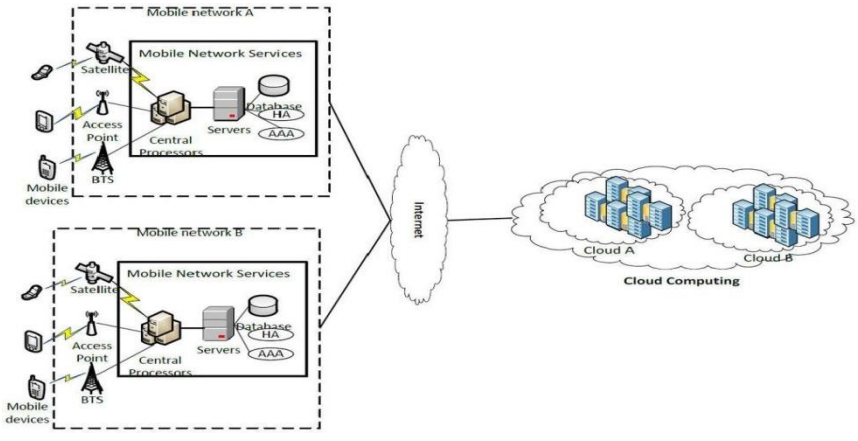

Figure 3.3 Mobile Cloud Computing Service Model

3.4.6 Mobile Cloud Computing Architecture

Mobile Cloud Computing (MCC) as a new technology of increasing interest from industry, academia and government has extended infrastructure, platform, and software services offered in a cloud to mobile users.

Mobile Cloud Computing (MCC) is the combination of cloud computing, mobile computing and wireless networks to bring rich computational resources to mobile users, network operators, as well as cloud computing providers.

The ultimate goal of MCC is to enable execution of rich mobile applications on a plethora of mobile devices, with a rich user experience.

MCC provides business opportunities for mobile network operators as well as cloud providers.

More comprehensively, MCC can be defined as a rich mobile computing technology that leverages unified elastic resources of varied clouds and network technologies toward unrestricted functionality, storage, and mobility to serve a multitude of mobile devices anywhere, anytime through the channel of Ethernet or Internet regardless of heterogeneous environments and platforms based on the pay-as-you-use principle.

3.5Challenges

Collaborating on Calendars, Schedules and Task Management Exploring Online Scheduling Applications

Exploring Online Planning and Task Management Collaborating on Event Management Collaborating on Contact Management Collaborating on Project Management Collaborating on Word Processing Collaborating on Databases Storing and Sharing Files

3.5.1 COLLABORATING ON CALENDARS, SCHEDULES AND TASK MANAGEMENT

Comptuer users have a practice of maintaining their schedules, events and appointments on PCs or Laptops. But those days have gone, now it can be easily maintained on the Web. Though Microsoft Outlook or Windows Calendar provides advanced features, it could be maintained on a single computer only and could not be tracked while you are traveling or away from your home or office PC or laptop. Therefore, computer users have starting using Web-based Calendars so that it can be accessed from anywhere even while traveling.

It helps to share easily with other users on any location and leads to do collaborative projects more efficiently.

Some of the Web-Based Calendars are given below:

3.5.2 Google Calendar

The Google Calendar is the most popular Web-based Calendar and associate with its search engine (calendar.google.com). It is free and easy to use. It helps to create a personal or shared

calendar or both which enables it to maintain schedules, events, and appointment of business, family, friends and community. The appointments can be directly entered on the Calendar. It displays in many views such as daily, weekly,
or monthly on a single page

Figure 3.4 Snapshot of Google Calendar

What types of calendars can you create with Google Calendar?

Personal calendars - like your default calendar

Public calendars - which others can access via the web

Friends' calendars - which you import from their Google Calendar web pages

Holiday calendars - which add national holidays to a basic calendar

Yahoo! Calendar

The primary competitor of Google Calendar is Yahoo Calendar. It looks alike and functions like a Google Calendar. It is hosted by Yahoo.

The advantages of Yahoo Calendar are:

Web-based

Free

Presence of Add task button

Share with others users in a collaborative environment

3.5.3Disadvantage:

1. Yahoo! Calendar allows you to create a single calendar only. All events, public and private, should be stored on this calendar only. You can't have multiple calendars for various functions.

One of Google Calendar's primary competitors is Yahoo! Calendar (calendar.yahoo.com), hosted by its search competitor Yahoo! This web-based calendar looks, feels, and functions quite similarly to Google Calendar, and is also free for anyone to use.

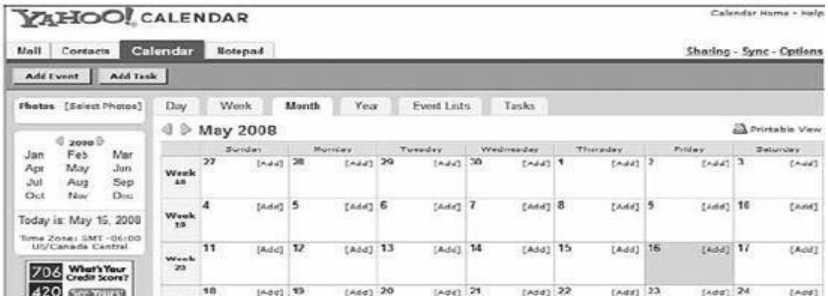

Figure 3.5 Snapshot of Yahoo Calendar

3.6Windows Live Calendar

The third largest search site offers a competitive Windows Live Calendar. It looks like both Google and Yahoo Calendar and also allows you to schedule meetings with other calendar and Hotmail users. It is possible to share calendar with authorized users for group collaboration.

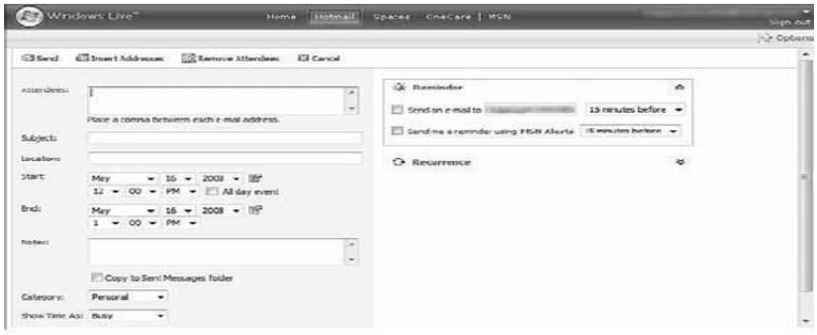

Figure 3.6 Snapshot of Windows Live Calender

3.6.1Apple MobileMe Calendar

Apple's MobileMe (www.me.com) is a new competitor in the web-based apps market. It includes online mail, contacts, and calendar, as well as an online photo gallery and file storage. The MobileMe Calendar is, of course, a web-based calendar that can beaccessed from any computer connected to the Internet, Mac or Windows. What makes it more unique and potentially more useful is that it can also be accessed from Apple's iPhone, which makes it a truly mobile calendar.

Other Calendars

Some more calendars are listed below:

Hunt Calendars (www.huntcal.com) –Useful features include email reminders lets you add web links and images to calendar events, which is fairly unique.

Famundo (www.famundo.com) - If you keep the schedule for a community group, check out a free web based calendar ideal for schools, churches, sports teams, and the like

eStudio Calendar (www.same-page.com/calendar-software.html) – designed specifically for business use. You get three types of calendars in a single interface:

Member Event calendar helps users manage their personal time, keep track of meetings with others, and so on.

Team Event calendar is used to schedule activities for a group, as wellas schedule facilities.

Supervisorcalendar provides reports to managers about business activities and schedules.

30Boxes (www.30boxes.com) - 30 ―boxes displayed on a typical monthly calendar

offers a slick interface for adding events

Trumba (www.trumba.com) offers web-based calendars ideal for community organizations, schools, and similar public entities. The company lets you embed individualized widgets your own website.

Calendars Net (www.calendars.net) - a free web-based calendar designed for companies or individuals who want to add interactive calendars to their websites.

Jotlet (www.jotlet.net) - a JavaScript API and library that you can use to build rich calendar functionality into any web page.

3.7 EXPLORING ONLINE SCHEDULING APPLICATIONS

As anyone in a large office knows, scheduling a meeting can be a frustrating experience. Online Scheduling web-based app takes much of the pain out of scheduling meetings, for both large and small groups. The typical app requires all users to enter their individual calendars beforehand. When you schedule a meeting, the app checks attendees' schedules for the first available free time for all. The app then generates automated email messages to inform attendees of the meeting request (and the designated time), followed by automatic confirmation emails when attendees accept the invitation. Professionals who schedule appointments with their clients—doctors, lawyers, hairdressers, and the like— face similar scheduling challenges. For this purpose, separate web-based appointment scheduling applications exist. These apps function similarly to traditional meeting schedulers, but with a focus on customer appointments.

3.7.1 Jiffle

Web-based solutions for meeting scheduling. Jiffle (www.jifflenow.com),

Schedules meetings, appointments, etc. for the enterprise environment.

To track employees' free time, it synchronizes seamlessly with both Microsoft Outlook and Google Calendar.

It also offers its own Jiffle Calendar application.

Jiffle allows the originating user to mark available time slots on his calendar, and then share them with proposed attendees via a Jiffle-generated email invitation. These attendees view the invitation, log in to the Jiffle website, and then select their preferred time slots from the ones proposed.Based on these responses, Jiffle picks the best time for the meeting and notifies all attendees via an automatic confirmation email.

For smaller companies, Jiffle is free for up to 10 meeting confirmations per month.

For larger companies, Jiffle Plus, Jiffle Pro, and Jiffle Corporate plans are available.

Presdo

- Jiffle, Presdo (www.presdo.com)
- a scheduling tool that isn't limited to a single company.
- lets you schedule meetings and events with anyone who has an email address.
- adding an event is as simple as entering a description into a box. You then enter the email addresses of other participants, and Presdo emails out the appropriate invites.
- When an attendee responds, he's automatically added to the event's guest list.

3.7.2 Diarised

- Diarised (www.diarised.com) - like Presdo
- a web-based meeting maker that users across different companies can use.
- It helps you pick the best time for a meeting by sending out emails to invitees, letting them choose the best times for them, and then sending you a summary of those best dates.
- You pick the final date, Diarised notifies everyone via email, and your meeting is scheduled.

3.7.3 Windows Live Events

- Microsoft's Windows Live Events (home.services.spaces.live.com/events/)
- a customized version of its Live Spaces offering; it lets Live Spaces users organize events and share activities between participants.
- To schedule an event, you set up a list of invitees and then send out a mass
- email with a link back to your Live Event site.
- Information about the event is posted on the site itself, which also serves as a place for attendees to come back after the event and share their photos, videos, and blog posts about the event.

3.7.4 Schedulebook

- Schedulebook (www.schedulebook.com)
- Offers several different types of webbased scheduling services. Depending on the application, you can use Schedulebook to schedule employees, customers, or other interested parties.
- The company's three offerings are
- Schedulebook Professionals, which is a business-oriented schedule/calendar/planning application
- Schedulebook Office, which schedules the use of any shared resource, such as company meeting rooms or even vacation homes
- Schedulebook Aviation, which is used by the aviation industry to schedule aircraft, flight training, and similar services

3.7.5 Appointment Quest

- Appointment Quest (www.appointmentquest.com)
- Designed to solve the scheduling problems of busy professionals.
- This application not only enables clients to make and you to accept appointments over the web, it also lets you manage personnel, schedules, and other calendar related items.

3.7.6 HitAppoint

- hitAppoint (www.hitappoint.com) enables online client booking.
- ideal for any business that requires the making of customer appointments—barbershops,hair salons, doctor and dentist offices, consultants, financial advisors, car repair shops, computer technicians, and the like.

3.8 EXPLORING ONLINE PLANNING AND TASK MANAGEMENT

Now let's pivot from schedules to tasks. Planning and task applications let you manage everything from simple to-do lists to complex group tasks, all over the Internet and collaboratively with other users.

3.8.1 iPrioritize

- Sharing to-do lists is important for families, community groups, and businesses.
- Your to-do list might be as simple as a grocery list or as complex as a list of activities for a community program or business project.
- iPrioritize (www.iprioritize.com) - a good basic to-do list manager.
- authorized users can create a new to-do list, add items to the list, prioritize tasks by dragging them up and
- down thelist, and mark items complete when finished.
- Web-based - so you can access your lists anytime and anyplace.

When you have a list, you can print it out, email it to someone else, subscribe to changes in the list via RSS, and even view lists on your mobile phone— which is a great way to consult your grocery list when at the supermarket!

Bla-Bla List

Bla-Bla List (www.blablalist.com) - another simple to-do list manager. Web-based - so you can access your lists from any location at any time.

You can even publish your lists via RSS so that family and coworkers can get instant updates.

3.8.2 Hiveminder

- Hiveminder (www.hiveminder.com)
- you can enter list items in a kind of freeform fashion, and it will help you create and prioritize lists based on your ―brain dumps.‖
- Remember the Milk
- When you need to ―remember the milk‖ at the grocery store, check out the aptly named Remember the Milk (www.rememberthemilk.com)

- web-based todo list manager.
- Once you create a list, you can arrange reminders via email, instant messaging, or text messages to your mobile phone.

3.8.3 Ta-da List

- Web-based to-do list manager.
- Ta-da List (www.tadalist.com)
- Lets you make all sorts of lists, share them with friends, family, and coworkers,and then check off items as they're completed.

3.8.4 Tudu List

- Tudu List (www.tudulist.com) a little different from other to-do list managers
- It also includes a web-based calendar.
- Items are added both to the appropriate to-do list and to your calendar, on the date they're due.
- TaskTHIS
- TaskTHIS (taskthis.darthapo.com)
- Offers the ability to add extended notes to any individual task.
- You can publish your tasks via RSS or share with others via the web.

3.8.5 Vitalist

- Vitalist (www.vitalist.com) - to-do list manager
- organizes all sorts of tasks and projects.
- It's unique in that it uses the Getting Things Done (GTD)
- workflow methodology popularized by management consultant David Allen.

3.8.6 TracksLife

- Trackslife (www.trackslife.com) - a database-oriented task manager.
- Each ―track‖ is a separate database that combines columns of money, numbers,words, paragraphs, and yes/no responses.
- The application sends out reminders of critical events via email or RSS.

3.8.7 Others

- Voo2Do (www.voo2do.com- more sophisticated priority management.
- Web-based application lets you set up different projects, organize tasks by project, track time spent and remaining on a given task or project, publish task lists, and even add tasks via email.

3.8.8 HiTask

Tasks are added to your calendar and color tagged for easy viewing.

The task manager and scheduler both utilize drag-and-drop editing, and you can share and assign tasks and projects to a group of people via the web.

- **Zoho Planner**
- Zoho Planner (planner.zoho.com) - perhaps the most sophisticated task planner
- With Zoho Planner, you create a new page for each project you're working on
- To that project, you add lists with individual to-dos within each list. Each list item can include extensive notes as well as images.
- You can share each project page with users you designate.
- Each todo item also appears on your central calendar.

3.9 . COLLABORATING ON CONTACT MANAGEMENT

Most technically adept people today keep their lists of friends, family, and business contacts in some sort of computer-based address book. Maybe it's in the Windows Address Book, maybe it's in Microsoft Outlook, but it's likely that you have all your contacts in some electronic file someplace on your computeR

1. Understanding Contact Management and CRM

Whether we're presenting simple address book–based applications or sophisticated sales automation programs, it all comes down to how the application uses the information you provide about a person—which is, in essence, contact management.

All about Contact Management

- Contact management is the act of storing information about friends, family, and business colleagues for easy retrieval at a later date.
- Contact management can be more involved and more useful than simple name/address storage.
- More sophisticated contact management applications
- Help you track all sorts of details about your friends and colleagues, from personal info (birth date, spouse's name, children's names, favorite restaurants, and the like) to business info (employer, department, job responsibilities, and so forth). These contact management systems typically integrate this personal/professional information with calendar functions and a task manager

Web-based contact management applications enable you to access your contact information from any computer connected to the Internet. Instead of storing personal contacts on your home PC and work contacts on your office computer, you store all your contacts in the cloud, where they can be accessed from both home and work.

3.9.1 All About CRM

Many businesses require more practical use of their contact information. It's not enough to have the equivalent of a digital Rolodex on hand; that contact information can be injected into various automated processes to help establish and maintain lasting and productive relationships with the

company’s customers.

This process of managing the needs, wants, and buying patterns of customers is referred to as customer relationship management. CRM helps companies understand and anticipate the needs of current and potential customers; it’s an essential tool for building strong customer relationships.

CRM software not only stores customer contact information, it also stores and analyzes all data relating to a particular customer, and then uses that data to help you determine how best to relate with that customer. For example, you can use a CRM program to discover which customers order the most from your company—and then trigger regular phone calls or emails to those customers. Or you can use CRM to find out which customers have the most contact with your technical support department, and then ward off future support calls by proactively sending out support info or scheduling a special support seminar.

3.9.2 **Exploring Contact Management and CRM Applications**

The line between contact management, CRM, and SFA applications is blurry enough to make clear distinctions impossible. To that end, we’ll look at all three types of applications in one long list—starting with the industry-leadingSalesforce.com, and proceeding in alphabetic order from there.

Salesforce.com

The most popular web-based contact management/CRM available today is offered by Salesforce.com (www.salesforce.com). In fact, the company offers several different cloud services:

Salesforce.com, a software-as-a-service CRM application designed for sales, marketing, customer service and other uses

Force.com, a platform-as-a-service application designed for developers who want to design or customizetheir own sales force apps

AppExchange, an online marketplace of add-on software for Salesforce.com, developed by independentcompanies

All these cloud services are buttressed by a robust community and support structure, including blogs, forums, education and training initiatives, and the like. The company’s primary application is the self-named Salesforce.com. The company offers a hosted collection of on-demand business applications that include the following:

Sales Force Automation, which includes activity management, channel and territory management, forecasting, mobile access, email templates,

and real-time analytics that help companies increase sales productivity and grow revenues _ Service & Support, a customer service solution for enterprise call centers

Partners, a partner relationship management application that enables collaboration and partnership with channel partners

Marketing, which includes tools to execute, manage, and analyze the results of multichannel marketing campaigns

Content, which enables companies to share documents and other content across the organization

Ideas, which help a company, build online communities with their customers, partners, and employees Analytics, which offers real-time reporting, calculations, and dashboards to help improve decision making and resource allocation

Connections

3.9.4 Web-based CRM solution

- Connections (www.bconnections.com)
- A contact management program augmented with essential CRM functions for small and medium-sized businesses.
- The connections application starts with a list of companies you do business with, and a list of contacts at those businesses.
- This contact information is hosted on the web and accessible from any Internet-connected computer.
- The application includes a web-based calendar that sales management can use to manage the activities of all their reps.
- It also tracks leads and sales opportunities, to help you better prepare sales forecasts.

3.9.5 Big Contacts

- Big Contacts (www.bigcontacts.com)
- A web-based contact manager designed for workgroups as small as 2 people or as large as 2,000.
- It features an address book, group calendar, task manager, and to-do lists.
- Its CRM functions include sales tracking, activity reports, team management, and mobile access.
- Pricing is on a per-user basis.

3.9.6 eStudio Contact Manager

- For more basic contact management, check out eStudio Contact Manager (www.same-page.com/contact-management.html).
- This is application is an online address book specifically designed for business contacts.
- The address book can be accessed by multiple users from any Internet-connected computer, making it ideal for real-time contact management for sales teams, project groups, and small businesses.

3.9.7 High-rise

- High-rise (www.highrisehq.com)
- A very sophisticated contact management application.
- Each contact record can include basic info (name, address, email, and so on), as well as notes,

file attachments, images, links to audio and video files, emails, and so on.

- You can even add tasks you need to get done (call, email, send a thank-you note, and so on) regarding this person; these tasks show up in the individual's contact page as well as in your master to-do list.
- Contact information (including individual notes and emails) can be aggregated by company on special company pages. Key information is summarized on your personal dashboard page, which provides a bird's-eye view of your latest activities and upcoming tasks.

3.9.8 Apple MobileMe Contacts

- MobileMe Contacts is Apple's new web based contact management application, useable by anyone with a Mac or Windows computer—as well as anyone using Apple's iPhone
- MobileMe Contacts (www.me.com) is a straight ahead contact management app with no CRM pretentions. It's essentially an address book stored in Apple's cloud that remains in sync with whatever device you use to access it.

3.9.9 MyEvents

- MyEvents (www.myevents.com) is a combination contact manager, web calendar, task manager, and online community builder. You store all your contacts online, where you can access them via any web browser or wireless device.
- The calendar function is ideal for both personal and group events, via shared public calendars. Plus you get online file storage and sharing, online digital photo albums, hosted web pages, and community bulletin boards and chat rooms.

3.9.10 Plaxo

- Plaxo (www.plaxo.com) is an odd little beast. At its heart, it's an online address book, with contact information stored in the clouds and accessible from any Internet-connected computer. But it's also been accused of being spyware (because its Outlook plug-in is installed automatically when you install various partner software, most notably AOL Instant Messenger).

3.9.11 People Matrix

- People Matrix (www.wolfereiter.com/PeopleMatrix.aspx)
- A web-based contact management application tweaked for human resources use.

3.9.12 Pipeline Deals

- Pipeline Deals (www.pipelinedeals.com)
- Offers an easy-to-use web-based CRM solution.
- The application lets you track contacts, leads, milestones, deal status, and other key data. As the name implies, Pipeline Deals is deal focused.

3.9.13 SalesBoom

- SalesBoom (www.salesboom.com) provides web-based CRM and back-office solutions, with different editions for different-sized businesses:

- Enterprise Edition, for larger organizations. Includes inventory management, product management,accounting, and human resources management solutions.
- Professional Edition, for medium-sized businesses. Includes marketing automation, sales force automation, and customer service and support solutions.
- Team Edition, for small businesses. Includes sales force automation, contact management, and customer service and support solutions. All of Salesroom's sales force automation solutions include lead management, contact management, account management, opportunity management, and forecasting features. The Enterprise Edition also includes quote management, contract management, commission's management, and a product database.

3.9.14 SalesJunction.com

- SalesJunction.com (www.salesjunction.com) offers a web-based CRM and SFA contact management system, priced on a per-user basis. Unique features include management of service cases, mass email sales campaigns, and sales pipelines.
- The company's Pro Edition also includes territory management functionality, which lets companies set up, assign, and work leads by territories. Individual users can be assigned to multiple territories, and managers can be assigned to manage as many territories as you like.

SalesNexus

- Offers Web-based contact management software SalesNexus (www.salesnexus.com).
- It was designed from the ground up around the needs of salespeople, sales management, and marketing professionals.
- Sales Nexus features include the ability to create and store proposals, estimates, quotes, and sales sheets; customized sales pipeline and activity reporting; management of automated email marketing campaigns; and automatic lead creation from website forms. In addition, Sales Nexus can capture and report the source of website leads.

3.9.15 Zoho CRM

- Our final contact management/CRM application is Zoho CRM, available in three different editions: Free Edition (for up to three users), Professional Edition, and Enterprise Edition. The application includes the following modules:
- Sales & Marketing, which integrates sales with campaigns, leads, sales pipeline, and forecasts
- Inventory Management, which provides a complete integrated inventory management system
- Customer Support & Service, which employs cases and solutions to integrate the customer support process with sales data
- Reports & Dashboards, which help you, analyze sales and marketing trends and key metrics.

SATHYABAMA
INSTITUTE OF SCIENCE AND TECHNOLOGY
(DEEMED TO BE UNIVERSITY)
Accredited "A" Grade by NAAC | 12B Status by UGC | Approved by AICTE
www.sathyabama.ac.in

SCHOOL OF ELECTRICAL AND ELECRONICS ENGINEERING

DEPARTMENT OF ELECTRONICS AND COMMUNICATION ENGINEERING

UNIT – IV – Cloud computing – SECA5303

UNIT 4

CLOUD DEPLOYMENT

Deploy application over cloud. Comparison among SAAS, PAAS, IAAS

4.0 List cloud deployment models.

Following are the four types of Cloud Deployment Models identified by NIST.

- Private cloud
- Community cloud
- Public cloud
- Hybrid cloud

4.1 Private cloud

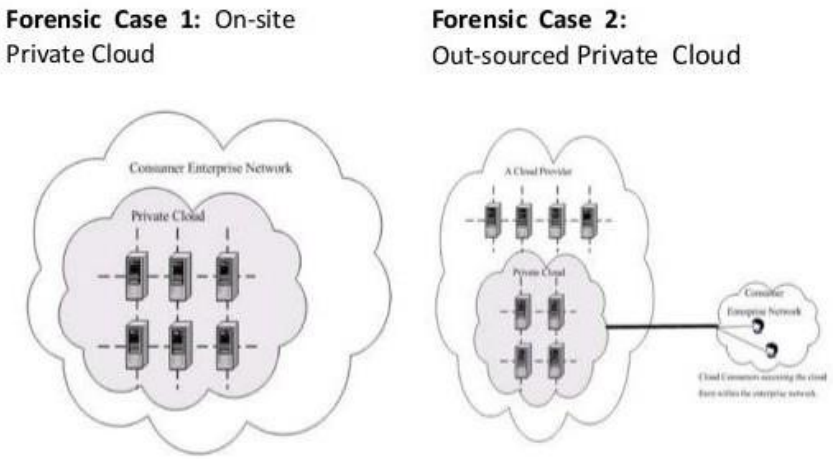

Figure 4.1 Schematic Diagram of Private Cloud

- The cloud infrastructure is operated solely for an organization.
- Contrary to popular belief, private cloud may exist off premises and can be managed by a third party. Thus, two private cloud scenarios exist, as follows:

4.1.1 On-site Private Cloud

- Applies to private clouds implemented at a customer's premises.

4.1.2Outsourced Private Cloud

- Applies to private clouds where the server side is outsourced to a hosting company.

4.1.3 Examples of Private Cloud:

- Eucalyptus
- Ubuntu Enterprise Cloud - UEC (powered by Eucalyptus)
- Amazon VPC (Virtual Private Cloud)
- VMware Cloud Infrastructure Suite
- Microsoft ECI data center.

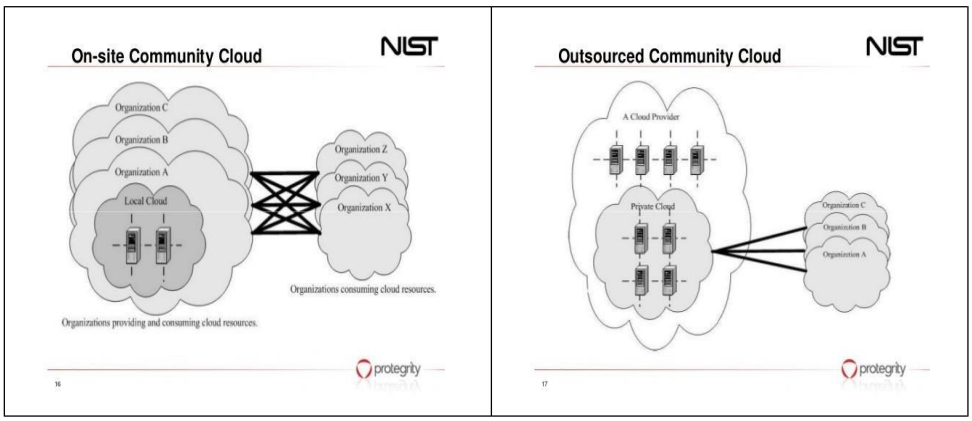

4.2 Community cloud

Figure 4.2 Schematic Diagram of Community Cloud

- The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations).
- Government departments, universities, central banks etc. often find this type of cloud useful.
- Community cloud also has two possible scenarios:

4.2.1On-site Community Cloud Scenario

- Applies to community clouds implemented on the premises of the customers composing a community cloud.

4.2.2 Outsourced Community Cloud

- Applies to community clouds where the server side is outsourced to a hosting company.

4.2.3Examples of Community Cloud:

- Google Apps for Government
- Microsoft Government Community Cloud

4.3 Public cloud

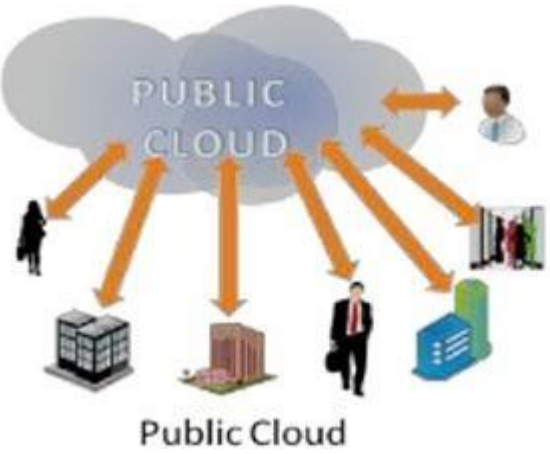

Figure 4.3 Schematic Diagram of Public Cloud

- The most ubiquitous, and almost a synonym for, cloud computing.
- The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

4.3.1 Examples of Public Cloud:

- Google App Engine
- Microsoft Windows Azure
- IBM Smart Cloud
- Amazon EC2

4.3.2 Hybrid Cloud

Figure 4.4 Schematic Diagram of Hybrid Cloud

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

4.3.3 Examples of Hybrid Cloud:

- Windows Azure (capable of Hybrid Cloud)
- VMware vCloud (Hybrid Cloud Services)

Cloud Computing Definition: Components

This slide is a visual representation of the NIST Cloud Computing Definition. We have grouped the NIST concepts into three areas-Deployment Models, Service Models, and Essential Characteristics. As you can see the key concepts in each area are Cloud characteristics On-demand self-service Broad network access Resource pooling Rapid elasticity Measured Service. Basic service models Software as a Service (SaaS) Platform as a Service (PaaS) Infrastructure as a Service (IaaS). Deployment models Private clouds Public clouds Hybrid clouds Community clouds Federated clouds, Interclouds

4.3.4 Cloud Service Models

There are actually more service models than the three (IaaS, PaaS, SaaS) widely in use today. Service models like Data Analytics as a Service and HPC/Grid as a Service are emerging as useful models. How one Selects the appropriate service model depends on factors such as availability of suitable application software Need for development and test environment, need for effective computing infrastructure control and management required distribution of data, services, and infrastructure, existence and complexity of enterprise IT infrastructure and datacenter/warehouse.

4.3.5 Use cases: Infrastructure as a Service (IaaS)

Now let us look at each of these Service Models individually.

First consider Infrastructure as a Service (IaaS). IaaS is easily used for

Website hosting, where a web server and operating system stack are put on VM's, where they can easily take advantage of cloud features such as easy scaling, global availability, managed environment, geographical load balancing, special content delivery front-end or infrastructure

Disaster Recovery preparedness, where a running server snapshot is made (including in memory image) which can be reconstituted in the event of a disaster

Test and Development where easy to create VM's can be used eliminating the need for creating test environment which may require complex infrastructure for testing new products under stress load. Actually any kind of Short-term collaborative projects, requiring a lab on demand

And of course a way to solve the needs of Cyclical and/or seasonal capacity, where the software architecture allows use of scale-out sa a way to handle increased capacity

4.3.6 Use cases: Platform as a Service (PaaS)

Now let us look at Platform as a Service (PaaS). Here, users have control over applications configuration, deployment and management. Instead of using a server metaphor, because they are given an Applications hosting and deployment platform abstracted from the server Providers usually provide all the details for the computing platform, connectivity, elasticity/scalability, backup etc.

PaaS is good for deploying applications which came from an "applications container" world before like J2EE or .NET. In fact PaaS systems such as RedHat OpenShift or Cloud Foundry from IBM or Pivotal are very much like J2EE, and Windows Azure from Microsoft is very much like .NET.

SaaS applications like Customer Relations Management (CRM) and Supply Chain Management (SCM) often come with a corresponding PaaS environment which allows for users to create their own extensions or variations on the SaaS application and host them there right "next" to the SaaS system. An example of this is the PaaS environment Force.com which complements the SaaS application Salesforce.com

4.3.7 Use cases: Cloud Application/Software as a Service (Apps/SaaS)

Now let's look at the Service Model of Cloud Application/Software as a Service (Apps/SaaS)

These are complete applications on the cloud ready for the user to use directly, usually through a browser, but in general through any "finished" user interface.

Email is a perfect example of a complete application one accesses through a browser. As an alternative example, the "cloud storage" applications are also SaaS with a very simple user interface –a "folder".

In the Mobile world SaaS is known as "an app", because the front end User Interface sits on the

phone, while the back end sits in the cloud. In fact while many Mobile apps don't look like it, they are built with the reconfigurable Mobile version of a browser called WebKit.

The SaaS providers can ensure good levels of security first by obtaining the special certification and compliance level needed in that industry and also ensuring that the underlying cloud provider obtains the special certification and compliance level needed in that industry

4.4 Cloud and High Performance Computing (HPC)

Scientists and Researchers traditionally build specialized computing systems such as Grids and Supercomputers to satisfy their needs for massive storage and computational capability. These specialized systems are very expensive.

Cloud systems have become very cost effective, comparatively, and also have been able to scale to much larger sizes than any of the biggest grids or supercomputers. As a result, Scientists and Researchers have been experimenting with using Clouds as a platform for HPC.

While the loosely coupled nature of Cloud Computing lends itself to Big Data problems (after all Internet Search was the "original" big data problem) the typical HPC problem requires closely coupled and even large coherent memory space capability. On the one hand Cloud engineers are finding ways to create specialized Cloud platforms which provide some of the capabilities HPC and Grid platforms. And on the other hand HPC and Grid scientists are finding algorithms more suited to the Cloud architecture. Together this is catalyzing an explosion of the scientific use of Clouds where previously only Grids and Supercomputers would suffice.

4.4.1 Computer Grids and Computer Clouds

Let's look more deeply at the concepts of Grids and how they relate to Cloud Computing. The Computer Grid concept was developed by the Scientific community to address needs for data-driven science where the size of the data and/or the size of the computation required is larger by orders of magnitude than the largest single "scale-up" computer. Grids have become standardized and interoperable and have been a very successful platform for science, as the slide details.

In fact, many would say that Computer Grids can be treated as a predecessor of Cloud Computing. However Computational grids focus on computationally intensive operations, which is a different architectural objective than Clouds. This is why Clouds have made adjustments to hopefully lure Grid scientists to try to use Clouds.

4.4.2 High Performance Computing (HPC) and Clouds

HPC libraries and applications are optimized for specific hardware/supercomputer. HPC user application needs sometimes bypass OS kernel and communicate directly with remote user processes and hardware/storage. Most HPC systems employ parallel file system to increase I/O bandwidth of computing

nodes. Standard Cloud is not designed for HPC but rather for Highly Scalability Computing (HSC). The slide details some of these scenarios.

4.4.3 HPC and Cloud

This illustration shows the Relationship between HPC, Cloud HSC, Cloud HPC, and Grid computing models.The programming model many Scientific problems use on Grids is called Message Passing Interface (MPI). Also, there is an assumption that the platform is implementing a clustering technique which in turn requires a special low latency network and special high performance I/O processes.Here you can see which technique is suited for which capability. It is still an art for to be able to run a HPC problem on a cloud and part of the battle is to find the right cloud platform for the type of problem at hand

4.4.4 HPC Cloud Use Cases – Vendors, Operators

Here are some examples of specialized Cloud platforms which are designed to serve the HPC/Grid community. Amazon has an offering which is physically separate from their main cloud. There are also add-on software products which further enhance the HPC capability on Amazon. There are also HPC clouds offered by and to research community Examples for XSEDE, EGI, and DEISA are detailed.

4.4.5 Amazon EC2 HPC

This slide has extremely detailed information about Amazon EC2 HPC, as Currently Amazon offers dedicated computation optimized VMs and "cluster instances" that deliver better performance to HPC users. They also offer access to GPU (Graphics Processing Units –used as Vector Machines with a special API). Please note the examples given.

Univa HPC with RightScale and Amazon AWS

This slide details the additional capability one can get on Amazon with Univa HPC and Rightscale.

Cloud Deployment Models

Now we turn to study the different Cloud Deployment Models

- Private Clouds
- Public Clouds
- Hybrid Clouds
- Community Clouds
- Federated Clouds
- Multi-clouds and Inter-clouds
- Private Cloud

First we consider the Private Cloud. The NIST definition is: the cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it

may exist on or off premises.

Private clouds are a choice for companies that already own datacenter and developed IT infrastructure and have particular needs around security or performance. They are a better choice for the company datacenter than Legacy servers in so many ways, bringing many benefits derived from virtualization and automation. However they also provide Challenges and disadvantages, mostly in that the enterprise needs to migrate or re-factor applications to take advantage of the Cloud automation.

4.4.6 Public Cloud

Now let us consider the Public Cloud. The NIST definition is: the cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider and is a form of providing public cloud services and a Cloud Service Providers business model Brings economy of scale in pooling datacenter resources, virtualization and on-demand provisioning Allows outsourcing enterprise IT infrastructure Solves/addresses disaster recovery problem Suitable for SME and agile companies.

However, the Data and processing environment are not under control of enterprises, which for applications or data with a security requirement may raise concerns. Furthermore, Service Providers can't be perfect, power outages, network issues, and so on can disrupt service.

Although the security and reliability of the public cloud will almost certainly exceed that of the private cloud due to the skill and staff size of the public cloud operator versus the enterprise itself, it is not 100% fail safe. Their SLA will certainly be problematic for many enterprises.

4.4.7 Hybrid Cloud

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

Hybrid cloud, while the most complicated configuration to manage, is also the most economical model for modern companies. It combines core cloud based enterprise infrastructure and high load tasks outsourcing to public clouds. It also Combines benefits of the controlled environment in private clouds and rapid elasticity of public clouds

However, it requires deeper enterprise IT/cloud modernization. Processes/workflow require re-engineering and re-architecting. And the challenges of getting seamless integration between the Private and the Public cloud can be solved but by custom work, and the issues of Compatibility, standardization and compliance are not turnkey yet.

4.4.8 Community Cloud

The Community Cloud, while used by several constituencies, is something slightly different from a Public Cloud, The NIST definition is – The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises. The Community Cloud then involves cooperation and integration of IT infrastructure and resources from multiple organizations. It May serve large inter-organizational projects. For example, the scientists form many organizations at the CERN Large Hadron Collider (HC) share a Community Cloud. It requires interoperability and compliance between member organizations and their resources, including Identity management.

4.4.9 Emerging models: Federated and Intercloud

There are some interesting emerging models. Cloud Researchers have suggested that, like the Internet, a mechanism should exist for users to be able to utilize multiple clouds from multiple providers or companies, and not have the details visible. In the Internet, the user sees a uniform and global topology. The user is unaware of which Internet Service Provider is hosting the web site which is he viewing, and that user is also very unaware how packets have traveled in between their browser and that web site.

We will now look at the models of Cloud federation, which allows inter-cloud resources sharing and combined provisioning.

There are two types of cloud federation; a Provider side federation for resources sharing and provisioning, and a User/customer side federation that allows creation of multi-provider heterogeneous cloud infrastructures.

The provider side federation is called “Intercloud”. The Intercloud deployment model provides a general framework for multi-provider heterogeneous cloud based services and infrastructures building and operation.

4.4.10 Use cases for Private clouds

Now we turn to use cases for each type of Cloud. First we consider Use cases for Private clouds. Private clouds will be chosen if there is a requirement for High data I/O and low network latency: disk intensive processes, wide sensor network, or process control. There are Legacy applications and in some cases special equipment requirements. There are Specialty hardware or configuration requirements, e.g. VM with 32+GB for in-memory data processing and governance or regulatory requirements.

Use cases for Public and Hybrid Clouds

Use cases for Public and Hybrid Clouds include, situations where there are Unpredictable growth: game or social websites, marketing campaigns. Where Cyclical: applications with regular daily or seasonal traffic fluctuation such as financial markets or eCommerce. Or for easily parallelized:

applications using batch processing, data analytics, media encoding Example use cases Web/social web and mobile applications Development and test, proof of concept Big Data analytics, business reporting.

4.5 Use cases for Hybrid clouds

Use cases for Hybrid clouds covers where a company combines both private and public clouds. For example, the Private cloud hosts regular workload and master processes, also security and compliance critical applications/tasks. The Public cloud hosts non-critical and not regular but computing intensive workload. This scenario requires compatibility between private and public cloud platforms.

Hybrid cloud supports the notion of Cloudbursting – term widely used by businesses to describe a situation when workload is temporarily migrated to cloud, extending and replicating the private cloud resources and VMs (using formula "buy the base, rent a spike")

Use cases for Intercloud and Multi-cloud

As to Use cases for Intercloud and Multi-cloud, the Intercloud deployment model provides a basis for provisioning heterogeneous multi-provider cloud based project oriented infrastructures on-demand

The first and most powerful idea is for the providers to connect, and provide a seamless integration of their public clouds, just as the internet is. Just as the Internet connects to nearly everything these days, so would the Intercloud.

The other use cases have to do with uniting various combinations of clouds (enteprrise/campus, groups of Community clouds, etc) in a transparent way.

It is clear that today Clouds are somewhat "walled gardens" and interoperability and federation will emerge in some fashion or another, eventually.

4.5.1 Use case: Intercloud infrastructure provisioning: Workflow => Logical (Cloud) Infrastructure (1)

Let us examine a use case, where I have a large application which can be mapped out in functionality by its Master Workflow. This is illustrated in the Slide on the top. I have several cloud service providers, as illustrated by the four entities on the bottom of the slide. Actually, a participant in this can be non-cloud, as long as it participates in the same federation protocols, as indicated by the terms "Resource Service Provider".

Through Federation, a "logical Intercloud" is formed, and the 4 different datacenters let's say, appear to part of one virtual cloud. This is indicated as the "cloud in the middle". To further complicate our example, the problem is being looked after by two universities and so the User Groups A and B represent traditional diagrams of a control function.

Use case: Intercloud infrastructure provisioning: Workflow => Logical (Cloud) Infrastructure (2)

Enterprise/Scientific workflow is mapped to heterogeneous cloud infrastructure containing IaaS, PaaS components.

This slide shows the mapping of the virtual Resources

Use case: Intercloud infrastructure provisioning: Logical Infrastructure

=> Network Infrastructure (1)

The application is mapped to run on the virtual VM's which are mapped to the actual underlying services. The two campuses in this example can monitor and manage the Workflow across VMs across the federated cloud. Here, the network of one for the cloud carriers is used to choreograph all that occurs.

Use case: Intercloud infrastructure provisioning: Logical Infrastructure

=> Network Infrastructure (2)

Or as in the previous example, but the orchestration occurs over the public internet. Either way the composite application is choreographed or the elements of application code are set to run on the public cloud.

These sets of slides meant to illustrate scenarios where our definition of "an application" is not constrained to one cloud or to our native service provider either. Intercloud allows for a composite application to be deployed as it makes sense, putting some components with one Cloud Service Provider and some with another.

Enterprise related use cases

Now we are going to consider a number of Enterprise related use cases.

Moving Enterprise IT infrastructure to cloud

Private cloud, Hybrid cloud, Virtual Private Cloud Website hosting on cloud Application development and testing Lab on-demand in cloud Computer Cluster on-demand

See also Univa, RightScale, Amazon use case above.

Use cases and business relationships

In considering Use cases and business relationships, there has been detailed analysis in looking at the important combinations possible, for example, considering classification Service models and deployment models, Stakeholders involvement and business relations, and Industry or community use cases.

What we get are scenarios like End users to Cloud, Enterprise to Cloud to End users, Enterprise to Cloud, Enterprise to Cloud to Enterprise, Private Cloud, Changing Cloud Providers, etc. The following slides examine several of these.

4.5.2 Website hosting in Cloud

The easiest and most used use case for Cloud is that of Website hosting. This use case shows how to deploy the infrastructure for this in the cloud.

In the Company's IT infrastructure, we need a Local database server, access to the various Department IT resources, and IT Administration capability. We link to the cloud using a Secure Virtual Private Network (VPN) or Virtual Private Cloud (VPC) connection, In the Cloud, we spin up the web site facilities (using

a 3-tier model) including Database, servers, storage; Web servers, Apps servers; and a Load balancer.

The illustration shows the topology for this use case. This use case leverages Cloud for outsourcing web-based customer/user facing services while keeping in house IT infrastructure and a set of resources to support internal/office operation and store critical data required for local company's services. It allows the enterprise to: automatically scale up and down cloud based resources that have cyclical or random demand increase, paying only for what you need and when you need; simplify geographical expansion and geographical load distribution; avoid building expensive physical private webserver and application servers, also projected for peak demand periods; maintain enterprise control and security over cloud resources using VPN access to cloud resources

4.5.3 Test and Development, Proof of Concept

A second illustrative use case is the one for Test and Development, Proof of Concept.

Using the cloud Eliminates need for creating test environment which may require complex infrastructure for testing new products under stress load It is a Cheap alternative for proof-of- concept services try out without capital investment

Application development and testing

Another excellent use case for the Cloud is to utilize it for Application development and testing. In the Company's IT infrastructure there is the Integrated Development Environment (IDevE), Quality Assurance (QA), Performance testing (PERF), and the Developers. In the Cloud based facilities one places a Replicated IDevE, QA, and PERF environment, along with Web servers, Apps servers, Database/Storage, and test/load generator. The company and the Cloud are connected via a Secure Virtual Private Network (VPN) or Virtual Private Cloud (VPC) connection.

Each time an expanded development and testing capability is needed , the cloud environment is spun up. Moving application development and testing to clouds allows for testing of close to real new applications infrastructure and load without creating expensive in-house test environment. Enterprise IT infrastructure needs only resources to support development with limited local test functionalities. Non-production test environment is moved from production facilities to clouds. Provide development and QA teams with on-demand close to real testing environment using full scale load. Use cloud platform based monitoring and performance measurement tools. Empower company to respond quickly and cost-effectively to market demands for test, trial and proof-of-concept activities.

4.5.4 Lab on-demand in cloud

As an extension to the Application development and testing Use Case, is to have a full Lab on-demand in the Cloud. It is nearly identical to the previous use case but in a more general way organizations can put a variety of test and measurement gear up in the cloud.

Suitable for both development lab and educational lab, fully functional virtual lab facilities, on- demand

provisioning, adopted to cyclical educational process, preset configuration, easy user management.

Computer Cluster on-demand

A final example use case is that of a Computer Cluster on-demand. This use case addresses enterprises need for Big Data analytics capabilities, makes use of cloud based Data Analytics facilities affordable for SME, solves Big Data storage issues, moves computational workflow/jobs management.

In the Company's IT infrastructure there are the Data repositories, Data servers, Data management system, Analytics capabilities repository , and Data visualization facilities. In the Cloud based facilities there are Data servers (SQL, NoSQL, etc), Data Analytics tools, High Performance Computing (HPC) cluster, and Scientific/Business Workflow Management. A High Performance Optical Network for fast data upload and download is used.

This allows one to:

Build proprietary computer cluster using general cloud compute an storage resources or use cloud based Data Analytics services

Makes HPC and Big Data analytics services accessible for SME

Reduce capex for big companies by keeping limited data analytics capabilities in-house outsourcing peak computation job to cloud

Move computational HPC workflow/jobs to cloud using the same scientific/business workflow management/monitoring system/tools Requires high-performance optical network, typically dedicated.

Wrap up and Take away

Clouds provide 3 basic service models and a number of deployment models that reflects the variety of business and service relations Selecting right service and deployment model is an important stage in cloud implementation by enterprise Enterprise may adopt gradual/staged cloud implementation from IaaS to PaaS/SaaS and from private cloud to hybrid and public cloud.

Demand from Big Data and data intensive applications facilitates demand for HPC services as a part of cloud infrastructure Clouds are not designed for HPC but rather as Highly Scalable Computing (HSC) environment Clouds can be effectively used for tasks parallelization (also called "embarrassing.

SATHYABAMA
INSTITUTE OF SCIENCE AND TECHNOLOGY
(DEEMED TO BE UNIVERSITY)
Accredited "A" Grade by NAAC | 12B Status by UGC | Approved by AICTE
www.sathyabama.ac.in

SCHOOL OF ELECTRICAL AND ELECRONICS ENGINEERING

DEPARTMENT OF ELECTRONICS AND COMMUNICATION ENGINEERING

UNIT – V – Cloud computing – SECA5303

UNIT 5

INTRODUCTION TO OPEN SOURCE IOT CLOUD PLATFORMS

Introduction to IoT cloud platforms like Open Shift, Kaa etc.

5.1 DynamoDB

- Amazon DynamoDB -- also known as Dynamo Database or DDB -- is a fully managed NoSQL database service provided by Amazon Web Services. DynamoDB is known for low latencies and scalability.
- According to AWS, DynamoDB makes it simple and cost-effective to store and retrieve any amount of data, as well as serve any level of request traffic.
- All data items are stored on solid-state drives, which provide high I/O performance and can more efficiently handle high-scale requests.
- An AWS user interacts with the service by using the AWS Management Console or a DynamoDB API.
- DynamoDB uses a NoSQL database model, which is nonrelational, allowing documents, graphs and columnar among its data models.
- A user stores data in DynamoDB tables, then interacts with it via GET and PUT queries, which are read and write operations, respectively.
- DynamoDB supports basic CRUD operations and conditional operations. Each DynamoDB query is executed by a primary key identified by the user, which uniquely identifies each item.

5.1.1 Scalability, availability and durability

- DynamoDB enforces replication across three availability zones for high availability, durability and read consistency.
- A user can also opt for cross-region replication, which creates a backup copy of a DynamoDB table in one or more global geographic locations.
- The DynamoDB scan API provides two consistency options when reading DynamoDB data:
    - Eventually consistent reads
    - Strongly consistent reads
- The former, which is the AWS default setting, maximizes throughput at the potential expense of not having a read reflect the latest write or update. The latter reflects all writes and updates.
- There are no DynamoDB limits on data storage per user, nor a maximum throughput per table.

5.1.2 Security

- Amazon DynamoDB offers Fine-Grained Access Control (FGAC) for an administrator to protect

data in a table.

- The admin or table owner can specify who can access which items or attributes in a table and what actions that person can perform.

- FGAC is based on the AWS Identity and Access Management service, which manages credentials and permissions.
- As with other AWS products, the cloud provider recommends a policy of least privilege when granting access to items and attributes.
- An admin can view usage metrics for DynamoDB with Amazon CloudWatch.

5.1.3 Additional DynamoDB features

- The DynamoDB Triggers feature integrates with AWS Lambda to allow a developer to code actions based on updates to items in a DynamoDB table, such as sending a notification or connecting a table to another data source.
- The developer associates a Lambda function, which stores the logic code, with the stream on a DynamoDB table.
- AWS Lambda then reads updates to a table from a stream and executes the function.
- The DynamoDB Streams feature provides a 24-hour chronological sequence of updates to items in a table.
- An admin can access the stream via an API call to take action based on updates, such as synchronizing information with another data store. An admin enables DynamoDB Streams on a per-table basis.

5.1.4 Write short note on Relational Database Service.

- Amazon Relational Database Service (Amazon RDS) makes it easy to set up, operate, and scale a relational database in the cloud.
- It provides cost-efficient and resizable capacity while automating time-consuming administration tasks such as hardware provisioning, database setup, patching and backups.
- It frees you to focus on your applications so you can give them the fast performance, high availability, security and compatibility they need.
- Amazon RDS is available on several database instance types optimized for memory, performance or I/O and provides you with six familiar database engines to choose from, including Amazon Aurora, PostgreSQL, MySQL, MariaDB, Oracle, and Microsoft SQL Server.
- You can use the AWS Database Migration Service to easily migrate or replicate your existing databases to Amazon RDS.

5.2 Advantages/Benefits

(i) Easy to Administer

Amazon RDS makes it easy to go from project conception to deployment. Use the AWS Management Console, the AWS RDS Command-Line Interface, or simple API calls to access the capabilities of a production-ready relational database in minutes. No need for infrastructure provisioning, and no need for installing and maintaining database software.

(ii) Highly Scalable

We can scale our database's compute and storage resources with only a few mouse clicks or an API call, often with no downtime. Many Amazon RDS engine types allow you to launch one or more Read Replicas to offload read traffic from your primary database instance.

(iii) Available and Durable

Amazon RDS runs on the same highly reliable infrastructure used by other Amazon Web Services. When you provision a Multi-AZ DB Instance, Amazon RDS synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Amazon RDS has many other features that enhance reliability for critical production databases, including automated backups, database snapshots, and automatic host replacement.

(iv) Fast

Amazon RDS supports the most demanding database applications. You can choose between two SSD- backed storage options: one optimized for high-performance OLTP applications, and the other for cost- effective general-purpose use. In addition, Amazon Aurora provides performance on par with commercial databases at $1/10^{th}$ the cost.

(v) Secure

Amazon RDS makes it easy to control network access to your database. Amazon RDS also lets you run your database instances in Amazon Virtual Private Cloud (Amazon VPC), which enables you to isolate your database instances and to connect to your existing IT infrastructure through an industry-standard encrypted IPsec VPN. Many Amazon RDS engine types offer encryption at rest and encryption in transit.

(vi) Inexpensive

You pay very low rates and only for the resources you actually consume. In addition, you benefit from the option of On-Demand pricing with no up-front or long-term commitments, or even

lower hourly rates via Reserved Instance pricing.

5.3 Redshift

- Perhaps one of the most exciting outcomes of the public cloud was addressing the shortcomings of traditional enterprise data warehouse (EDW) storage and processing. The fast provisioning, commodity costs, infinite scale, and pay-as-you-grow pricing of public cloud are a natural fit for EDW needs, providing even the smallest of users the ability to now get valuable answers to BI questions.
- **Amazon Redshift** is one such system built to address EDW needs, and it boasts low costs, an easy SQL- based access model, easy integration to other Amazon Web Services (AWS) solutions, and most importantly, high query performance.
- Amazon Redshift gets its name from the astronomical phenomenon noticed by Hubble, which explained the expansion of the universe. By adopting the Amazon Redshift moniker, AWS wanted to relay to customers that the service was built to handle the perpetual expansion of their data.

Amazon Redshift Architecture:

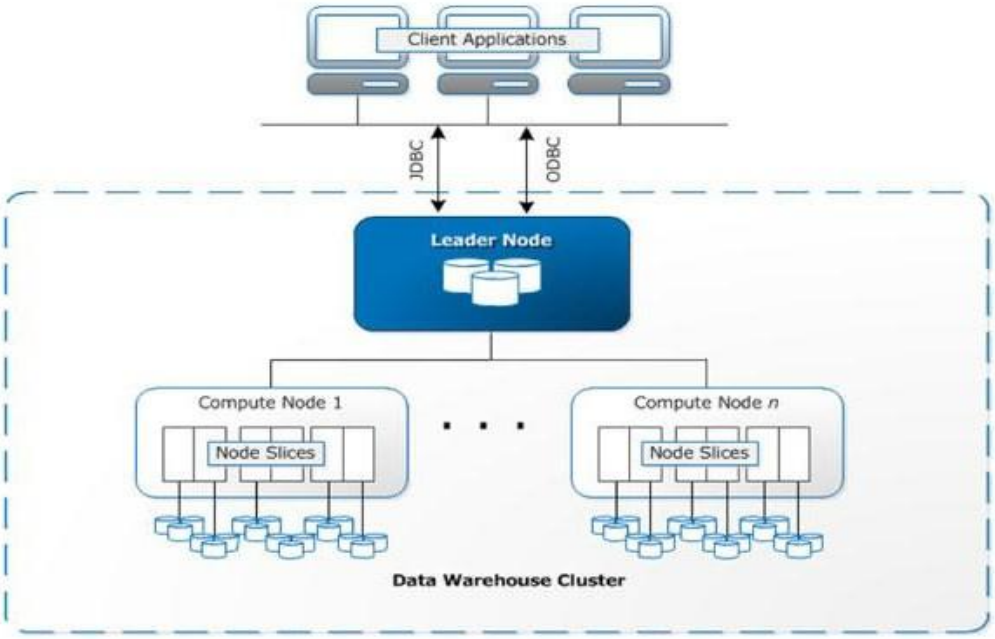

Figure 5.1 Schematic diagram of Amazon Redshift Architecture

- An Amazon Redshift cluster consists of one leader node (which clients submit queries to) and one or more follower (or "compute") nodes, which actually perform the queries on locally stored data.
- By allowing for unlimited expansion of follower nodes, Amazon Redshift ensures that customers can continue to grow their cluster as their data needs grow.
- Customers can start with a "cluster" as small as a single node (acting as both leader and follower), and for the smallest supported instance type (a DW2), that could be as low cost as $0.25/hour or about
- $180/month. By using "Reservations" (paying an up-front fee in exchange for a lower hourly running cost) for the underlying instances, Amazon Redshift can cost as little as $1,000/TB/year — upwards of one-fifth to one-tenth of the cost of a traditional EDW.
- Because Amazon Redshift provides native Open Database Connectivity (ODBC) and Database Connectivity (JDBC) connectivity (in addition to PostgresSQL driver support), most third-party BI tools (like Tableu, Qlikview, and MicroStrategy) work right out of the box. Amazon Redshift also uses the ubiquitous Structured Query Language (SQL) language for queries, ensuring that your current resources can quickly and easily become productive with the technology.
- Amazon Redshift was custom designed from the ParAccel engine — an analytic database which

used columnar storage and parallel processing to achieve very fast I/O.

- Columns of data in Amazon Redshift are stored physically adjacent on disk, meaning that queries and scans on those columns (common in online analytical processing [OLAP] queries) run very fast.
- Additionally, Amazon Redshift uses 10GB Ethernet interconnects, and specialized EC2 instances (with between three and 24 spindles per node) to achieve high throughput and low latency.
- For even faster queries, Amazon Redshift allows customers to use column-level compression to both greatly reduce the amount of data that needs stored, and reduce the amount of disk I/O.
- Amazon Redshift, like many of AWS's most popular services, is also fully managed, meaning that low- level, time-consuming administrative tasks like OS patching, backups, replacing failed hardware, and software upgrades are handled automatically and transparently.
- With Amazon Redshift, users simply provision a cluster, load it with their data, and begin executing queries. All data is continuously, incrementally, automatically backed up in the highly durable S3, and enabling disaster recovery across regions can be accomplished with just a few clicks.
- Spinning a cluster up can be as simple as a few mouse clicks, and as fast as a few minutes.
- A very exciting aspect of Amazon Redshift, and something that is not possible in traditional EDWs, is the ability to easily scale a provisioned cluster up and down.
- In Amazon Redshift, this scaling is transparent to the customer—when a resize is requested, data is copied in parallel from the source cluster (which continues to function in read-only mode) to a new cluster, and once all data is live migrated, DNS is flipped to the new cluster and the old cluster is de- provisioned.
- This allows customers to easily scale up and down, and each scaling event nicely re-stripes the data across the new cluster for a balanced workload.
- Amazon Redshift offers mature, native, and tunable security. Clusters can be deployed into a Virtual Private Cloud (VPC), and encryption of data is supported via hardware accelerated AES-256 (for data at rest) and SSL (for data on the wire).
- Compliance teams will be pleased to learn that users can manage their own encryption keys via AWS's Hardware Security Module (HSM) service, and that Amazon Redshift provides a full audit trail of all SQL connection attempts, queries, and modifications of the cluster.

5.4 ElastiCache

- ElastiCache is a web service that makes it easy to set up, manage, and scale a distributed in-memory data store or cache environment in the cloud.
- It provides a high-performance, scalable, and cost-effective caching solution, while removing the

complexity associated with deploying and managing a distributed cache environment.

- With ElastiCache, you can quickly deploy your cache environment, without having to provision hardware or install software.
- You can choose from Memcached or Redis protocol-compliant cache engine software, and let ElastiCache perform software upgrades and patch management for you.
- For enhanced security, ElastiCache can be run in the Amazon Virtual Private Cloud (Amazon VPC) environment, giving you complete control over network access to your clusters.
- With just a few clicks in the AWS Management Console, you can add or remove resources such as nodes, clusters, or read replicas to your ElastiCache environment to meet your business needs and application requirements.
- Existing applications that use Memcached or Redis can use ElastiCache with almost no modification.
- Your applications simply need to know the host names and port numbers of the ElastiCache nodes that you have deployed.
- The ElastiCache Auto Discovery feature for Memcached lets your applications identify all of the nodes in a cache cluster and connect to them, rather than having to maintain a list of available host names and port numbers.
- In this way, your applications are effectively insulated from changes to node membership in a cluster.
- ElastiCache has multiple features to enhance reliability for critical production deployments:
    - Automatic detection and recovery from cache node failures.
    - Multi-AZ with Automatic Failover of a failed primary cluster to a read replica in Redis clusters that support replication (called replication groups in the ElastiCache API and AWS CLI.
    - Flexible Availability Zone placement of nodes and clusters.
    - Integration with other AWS services such as Amazon EC2, Amazon CloudWatch, AWS CloudTrail, and Amazon SNS to provide a secure, high-performance, managed in-memory caching solution.

5.5 High performance AWS Networking.

- High performance AWS Networking is nothing but use of various network services provided by AWS for better performance.
- AWS Networking include following services:

1. Private DNS Servers
    - The Private DNS are name servers that reflect your domain name rather than our default ones.
    - Having private nameservers could be useful if you intend to resell hosting services or want to brand your business.
    - Also, when using Private DNS, if a domain name is migrated to another server, there is no need to change any nameservers and the domain names will automatically point to the new location.
2. Virtual Private Clouds (Explain Earlier),
3. Cloud Models (Explain Earlier), etc.

5.6 Cloud watch and cloud formation

5.6.1 Cloud watch

- Amazon CloudWatch is a monitoring service for AWS cloud resources and the applications you run on AWS.
- You can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, set alarms, and automatically react to changes in your AWS resources.
- Amazon CloudWatch can monitor AWS resources such as Amazon EC2 instances, Amazon DynamoDB tables, and Amazon RDS DB instances, as well as custom metrics generated by your applications and services, and any log files your applications generate.
- You can use Amazon CloudWatch to gain system-wide visibility into resource utilization, application performance, and operational health.
- You can use these insights to react and keep your application running smoothly.

5.6.2 Cloud formation

- AWS CloudFormation provides a common language for you to describe and provision all the infrastructure resources in your cloud environment.
- CloudFormation allows you to use a simple text file to model and provision, in an automated and secure manner, all the resources needed for your applications across all regions and accounts.
- This file serves as the single source of truth for your cloud environment.
- AWS CloudFormation is available at no additional charge, and you pay only for the AWS resources needed to run your applications.

5.6.3 Formation MODEL IT ALL

- AWS CloudFormation allows you to model your entire infrastructure in a text file. This template

becomes the single source of truth for your infrastructure. This helps you to standardize infrastructure components used across your organization, enabling configuration compliance and faster troubleshooting.

AUTOMATE AND DEPLOY

- AWS CloudFormation provisions your resources in a safe, repeatable manner, allowing you to build and rebuild your infrastructure and applications, without having to perform manual actions or write custom scripts. CloudFormation takes care of determining the right operations to perform when managing your stack, and rolls back changes automatically if errors are detected.

IT'S JUST CODE

- Codifying your infrastructure allows you to treat your infrastructure as just code. You can author it with any code editor, check it into a version control system, and review the files with team members before deploying into production.

Disaster Recovery (DR)

- Our data is the most precious asset that we have and protecting it is our top priority.
- Creating backups of our data to an off shore data center, so that in the event of an on premise failure we can switch over to our backup, is a prime focus for business continuity.
- As AWS says, 'Disaster recovery is a continual process of analysis and improvement, as business and systems evolve. For each business service, customers need to establish an acceptable recovery point and time, and then build an appropriate DR solution.'
- Backup and DR on Cloud reduces costs by half as compared to maintaining your own redundant data centers. And if you think about it, it's really not that surprising.
- Imagine the kind of cost you would entail in buying and maintaining servers and data centers, providing secure and stable connectivity and not to mention keeping them secure.
- You would also be underutilizing severs; and in times of unpredictable traffic rise it would be strenuous to set up new ones. To all these cloud provides a seamless transition reducing cost dramatically.

5.7 Standard Approaches of Backup and Disaster Recovery Using Amazon Cloud

5.7.1 Backup and Recovery

- To recover your data in the event of any disaster, you must first have your data periodically backed up from your system to AWS.
- Backing up of data can be done through various mechanisms and your choice will be based on the RPO (Recovery Point Objective- So if your disaster struck at 2 pm and your RPO is 1 hr, your Backup & DR will restore all data till 1 pm.) that will suit your business needs.

- AWS offers AWS Direct connect and Import Export services that allow for faster backup.
- For example, if you have a frequently changing database like say a stock market, then you will need a very high RPO. However if your data is mostly static with a low frequency of changes, you can opt for periodic incremental backup.
- Once your backup mechanisms are activated you can pre-configure AMIs (operating systems & application software).
- Now when a disaster strikes, EC2 (Elastic Compute Capacity) instances in the Cloud using EBS (Elastic Block Store) coupled with AMIs can access your data from the S3 (Simple Storage Service) buckets to revive your system and keep it going.

5.7.2 Pilot Light Approach

- The name pilot light comes from the gas heater analogy. Just as in a heater you have a small flame that is always on, and can quickly ignite the entire furnace; a similar approach can be thought of about your data system.
- In the preparatory phase your on premise database server mirrors data to data volumes on AWS. The database server on cloud is always activated for frequent or continuous incremental backup.
- This core area is the pilot from our gas heater analogy. The application and caching server replica environments are created on cloud and kept in standby mode as very few changes take place over time.
- These AMIs can be updated periodically. This is the entire furnace from our example. If the on premise system fails, then the application and caching servers get activated; further users are rerouted using elastic IP addresses to the ad hoc environment on cloud. Your Recovery takes just a few minutes

5.7.3 Warm Standby Approach

- This Technique is the next level of the pilot light, reducing recovery time to almost zero.
- Your application and caching servers are set up and always activated based on your business critical activities but only a minimum sized fleet of EC2 instances are dedicated.
- The backup system is not capable of handling production load, but can be used for testing, quality assurance and other internal uses.
- In the event of a disaster, when your on premise data center fails, two things happen.
- Firstly multiple EC2 instances are dedicated (vertical and horizontal scaling) to bring your application and caching environment up to production load. ELB and Auto Scaling (for distributing traffic) are used to ease scaling up.
- Secondly using Amazon Route 53 user traffic is rerouted instantly using elastic IP addresses and

there is instant recovery of your system with almost zero down time.

5.7.4 Multi-Site Approach

- Well this is the optimum technique in backup and DR and is the next step after warm standby.
- All activities in the preparatory stage are similar to a warm standby; except that AWS backup on Cloud is also used to handle some portions of the user traffic using Route 53.
- When a disaster strikes, the rest of the traffic that was pointing to the on premise servers are rerouted to AWS and using auto scaling techniques multiple EC2 instances are deployed to handle full production capacity.
- You can further increase the availability of your multi-site solution by designing Multi-AZ architectures.