



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 COMPUTING
DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING

CLOUD COMPUTING – SBSA1402

COURSE OBJECTIVES

- To articulate the main concepts, key technologies, strengths, and limitations of cloud computing
- To learn how to use Cloud Services.
- To implement Virtualization
- To understand the basic concepts of cloud computing and cloud enables.
- To understand cloud services and Multi-tenancy computing
- To study about various models of cloud environments and virtualization.
- To manage the cloud computing infrastructure with security.
- To gain knowledge of cloud and to understand about Fog, edge computing.

COURSE OUTCOMES

CO1 - Analyze the Cloud computing setup with its vulnerabilities and applications using different architectures.

CO2 - Implement and install the cloud tools to make enable the cloud computing infrastructures.

CO3 - Apply and design suitable Virtualization concept, Cloud Resource Management and collaboration services.

CO4 - Create combinatorial auctions for cloud resources and services for computing clouds
Develop and make cloud services as commercial.

CO5 - Assess cloud and cloud to Fog with IoT.

CO6 - Ability to, understand fog computing architecture

SYLLABUS

UNIT I

(9 Hrs)

The vision of cloud computing - The cloud computing reference model - Characteristics and benefits - Challenges ahead - Historical developments - Distributed systems -Virtualization - Building cloud computing environments – Application development - Infrastructure and system development - Computing platforms and technologies.Parallel vs. distributed computing -

Elements of parallel computing - Hardware architectures for parallel processing Approaches to parallel programming - Laws of caution.

Unit-II (9 Hrs)

Introduction - The cloud reference model - Types of clouds - Economics of the cloud
Virtualization: Introduction - Characteristics of virtualized environments - Taxonomy of virtualization techniques - Virtualization and cloud computing - Pros and cons of virtualization - Technology example: VMware: full virtualization.

Unit-III (9 Hrs)

Anatomy of the Aneka container - Introducing parallelism for single-machine computation - Programming applications with threads - Multithreading with aneka - Programming applications with aneka threads. Cloud computing economics: Cloud infrastructure - Economics of private clouds - Software productivity in the cloud - Economies of scale: public vs. private clouds.

Unit-IV (9 Hrs)

Multi-entity support - Multi-schema approach - Multi-tenancy using cloud data stores - Data access control for enterprise applications. Data in the cloud: Relational databases - Cloud file systems: GFS and HDFS - BigTable, HBase - Cloud data stores: Datastore and SimpleDB.

Unit-V (9 Hrs)

Amazon web services: Compute services - Storage services - Communication services - Additional services. Google AppEngine: Architecture and core concepts - Application life cycle - Cost model – Observations. Microsoft azure: Azure core concepts - SQL azure - Windows azure platform appliance.

Max:45 Hrs



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 COMPUTING
DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING

UNIT – 1-CLOUD COMPUTING – SBSA1402

UNIT – 1

The vision of cloud computing - The cloud computing reference model - Characteristics and benefits - Challenges ahead - Historical developments - Distributed systems - Virtualization - Building cloud computing environments – Application development - Infrastructure and system development - Computing platforms and technologies. Parallel vs. distributed computing - Elements of parallel computing - Hardware architectures for parallel processing Approaches to parallel programming - Laws of caution.

Cloud computing Definition:

The cloud is a large group of interconnected computers. These computers can be personal computers or network servers; they can be public or private.

Cloud computing is a technology that uses the internet and central remote servers to maintain data and applications.

Eg: Yahoo email or Gmail etc

Cloud Computing Reference Model:

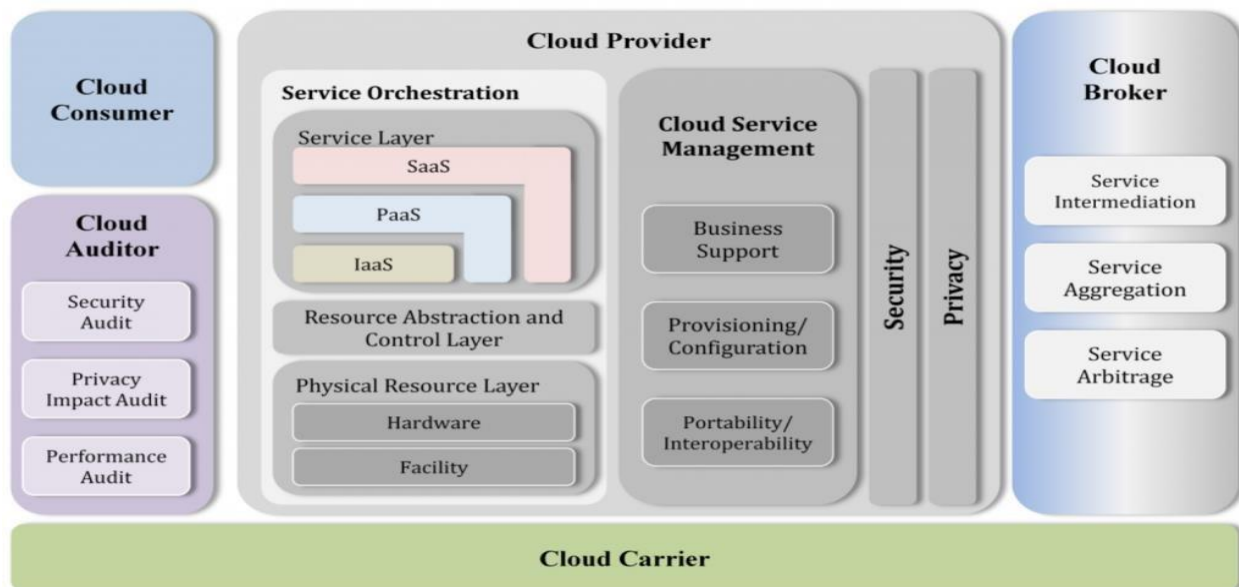


Fig 1 : Cloud Reference Model

The NIST cloud computing reference architecture defines five major actors:

- 1.Cloud consumer
- 2.Cloud provider
- 3.Cloud carrier
- 4.Cloud auditor and
- 5.Cloud broker.

Each actor is an entity (a person or an organization) that participates in a transaction or process and/or performs tasks in cloud computing.

Actors in Cloud Computing

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

Table 1: Actors of Cloud

Interactions among the actors:

A cloud consumer may request cloud services from a cloud provider directly or via a cloud broker.

A cloud auditor conducts independent audits and may contact the others to collect necessary information.

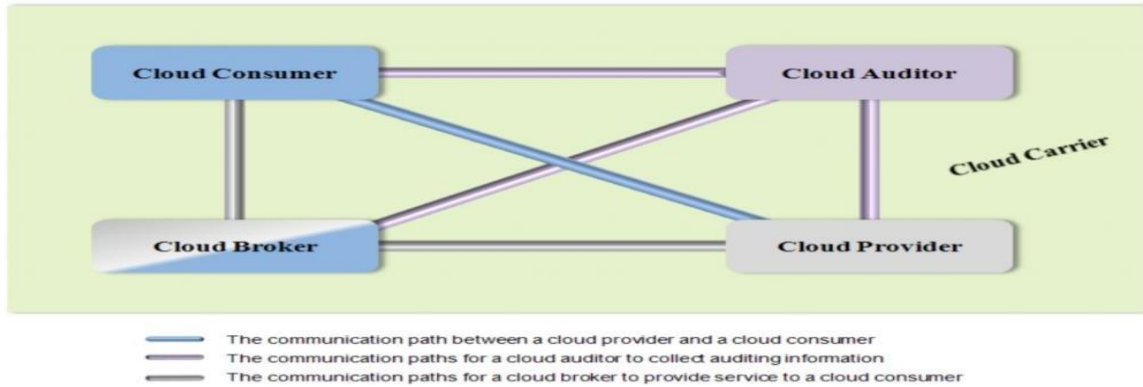


Fig 2 : Cloud Actor Interactions

- Cloud carriers provide the connectivity and transport of cloud services from cloud providers to cloud consumers.
- A cloud provider participates in and arranges for two unique service level agreements (SLAs), one with a cloud carrier (e.g. SLA2) and one with a cloud consumer (e.g. SLA1).
- A cloud provider arranges service level agreements (SLAs) with encrypted connections to ensure the cloud services are consumed at a consistent level according to the contractual obligations with the cloud consumers.
- A cloud service, a cloud auditor conducts independent assessments of the operation and security of the cloud service implementation. The audit may involve interactions with both the Cloud Consumer and the Cloud Provider




Fig 3: Usage Scenario for Cloud Carrier

Characteristics of cloud computing

1. Cloud computing is user-centric.

Once users are connected to the cloud, whatever is available (stored) in cloud as documents, messages, images, applications, it becomes users own.

Users can also share it with others.

2. Cloud computing is task-centric.

It focusing on the application for users on what users need to done and how the application can do it for users. Eg collage, Animation, Movie in Google Photos

3. Cloud computing is powerful.

Connecting hundreds or thousands of computers together in a cloud creates a wealth of computing power impossible with a single desktop PC.

4. Cloud computing is accessible.

Data is stored in the cloud; users can instantly retrieve more information from multiple repositories. You're not limited to a single source of data.

5. Cloud computing is intelligent.

With all the various data stored on the computers in a cloud, data mining and analysis are necessary to access that information in an intelligent manner.

6. Cloud computing is programmable.

Many of the tasks necessary with cloud computing must be automated.

For example, cloud must be replicated on other computers in the cloud. If that one computer goes offline, the cloud's programming automatically redistributes that computer's data to a new computer in the cloud.

Examples of cloud computing applications: Google Docs & Spreadsheets, Google Calendar, Gmail, Picasa.

Pros and Cons of Cloud Computing Advantages

- Lower-Cost Computers for Users
- Improved Performance
- Lower IT Infrastructure Costs
- Fewer Maintenance Issues

- Lower Software Costs
- Instant Software Updates
- Increased Computing Power
- Unlimited Storage Capacity
- Increased Data Safety
- Improved Compatibility Between Operating Systems
- Improved Document Format Compatibility
- Easier Group Collaboration
- Universal Access to Documents
- Latest Version Availability
- Removes the Tether to Specific Devices

Dis-Advantages

- Requires a Constant Internet Connection
- Doesn't Work Well with Low-Speed Connections
- Can Be Slow
- Features Might Be Limited
- Stored Data Might Not Be Secure
- Problem will arise If Data loss occurs

Historical Developments

Client/Server Computing: Centralized Applications and Storage

- In *client/server* model all the software applications, data, and the control resided on huge mainframe computers, known as *servers*.
- If a user wanted to access specific data or run a program, he had to connect to the mainframe, gain appropriate access, and then do his business.
- Users connected to the server via a computer terminal, called a workstation or *client*.

Drawbacks in client /server Model

- Even on a mainframe computer, processing power is limited.
- Access was not immediate nor could two users access the same data at the same time. When multiple people are sharing a single computer, you have to wait for your turn.
- There isn't always immediate access in a client/server environment.

- So the client/server model, while providing similar centralized storage, differed from cloud computing in that it did not have a user-centric focus. It was not a user-enabling environment.

Peer-to-Peer Computing: Sharing Resources

- P2P computing defines a network architecture in which each computer has equivalent capabilities and responsibilities.
- In the P2P environment, every computer is a client *and* a server; there are no masters and slaves.
- P2P enables direct exchange of resources and services.
- There is no need for a central server, because any computer can function in that capacity when called on to do so.
- P2P was a decentralizing concept. Control is decentralized, with all computers functioning as equals. Content is also dispersed among the various peer computers.

Distributed Computing: Providing More Computing Power

- One of the **subsets** of the P2P model.
- *distributed computing*, where idle PCs across a network or Internet are tapped to provide computing power for large, processor-intensive projects.

Collaborative Computing: Working as a Group

- Multiple users to work simultaneously on the same computer-based project called collaborative computing.
- The goal was to enable multiple users to collaborate on group projects online, in real time. To collaborate on any project, users must first be able to talk to one another.
- Most collaboration systems offer the complete range of audio/video options, for full-featured multiple-user video conferencing.
- In addition, users must be able to share files and have multiple users work on the same document simultaneously.
- Real-time white boarding is also common, especially in corporate and education environments.

Cloud Computing: The Next Step in Collaboration

- With the growth of the Internet, there was no need to limit group collaboration to a single enterprise's network environment.
- Users from multiple locations within a corporation, and from multiple organizations, desired to collaborate on projects that crossed company and geographic boundaries.
- To do this, projects had to be housed in the “cloud” of the Internet, and accessed from any Internet-enabled location.

Distributed System

- A distributed system contains multiple nodes that are physically separate but linked together using the network.
- All the nodes in this system communicate with each other and handle processes in tandem. Each of these nodes contains a small part of the distributed operating system software.




Fig 4: Distributed Systems

Types of Distributed Systems

- The nodes in the distributed systems can be arranged in the form of client/server systems or peer to peer systems. Details about these are as follows:

Client/Server Systems

- In client server systems, the client requests a resource and the server provides that resource.
- A server may serve multiple clients at the same time while a client is in contact with only one server.
- Both the client and server usually communicate via a computer network and so they are a part of distributed systems.

Peer to Peer Systems

- The peer to peer systems contains nodes that are equal participants in data sharing.
- All the tasks are equally divided between all the nodes. The nodes interact with each other as required as share resources. This is done with the help of a network.

Advantages of Distributed Systems

- can easily share data with other nodes.
- More nodes can easily be added to the distributed system i.e. it can be scaled as required.
- Failure of one node does not lead to the failure of the entire distributed system. Other nodes can still communicate with each other.
- Resources like printers can be shared with multiple nodes rather than being restricted to just one.

Disadvantages of Distributed Systems

- It is difficult to provide adequate security in distributed systems because the nodes as well as the connections need to be secured.
- Some messages and data can be lost in the network while moving from one node to another.
- The database connected to the distributed systems is quite complicated and difficult to handle as compared to a single user system.
- Overloading may occur in the network if all the nodes of the distributed system try to send data at once.

Virtualization

- **Virtualization** is a technique, which allows to share single physical instance of an application or resource among multiple organizations or tenants (customers).
- Creating a virtual machine over existing operating system and hardware is referred as Hardware Virtualization. Virtual Machines provide an environment that is logically separated from the underlying hardware.

- The machine on which the virtual machine is created is known as host machine and virtual machine is referred as a guest machine. This virtual machine is managed by a software or firmware, which is known as hypervisor.

Hypervisor

- The hypervisor is a firmware or low-level program that acts as a Virtual Machine Manager.




Fig 5: Type 1 Hypervisor

There are two types of hypervisor:

Type 1 hypervisor executes on bare system.

Eg LynxSecure, RTS Hypervisor, Oracle VM, Sun xVM Server, VirtualLogic VLX.

- The type1 hypervisor does not have any host operating system because they are installed on a bare system.
- Type 2 hypervisor is a software interface that emulates the devices with which a system normally interacts.

Eg Containers, KVM, Microsoft Hyper V, VMWare Fusion, Virtual Server 2005 R2, Windows Virtual PC and VMWare workstation 6.0.

Building cloud computing Environments

Cloud Computing Services

The three major Cloud Computing Offerings are

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

- Infrastructure as a Service (IaaS)

SaaS (Software as a Service)

- SaaS or software as a service is a software distribution model in which applications are hosted by a vendor or service provider and made available to customers over a network (internet).

PaaS (Platform as a Service)

- Platform as a service, is referred as PaaS, it provides a platform and environment to allow developers to build applications and services.

This service is hosted in the cloud and accessed by the users via internet.

IaaS (Infrastructure as a Service)

- IaaS (Infrastructure As A Service) is one of the fundamental service model of cloud computing alongside PaaS(Platform as a Service).
- It provides access to computing resources in a virtualized environment “the cloud” on internet.
- It provides computing infrastructure like virtual server space, network connections, bandwidth, load balancers and IP addresses.
- The pool of hardware resource is extracted from multiple servers and networks usually distributed across numerous data centers.
- This provides redundancy and reliability to IaaS.




Fig 6: Cloud Environment

(PaaS) that are available for the users to build and host an application are:

- Amazon Web Services: AWS offers virtual compute, storage, networking & complete computing stacks.
- It is known for its on-demand services namely Elastic Compute Cloud (EC2) and Simple Storage Service (S3).
- Google AppEngine: Launched in 2008, it provides applications (**SaaS**) and raw hardware (IaaS).
- App Engine managed infrastructure, provides a development platform to create apps, leveraging Google's infrastructure as a hosting platform.

Microsoft Azure:

- It is also a scalable runtime environment for web & distributed applications.
- it provides additional services such as support for storage (relational data & blobs), networking, caching, content delivery & others.

Cloud Application Development

- **Cloud application**, or **cloud app**, is a **software** program where **cloud**-based and local components work together. This model relies on remote servers for processing logic that is accessed through a web browser with a continual internet connection.
- Cloud application servers typically are located in a remote data center operated by a third- party cloud services infrastructure provider. Cloud-based application tasks may encompass email, file storage and sharing, order entry, inventory management, word processing, customer relationship management (CRM), data collection, or financial accounting features.

Benefits of cloud apps

- Fast response to business needs. Cloud applications can be updated, tested and deployed quickly, providing enterprises with fast time to market and agility. This speed can lead to culture shifts in business operations.
- **Simplified operation.** Infrastructure management can be outsourced to third-party cloud providers.
- **Instant scalability.** As demand rises or falls, available capacity can be adjusted.
- **API use.** Third-party data sources and storage services can be accessed with an application programming interface (API). Cloud applications can be kept smaller by

using APIs to hand data to applications or API-based back-end services for processing or analytics computations, with the results handed back to the cloud application. Vetted APIs impose passive consistency that can speed development and yield predictable results.

- **Gradual adoption.** Refactoring legacy, on-premises applications to a cloud architecture in steps, allows components to be implemented on a gradual basis.
- **Reduced costs.** The size and scale of data centers run by major cloud infrastructure and service providers, along with competition among providers, has led to lower prices. Cloud-based applications can be less expensive to operate and maintain than equivalents on-premises installation.
- **Improved data sharing and security.** Data stored on cloud services is instantly available to authorized users. Due to their massive scale, cloud providers can hire world-class security experts and implement infrastructure security measures that typically only large enterprises can obtain. Centralized data managed by IT operations personnel is more easily backed up on a regular schedule and restored should disaster recovery become necessary.

How cloud apps work

- Data is stored and compute cycles occur in a remote data center typically operated by a third- party company. A back end ensures uptime, security and integration and supports multiple access methods.
- Cloud applications provide quick responsiveness and don't need to permanently reside on the local device. They can function offline, but can be updated online.
- While under constant control, cloud applications don't always consume storage space on a computer or communications device. Assuming a reasonably fast internet connection, a well- written cloud application offers all the interactivity of a desktop application, along with the portability of a web application.

Cloud Computing Platforms and Technologies

Amazon Web Services (AWS) –

- AWS provides different wide-ranging clouds **IaaS services**, which ranges from virtual compute, storage, and networking to complete computing stacks.

- AWS is well known for its storage and compute on demand services, named as Elastic Compute Cloud (EC2) and Simple Storage Service (S3).
- EC2 offers customizable virtual hardware to the end user which can be utilized as the base infrastructure for deploying computing systems on the cloud.
- S3 is well ordered into buckets which contain objects that are stored in binary form and can be grown with attributes. End users can store objects of any size, from basic file to full disk images and have them retrieval from anywhere.
- In addition, EC2 and S3, a wide range of services can be leveraged to build virtual computing systems including: networking support, caching system, DNS, database support, and others.



Fig 7: AWS

Google AppEngine

- Google AppEngine is a scalable runtime environment frequently dedicated to executing web applications.
- These utilize the benefits of the large computing infrastructure of Google to dynamically scale as per the demand.
- AppEngine offers both a secure execution environment and a collection of which simplifies the development of scalable and high-performance Web applications.
- These **services include**: in-memory caching, scalable data store, job queues, messaging, and cron tasks.
- Currently, the supported programming languages are Python, Java, and Go.
- Microsoft Azure is a Cloud operating system and a platform in which users can develop the applications in the cloud.
- Azure provides a set of services that support storage, networking, caching, content delivery, and others.

Hadoop

- Hadoop is an implementation of MapReduce, an application programming model which is developed by Google.
- This model provides two fundamental operations for data processing: map and reduce.

Force.com and Salesforce.com

–

- Force.com is a Cloud computing platform at which users can develop social enterprise applications.
- The platform is the basis of SalesForce.com – a Software-as-a-Service solution for customer relationship management.
- Force.com allows creating applications by composing ready-to-use blocks: a complete set of components supporting all the activities of an enterprise are available.
- From the design of the data layout to the definition of business rules and user interface is provided by Force.com as a support.
- This platform is completely hosted in the Cloud, and provides complete access to its functionalities, and those implemented in the hosted applications through Web services technologies.

Difference between Parallel Computing and Distributed Computing

Parallel Computing:

- In parallel computing multiple processors perform multiple tasks assigned to them simultaneously.
- Memory in parallel systems can either be shared or distributed. Parallel computing provides concurrency and saves time and money. Distributed Computing:
- In distributed computing we have multiple autonomous computers which seem to the user as a single system.
- In distributed systems there is no shared memory and computers communicate with each other through message passing.
- In distributed computing a single task is divided among different computers.

Difference between Parallel Computing and Distributed Computing:

S.NO	PARALLEL COMPUTING	DISTRIBUTED COMPUTING
1	Many operations are performed simultaneously	System components are located at different locations
2	Single computer is required	Uses multiple computers
3	Multiple processors perform multiple operations	Multiple computers perform multiple operations
4	It may have shared or distributed memory	It have only distributed memory
5	Processors communicate with each other through bus	Computers communicate with each other through message passing.
6	Improves the system performance	Improves system scalability, fault tolerance and resource sharing capabilities

Table 2: Parallel vs Distributed Computing

Elements of parallel computing

- The primary goal of parallel computing is to increase the computational power available to your essential applications.
- Typically, This infrastructure is where the set of processors are present on a server, or separate servers are connected to each other to solve a computational problem.
- In the earliest computer software, that executes a single instruction (having a single Central Processing Unit (CPU)) at a time that has written for serial computation. A Problem is broken down into multiple series of instructions, and that Instructions executed one after another. Only one computational instruction complete at a time.

Main Reasons to use Parallel Computing is that:

1. Save time and money.
2. Solve larger problems.

3. Provide concurrency.
4. Multiple execution units

Types of parallel computing

Bit-level parallelism

In the Bit-level parallelism every task is running on the processor level and depends on processor word size (32-bit, 64-bit, etc.) and we need to divide the maximum size of instruction into multiple series of instructions in the tasks. For Example, if we want to do an operation on 16-bit numbers in the 8-bit processor, then we would require dividing the process into two 8 bit operations.

Instruction-level parallelism (ILP)

Instruction-level parallelism (ILP) is running on the hardware level (dynamic parallelism), and it includes how many instructions executed simultaneously in a single CPU clock cycle.

Data Parallelism

The multiprocessor system can execute a single set of instructions (SIMD), data parallelism achieved when several processors simultaneously perform the same task on the separate section of the distributed data.

Task Parallelism

Task parallelism is the parallelism in which tasks are splitting up between the processors to perform at once.

Hardware architecture of parallel computing –

The hardware architecture of parallel computing is disturbed along the following categories as given below :

1. Single-instruction, single-data (SISD) systems
2. Single-instruction, multiple-data (SIMD) systems
3. Multiple-instruction, single-data (MISD) systems

Multiple-instruction, multiple-data (MIMD) systems

Part – A			
Q.No	Questions	Competence	BT Level
1	Define Cloud Computing	Remember	BT1
2	List the major Actors of Cloud Computing Reference Architecture	Remember	BT1
3	Recognize the characteristics of cloud computing	Remember	BT1
4	Differentiate client server and distributed computing	Understand	BT2
5	List the types of Hypervisors	Remember	BT1
6	Identify the available Cloud Computing Services	Remember	BT1
7	Differentiate parallel computing and distributed computing	Understand	BT2
8	Explain the types of parallel computing	Remember	BT1
9	Explain the advantages of using parallel computing	Remember	BT1

Part – B			
Q.No	Questions	Competence	BT Level
1	Explain about Cloud Computing Reference Model	Understand	BT2
2	Illustrate Virtualization and the types of hypervisors	Understand	BT2
3	Summarize the services of Cloud Computing	Understand	BT2
4	Outline the Cloud Computing Platforms and Technologies	Understand	BT2
5	Explain about Hardware Architecture of Parallel computing	Understand	BT2



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 COMPUTING
DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING

UNIT – II-CLOUD COMPUTING – SBSA1402

UNIT-II

*Introduction - The cloud reference model - Types of clouds - Economics of the cloud
Virtualization: Introduction – Characteristics of virtualized environments - Taxonomy of
virtualization techniques - Virtualization and cloud computing - Pros and cons of
virtualization - Technology example: VMware: full virtualization.*

Types of Cloud:

- Private,
- Public
- Hybrid
- Community

Public cloud:

Public cloud are managed by third parties which provide cloud services over the internet to public, these services are available as pay-as-you-go billing mode.

They offer solutions for minimizing IT infrastructure costs and act as a good option for handling peak loads on the local infrastructure. They are a good option for small enterprises, which are able to start their businesses without large upfront investments by completely relying on public infrastructure for their IT needs.

A fundamental characteristic of public clouds is multitenancy. A public cloud is meant to serve multiple users, not a single customer. A user requires a virtual computing environment that is separated, and most likely isolated, from other users.

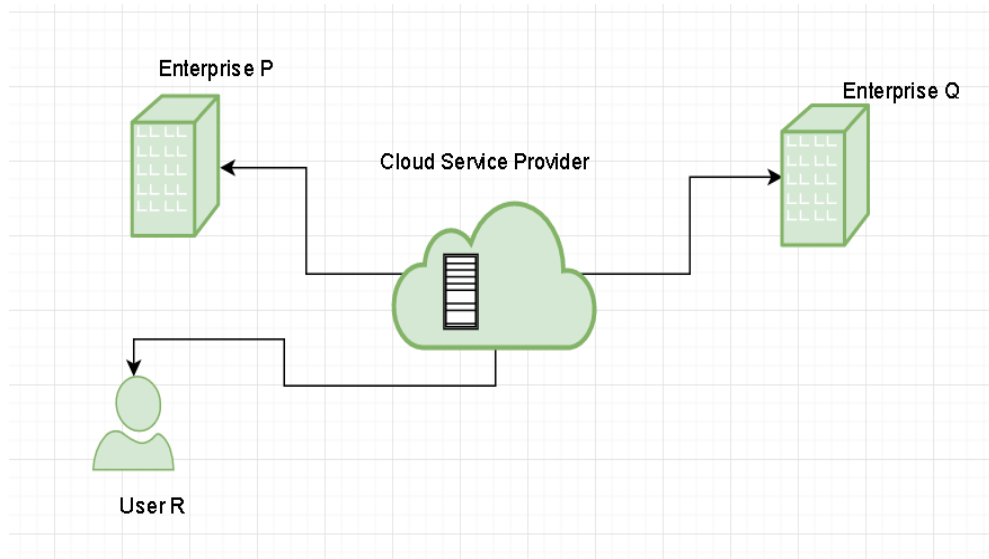


Fig 1: Cloud Service Providers

Private cloud :

Private clouds are distributed systems that work on a private infrastructure and providing the users with dynamic provisioning of computing resources. Instead of a pay-as-you-go model as in public clouds, there could be other schemes in that take into account the usage of the cloud and proportionally billing the different departments or sections of an enterprise.

The advantages of using a private cloud are:

- Customer information protection: In private cloud security concerns are less since customer data and other sensitive information does not flow out of a private infrastructure.
- Infrastructure ensuring SLAs: Private cloud provides specific operations such as appropriate clustering, data replication, system monitoring and maintenance, and disaster recovery, and other uptime services.
- Compliance with standard procedures and operations: Specific procedures have to be put in place when deploying and executing applications according to third-party compliance standards. This is not possible in case of public cloud.

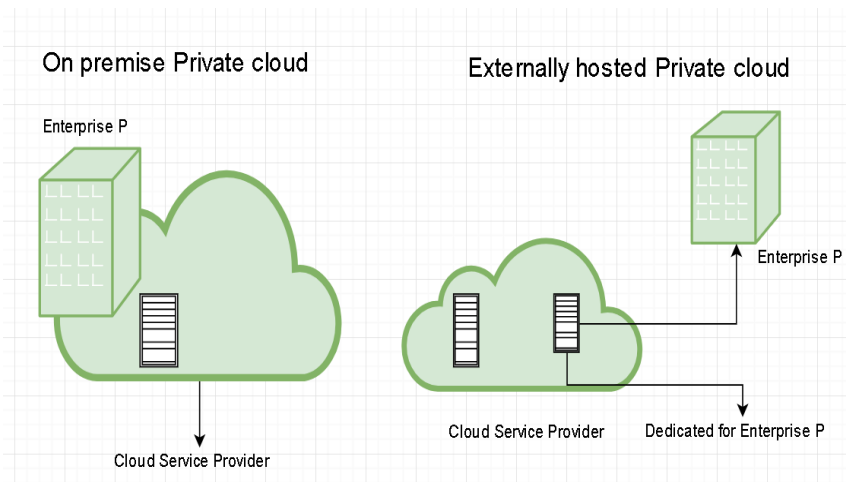


Fig 2: On-premise and Externally hosted Clouds

Hybrid cloud:

- Hybrid cloud is a heterogeneous distributed system resulted by combining facilities of public cloud and private cloud. For this reason they are also called heterogeneous clouds.
- A major drawback of private deployments is the inability to scale on demand and to efficiently address peak loads. Here public clouds are needed. Hence, a hybrid cloud takes advantages of both public and private cloud.

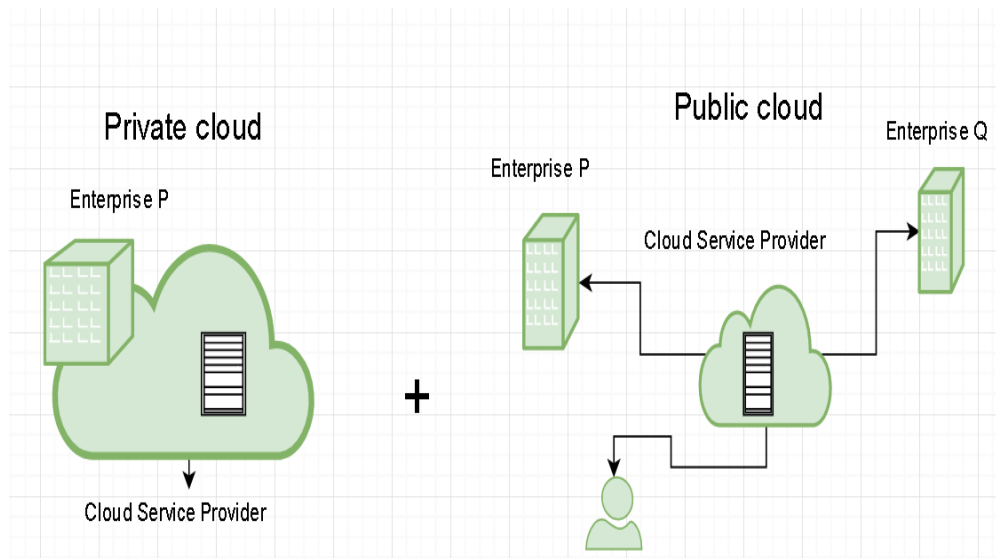


Fig 3: Public and Private Clouds

Community cloud:

- Community clouds are distributed systems created by integrating the services

of different clouds to address the specific needs of an industry, a community, or a business sector.

- In community cloud, the infrastructure is shared between organization which have shared concerns or tasks. The cloud may be managed by an organization or a third party.

Sectors that use community clouds are:

- Media industry: Media companies are looking for quick, simple, low-cost way for increasing efficiency of content generation. Most media productions involve an extended ecosystem of partners. In particular, the creation of digital content is the outcome of a collaborative process that includes movement of large data, massive compute-intensive rendering tasks, and complex workflow executions.
- Healthcare industry: In healthcare industry community clouds are used to share information and knowledge on the global level with sensitive data in the private infrastructure.
- Energy and core industry: In these sectors, the community cloud is used to cluster set of solution which collectively addresses management, deployment, and orchestration of services and operations.
- Scientific research: In this organization with common interests of science share large distributed infrastructure for scientific computing.

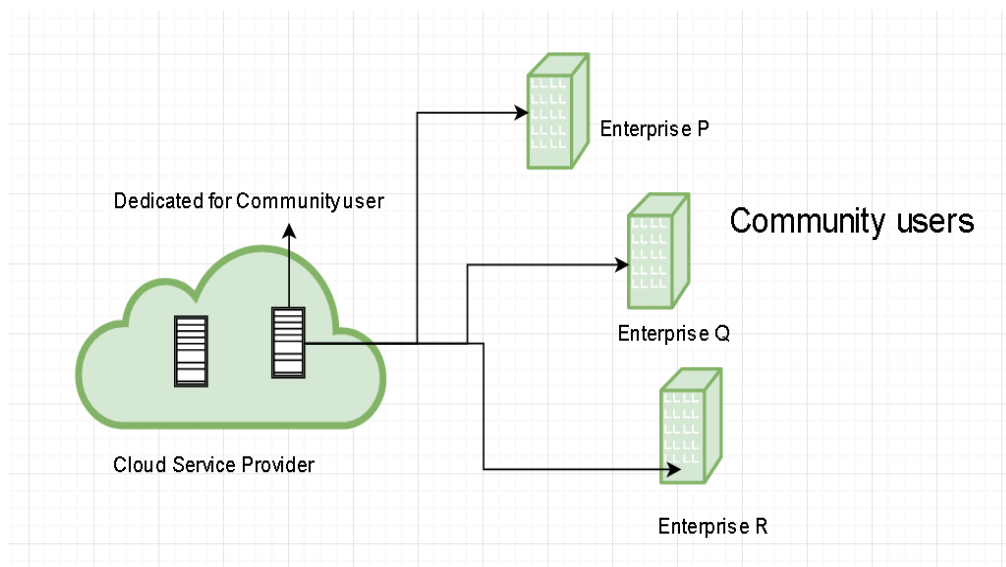


Fig 4: Community Clouds

Economics of Cloud Computing

- Economics of Cloud Computing is based on the PAY AS YOU GO method.
- Users/Customers must have to pay only for their way of usage of the cloud services.
- It is definitely beneficial for the users. So that Cloud is economically very convenient for all.
- Another side is to eliminate some indirect cost which is generated by assets such as license of software and their support.
- In the cloud, users can use software application on subscription basis without any cost because the property of the software providing service remains to the cloud provider.

Economical background of cloud is more useful for developers in the following ways:

- Pay as you go model offered by cloud providers.
- Scalable and Simple.

Cloud Computing Allows:

- Reduces the capital Costs of infrastructure. Removes the maintenance cost.
- Removes the administrative cost.

What is Capital Cost?

- The cost occurred in the purchasing infrastructure or the assets that are important in the production of goods. It takes a long time to generate profit.
- In the case of start-ups, there is no extra budget for the infrastructure and its maintenance. So cloud can minimize expenses of any small organization in terms of economy. It leads to the developers can only focus on the development logic and not on the maintenance of the infrastructure.

There are **three different Pricing** Strategies which are introduced by the Cloud Computing:

- Tiered Pricing,
- Per-unit Pricing, and
- Subscription based Pricing.

Tiered Pricing:

- Cloud Services are offered in various tiers. Each tier offers fix service agreements at specific cost. Amazon EC2 uses this kind of pricing.

Per-unit Pricing:

- The model is based upon the unit specific service concept. Data transfer and memory allocation includes in this model for specific units. GoGrid uses this kind of pricing in terms of RAM/hour.

Subscription based Pricing:

- In this model users are paying periodic subscription fees for the usage of software. So these models give more flexible solutions to the cloud economy.

Characteristics of Virtualization

- **Partitioning:**

In virtualization, many applications and operating systems (OSes) are supported in a single physical system by *partitioning* (separating) the available resources.

- **Isolation:**

Each virtual machine is isolated from its host physical system and other virtualized machines. Because of this isolation, if one virtual-instance crashes, it doesn't affect the other virtual machines. In addition, data isn't shared between one virtual container and another.

- **Encapsulation:**

A virtual machine can be represented (and even stored) as a single file, so you can identify it easily based on the service it provides. In essence, the encapsulated process could be a business service. This encapsulated virtual machine can be presented to an application as a complete entity. Therefore, encapsulation can protect each application so that it doesn't interfere with another application.

Taxonomy of virtualization techniques

Five Types of Virtualization

Hardware Virtualization

- Normally, a server devotes complete control of its hardware resources (cpu, RAM, and storage) to the actions of a single operating system.
- When you virtualize your hardware, it means that a program called a **hypervisor** manages the hardware's resources and divides them among different isolated

operating systems, referred to as “virtual machines.”



Fig 5: Vitual Environment

Advantage:

It prevents a single application crash running on the server.

Portability - With modern hypervisors an IT team can migrate running virtual machines from one server to another - allowing them to do updates and maintenance without ever having downtime for the application.

Desktop Virtualization

- Desktop Virtualization separates the desktop environment from the physical device that is used to access it.
- Advantage
- The desktops are accessed remotely
- This creates a lot of flexibility for employees to work from home

Application virtualization

- The user will interact with the native application on the client device using tablets or smartphones.

Advantage

Efficiency- In many situations where remote applications are useful, applications are redundant across employees.



Fig 6: Application Virtualization

Storage Virtualization

- Storage virtualization improves storage flexibility by creating a unified virtual pool of storage from physical storage devices in a network.
- It presents all physical storage in a cluster as a single shared group - visible to all servers.



Fig 7: Storage Virtualization

Network Virtualization

- Network virtualization pools resources from all physical networking equipment and presents it to virtual machines and applications as a single virtual network.

Advantage

It reduces the provisioning time for new network architectures.

Pros and cons of virtualization The Advantages of Virtualization

- It is cheaper.
- Virtualization doesn't require actual hardware components to be used or installed; IT infrastructures find it to be a cheaper system to implement.
- It keeps costs predictable.

- Because third-party providers typically provide virtualization options, individuals and corporations can have predictable costs for their information technology needs.
- For example: the cost of a Dell PowerEdge T330 Tower Server, at the time of writing, is \$1,279 direct from the manufacturer. In comparison, services provided by Bluehost Web Hosting can be as slow as \$2.95 per month.

It reduces the workload.

- Most virtualization providers automatically update their hardware and software that will be utilized. Instead of sending people to do these updates locally, they are installed by the third-party provider.
- This allows local IT professionals to focus on other tasks and saves even more money for individuals or corporations.

It offers a better uptime.

- Uptime has improved dramatically. Some providers offer an uptime that is 99.9999%. Even budget-friendly providers offer uptime at 99.99% today.

It allows for faster deployment of resources.

- Resource provisioning is fast and simple when virtualization is being used.
- There is no longer a need to set up physical machines, create local networks, or install other information technology components.
- As long as there is at least one point of access to the virtual environment, it can be spread to the rest of the organization.

It promotes digital entrepreneurship.

- Before virtualization occurred on a large scale, digital entrepreneurship was virtually impossible for the average person.

It provides energy savings.

- For most individuals and corporations, virtualization is an energy-efficient system. Because there aren't local hardware or software options being utilized, energy consumption rates can be lowered.
- Instead of paying for the cooling costs of a data centre and the operational costs of

equipment, funds can be used for other operational expenditures over time to improve virtualization's overall ROI.

The Disadvantages of Virtualization

It can have a high cost of implementation.

- The cost for the average individual or business when virtualization is being considered will be quite low.
- For the providers of a virtualization environment, however, the implementation costs can be quite high.
- Hardware and software are required at some point and that means devices must be developed, manufactured, or purchased for implementation.

It still has limitations.

- Not every application or server is going to work within an environment of virtualization.
- That means an individual or corporation may require a hybrid system to function properly.
- This still saves time and money in the long run, but since not every vendor supports virtualization and some may stop supporting it after initially starting it, there is always a level of uncertainty when fully implementing this type of system.

It creates a security risk.

- Information is our modern currency. If you have it, you can make money. If you don't have it, you'll be ignored.
- Because data is crucial to the success of a business, it is targeted frequently.
- The average cost of a data security breach in 2017, according to a report published by the Ponemon Institute, was \$3.62 million

It creates an availability issue.

- The primary concern that many have with virtualization is what will happen to their work should their assets not be available.
- If an organization cannot connect to their data for an extended period of time, they will struggle to compete in their industry. And, since availability is controlled by third-party

providers, the ability to stay connected is not in one's control with virtualization.

It creates a scalability issue.

- Although you can grow a business or opportunity quickly because of virtualization, you may not be able to become as large as you'd like.
- You may also be required to be larger than you want to be when first starting out. Because many entities share the same resources, growth creates lag within a virtualization network. One large presence can take resources away from several smaller businesses and there would be nothing anyone could do about it.

It requires several links in a chain that must work together cohesively.

- If you have local equipment, then you are in full control of what you can do. With virtualization, you lose that control because several links must work together to perform the same task.

It takes time.

- Although you save time during the implementation phases of virtualization, it costs users time over the long-run when compared to local systems.
- That is because there are extra steps that must be followed to generate the desired result.

Types of Clouds

- There are four different cloud models that you can subscribe according to business needs:

Private Cloud:

- Computing resources are deployed for one particular organization. This method is more used for intra-business interactions.
- The computing resources can be governed, owned and operated by the same organization.

Community Cloud:

- Computing resources are provided for a community and organizations.

Public Cloud:

- Computing resource is owned, governed and operated by government, an academic or business organization.

Hybrid Cloud:

- This type of cloud can be used for both type of interactions - B2B (Business to Business)

or B2C(Business to Consumer).

- This deployment method is called hybrid cloud as the computing resources are bound together by different clouds.

Full Virtualization

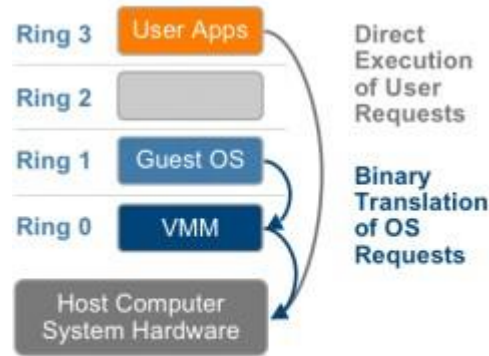


Figure 5 – The binary translation approach to x86 virtualization

Fig 8: Full Virtualization

using Binary Translation VMware can virtualize any x86 operating system using a combination of binary translation and direct execution techniques.

- Meanwhile, user level code is directly executed on the processor for high performance virtualization.
- Each virtual machine monitor provides each Virtual Machine with all the services of the physical system, including a virtual BIOS, virtual devices and virtualized memory management.
- This combination of binary translation and direct execution provides Full Virtualization as the guest OS is fully abstracted (completely decoupled) from the underlying hardware by the virtualization layer.
- The guest OS is not aware it is being virtualized and requires no modification.
- Full virtualization is the only option that requires no hardware assist or operating system assist to virtualize sensitive and privileged instructions.
- The hypervisor translates all operating system instructions on the fly and caches the results for future use, while user level instructions run unmodified at native speed.

- Full virtualization offers the best isolation and security for virtual machines, and simplifies migration and portability as the same guest OS instance can run virtualized or on native hardware.

Part – A			
Q.No	Questions	Competence	BT Level
1	List the advantages of using Private Cloud	Remember	BT1
2	Compare Public and Private Cloud	Understand	BT2
3	Show the characteristics of Virtualization	Remember	BT1
4	List types different of Virtualization	Remember	BT1
5	Show the disadvantages of virtualization	Remember	BT1
6	Mention the types of cloud models	Understand	BT2
7	Define Capital costs	Remember	BT1
8	Tell the Pricing strategies in cloud computing	Remember	BT1
9	Define the Economics of Cloud Computing	Remember	BT1

Part – B			
Q.No	Questions	Competence	BT Level
1	Explain about Full Virtualization with an example	Understand	BT2
2	Mention the types of clouds with an example	Understand	BT2
3	Describe about Type 1 and Type 2 hypervisor	Understand	BT2
4	Describe more on Economics of Cloud Computing	Understand	BT2



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 COMPUTING
DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING

UNIT- III - CLOUD COMPUTING – SBSA1402

UNIT - III

Anatomy of the Aneka container - Introducing parallelism for single-machine computation - Programming applications with threads - Multithreading with aneka - Programming applications with aneka threads. Cloud computing economics: Cloud infrastructure - Economics of private clouds - Software productivity in the cloud - Economies of scale: public vs. private clouds.

Anatomy of the Aneka container

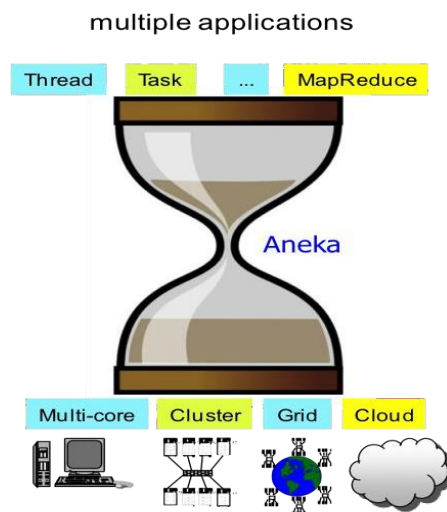


Fig 1: Anatomy of Aneka Containers

Aneka Architecture

- Aneka is a platform and a framework for developing distributed applications on the Cloud.
- It harnesses the spare CPU cycles of a heterogeneous network of desktop PCs and servers or datacenters on demand.
- Aneka provides developers with a rich set of APIs for transparently exploiting such resources and expressing the business logic of applications by using the preferred programming abstractions.
- System administrators can leverage on a collection of tools to monitor and control the deployed infrastructure. This can be a public cloud available to anyone through the Internet, or a private cloud constituted by a set of nodes with restricted access.
- The Aneka based computing cloud is a collection of physical and virtualized resources connected through a network, which are either the Internet or a private intranet.

- Each of these resources hosts an instance of the Aneka Container representing the runtime environment where the distributed applications are executed.

- The container provides the basic management features of the single node and leverages all the other operations on the services that it is hosting.

The services are broken up into

- Fabric,

- Foundation, and

- Execution services.

- Fabric services directly interact with the node through the Platform Abstraction Layer (PAL) and perform hardware profiling and dynamic resource provisioning.

- Foundation services identify the core system of the Aneka middleware, providing a set of basic features to enable Aneka containers to perform specialized and specific sets of tasks.

- Execution services directly deal with the scheduling and execution of applications in the Cloud.

- One of the key features of Aneka is the ability of providing different ways for expressing distributed applications by offering different programming models;

- Execution services are mostly concerned with providing the middleware with an implementation for these models.

- Additional services such as persistence and security are transversal to the entire stack of services that are hosted by the Container.

- At the application level, a set of different components and tools are provided to:

- 1) Simplify the development of applications (SDK);

- 2) Porting existing applications to the Cloud; and

- 3) Monitoring and managing the Aneka Cloud.

A common deployment of Aneka is presented at the side. An Aneka based Cloud is constituted by a set of interconnected resources that are dynamically modified according to the user needs by using resource virtualization or by harnessing the spare CPU cycles of desktop machines.

If the deployment identifies a private Cloud all the resources are in house, for example within the enterprise. This deployment is extended by adding publicly available resources on demand or by interacting with other Aneka public clouds providing computing resources connected over the Internet.

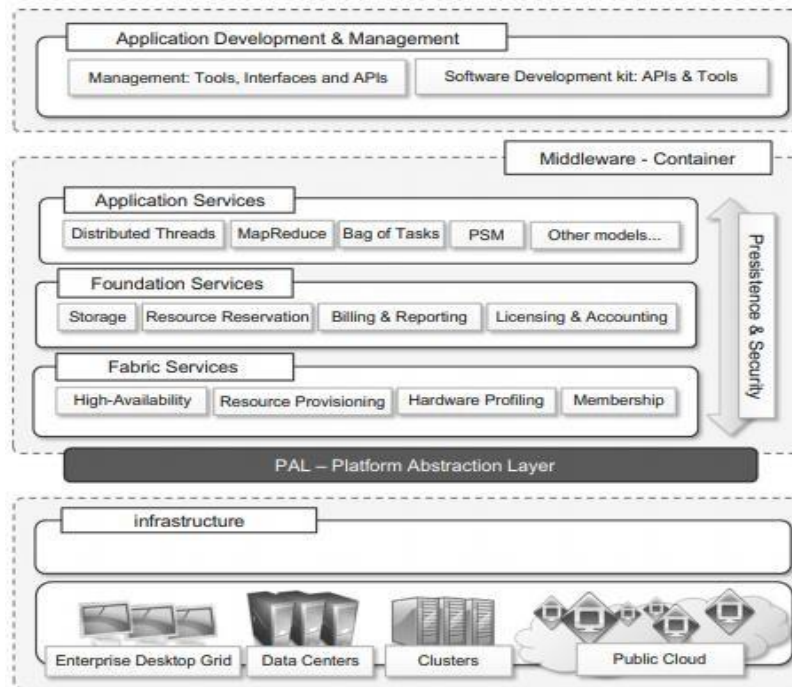


Fig 2: Architecture of Aneka

Parallelism for single machine computation

- Parallelism has been a technique for improving the performance of computers since the early 1960's, when Burroughs Corporation designed the D825, the first MIMD multiprocessor ever produced.
- From there on, a variety of parallel strategies have been developed.
- In particular, multiprocessing, which is the use of multiple processing units within a single machine, has gained a good deal of interest and gave birth to several parallel architectures.
- One of the most important distinctions is made in terms of the symmetry of processing units.
- Asymmetric multi-processing involves the concurrent use of different processing units that are specialized to perform different functions.
- Symmetric multiprocessing features the use of similar or identical processing units to share the computation load.
- Other examples, are non uniform memory access (NUMA) and clustered multi processing, which, respectively, define a specific architecture for accessing a shared memory between processors and the use of multiple computers joined together as a single virtual computer.

- Symmetric and asymmetric multiprocessing are the techniques used to increase the performance of commodity computer hardware.

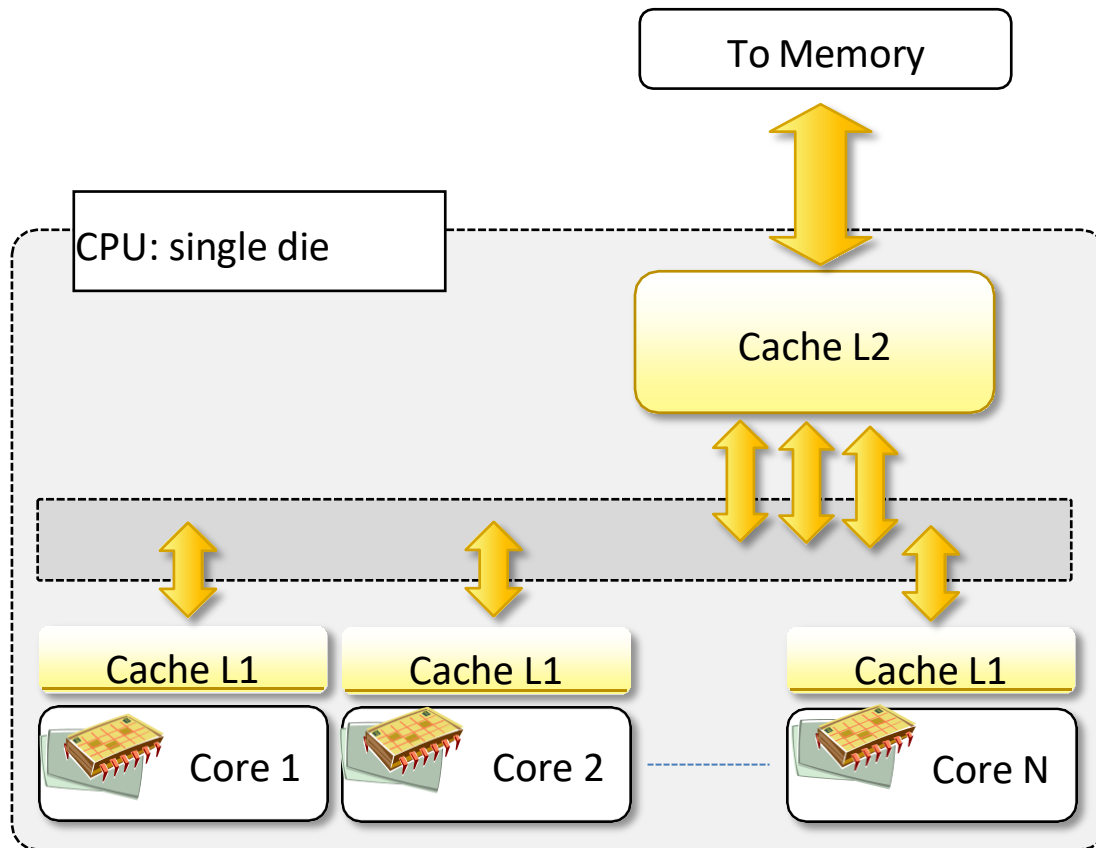


Fig 3: Core usage of Aneka

- The introduction of graphical processing units(GPUs), which are de facto processors, is an application of asymmetric processing, whereas multi core technology is the latest evolution of symmetric multiprocessing.
- Multiprocessor and especially multi core technologies are now of fundamental importance because of the physical constraint imposed on frequency scaling, which has been the common practice for performance gain in recent years.
- It became no longer possible to increase the frequency of the processor clock without paying in terms of power consumption and cooling, and this condition became unsustainable in May 2004, when Intel officially cancelled the development of two new microprocessors in favor of multi core development.
- This date is generally considered the end of the frequency-scaling era and the beginning of multi core technology. Other issues also determined the end of frequency scaling, such as the

continuously increasing gap between processor and memory speeds and the difficulty of increasing the instruction-level parallelism in order to keep a single high-performance core busy.

- Multicore systems are composed of a single processor that features multiple processing cores that share the memory.
- Each core has generally its own L1 cache, and the L2 cache is common to all the cores, which connect to it by means of a shared bus.
- Dual- and quad-core configurations are quite popular nowadays and constitute the standard hardware configuration for commodity computers.
- Architectures with multiple cores are also available but are not designed for the commodity market.
- Multicore technology has been used not only as a support for processor design but also in other devices, such as GPUs and network devices, thus becoming a standard practice for improving performance.

Multitasking and multithreading can be implemented on top of computer hardware that is constituted of a single processor and a single core, as was the common practice before the introduction of multi core technology.

In this case, the operating system gives the illusion of concurrent execution by interleaving the execution of instructions of different processes, and of different threads within the same process. This is also the case in multiprocessor/multicore systems, since the number of threads or processes is higher than the number of processors or cores.

Nowadays, almost all the commonly used operating systems support multitasking and multithreading.

Moreover, all the mainstream programming languages incorporate the abstractions of process and thread within their APIs, whereas direct support of multiple processors and cores for developers is very limited and often reduced and confined to specific libraries, which are available for a subset of the programming languages such as C/C++.

Programming Applications with Threads

- Modern applications perform multiple operations at the same time. Developers organize programs in terms of threads in order to express intra process concurrency.
- The use of threads might be implicit or explicit .

- Implicit threading happens when the underlying APIs use internal threads to perform specific tasks supporting the execution of applications such as graphical user interface (GUI) rendering, or garbage collection in the case of virtual machine-based languages.
- Explicit threading, is characterized by the use of threads within a program by application developers, who use this abstraction to introduce parallelism.

Thread:

- A thread identifies a single control flow, which is a logical sequence of instructions, within a process.
- By logical sequence of instructions, we mean a sequence of instructions that have been designed to be executed one after the other one.
- More commonly, a thread identifies a kind of yarn, that is used for sewing, and the feeling of continuity that is expressed by the interlocked fibers of that yarn is used to recall the concept that the instructions of thread express a logically continuous sequence of operations.
- Operating systems that supports multithreading identify threads as the minimal building blocks for expressing running code.
- This means that, despite their explicit use by developers, any sequence of instruction that is executed by the operating system is within the context of a thread.
- As a consequence, each process contains at least one thread but, in several cases, is composed of many threads having variable lifetimes. Threads within the same process share the memory space and the execution context; besides this, there is no substantial difference between threads belonging to different processes.
- In a multitasking environment the operating system assigns different time slices to each process and interleaves their execution.
- The process of temporarily stopping the execution of one process, saving all the information in the registers (and in general the state of the CPU in order to restore it later), and replacing it with the information related to another process is known as a context switch.
- This operation is generally considered demanding, and the use of multithreading minimizes the latency imposed by context switches, thus allowing the execution of multiple tasks in a lighter fashion.

- The state representing the execution of a thread is minimal compared to the one describing a process.

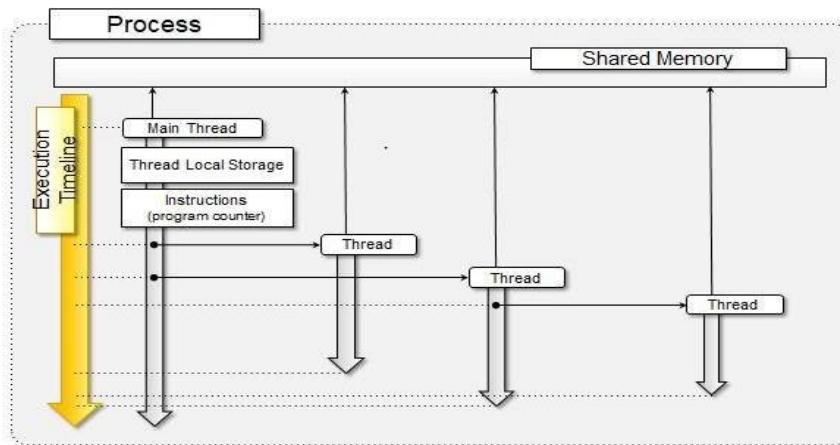


Fig 4: Threading concepts

Multithreading with Aneka

As applications become increasingly complex, there is greater demand for computational power that can be delivered by a single multi-core machine.

- Often, this demand cannot be addressed with the computing capacity of a single machine.
- It is then necessary to leverage distributed infrastructures such as Clouds.
- Decomposition techniques can be applied to partition a given application into several units of work that, rather than being executed as threads on a single node, can be submitted for execution by leveraging Clouds.

Even though a distributed facility can dramatically increase the degree of parallelism of applications, its use comes with a cost in term of application design and performance.

- For example, since the different units of work are not executing within the same process space but on different nodes both the code and the data needs to be moved to a different execution context.
- The same happens for results that need to be collected remotely and brought back to the master process.
- Moreover, if there is any communication among the different workers it is necessary to redesign the communication model eventually by leveraging the APIs provided by the middleware if any.
- In other words, the transition execution to a distributed from a single execution is process multi-threaded not transparent and application redesign and re-implementation are often required.

The amount of effort required to convert an application often depends on the facilities offered by the middleware managing the distributed infrastructure.

- Aneka, as a middleware for managing clusters, Grids, and Clouds, provides developers with advanced capabilities for implementing distributed applications.
- In particular, it takes traditional thread programming a step further. It lets you write multi-threaded applications in the traditional way, with the added twist that each of these threads can now be executed outside the parent process and on a separate machine.
- In reality, these “threads” are independent processes executing on different nodes, and do not share memory or other resources, but they allow you to write applications using the same thread constructs for concurrency and synchronization as with traditional threads.
- Aneka threads, as they are called, let you easily port existing multi-threaded compute intensive applications to distributed versions that can run faster by utilizing multiple machines simultaneously, with a minimum conversion effort.

Economics of cloud computing

Cloud computing economics depends on four customer population metrics:

1. Number of Unique Customer Sets
2. Customer Set Duty Cycles
3. Relative Duty Cycle Displacement
4. Customer Set Load

These above metrics allow the cloud-providers to use less IT resources and obtain maximum IT resource.

Proper balancing & handling of these above resources can save up to 30% of the IT resources. For economic planning, Booz Allen Hamilton, a leading strategy, and technology consulting firm, has a clear understanding of and models for effective Cloud Computing-based life-cycle cost and economic modeling.

This firm addresses every aspect of cost related efforts & questions such as:

- If you are migrating current systems to a cloud, how will you handle (and cost) the migration?
- Will you migrate all IT related tasks into the cloud, or it will be partially migrated?
- What IT chores should remain in the current environment?

- Will you use a public cloud, or you will need a private cloud?
 - If you are migrating IT requirements into a cloud, how will you handle and budget for the short-term operations of distributed and cloud infrastructures?
 - What can special-purpose computing tasks be enabled in the cloud technology that wasn't possible in the current environment?
 - Can the cloud provide levels of service commensurate with existing service-level agreements?
- Cloud technology provides users with strong facilities & economic incentives. The selection for implementing the private, public, community or hybrid cloud solely depends on the customer's specification for applications they want to use, the performance they need, the security they want to take & compliance requirement. Proper deployment model can save monetary value as well as time in a significant manner, provide better IT services & provide a higher level of reliability.

Achieve Cloud Economics for Operation & Services:

Users get frustrated when they need to secure their resources to business requirements. Even if they're willing to pay the attributed cost, they may find that technology cannot deliver resources to address their needs. Frustrated by the unavailability of resources required to fulfill their needs, and for that, they opt for the options to obtain resources are as follows:

To obtain sufficient resources for enabling the business opportunities

Hire a senior IT manager or an IT management team to put their request. This may help in letting the business proceed forward with the full initiative. But doing this fosters disrespect for the existing cloud process.

- Request to the senior corporate manager for raising the technology and resources (in case of cloud-security, or cloud-infrastructure).

Economic Characteristics of Cloud

- Scalability: Access to unlimited computer resources without thinking about the economic aspects. This feature needs planning & provisioning.
- Low Entry barrier: Users can gain access to systems for small investments also; which allows the offer to access global resources to small ventures.

- Flexibility: Cloud provides high economic elasticity. Users can re-size their resources based on their need. This feature allows optimizing the system & captures all possible requirements.
- Utility: As cloud providers 'pay-as-you-go' model, users can match their needs & resources on an ongoing basis. This eliminates waste and added benefits of shifting risks from the client.

Cloud infrastructure:

It refers to the software along with the hardware components such as storage drive, hardware, servers, virtual software, other cloud management software, and other networking devices; All work together to support the computing requirement of the cloud computing model. Moreover, the cloud technology holds a software abstraction layer that virtualizes the cloud resource & presents them to users locally.

Cloud Infrastructure Management Interface (CIMI) is an open standard API that is used to manage the cloud infrastructure. It enables its users to handle the entire cloud infrastructure easily by providing a means to interact with the provider & their consumer or developer.

The hypervisor can be defined as the firmware (a permanent set of instruction or code programmed into the read-only memory & is a low-level program) that acts as a manager for the virtual machine. It is also called Virtual Machine Monitor (VMM) which creates & runs the virtual machine. It provides the guest OS with a virtual operating platform to manages the execution of other applications.

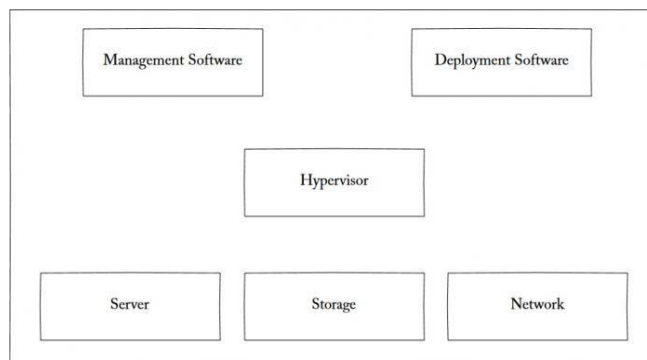


Fig 5: Hypervisor and Interactions

Types of Hypervisor:

- Native Hypervisor
- Hosted Hypervisor

In cloud technology, virtualized resources are kept & maintained by the service provider or the department of IT; these resources comprise of servers, memory, network switches, firewalls, load- balancers & storage. In the cloud computing architecture, the cloud infrastructure referred to the back-end components of the cloud.

Management Software firstly helps to configure the infrastructure then maintaining it. The Deployment software, on the other hand, is used to deploy & combine all applications on the cloud.

Network, as we all know is the key part of cloud technology allowing users to connect to the cloud via the internet. Multiple copies of data are kept stored in the cloud. This is because, if any storage resource fails - then the data can be extracted from another one. So, storage is another essential component of cloud infrastructure.

Server helps to handle & compute all cloud resources & offer services like allocation, de-allocation, sharing and monitoring the cloud resources and in some cases used to provide security.

Restrictions or Limitations of Cloud Infrastructures

The limitations of cloud technology concerning infrastructure are:

- Scalability
- Intelligent Monitoring
- Security

Software Productivity in the Cloud**Cloud-Based Software's Advantage**

- Better productivity
- Efficiency for employees
- Lower spending on hardware
- Reduced energy costs

- Flexibility
- Scalability and future focus

Disadvantages of Cloud-Based Productivity Software

- 1)Downtime
- 2)Security
- 3)Cost

Google Office Apps for Business

•Google is offering the best cloud-based productivity software names as Google Office Applications. When you talk about Cloud software, it is impossible for Google to be left behind. With Google drive, you will be able to avail of around 15 GB of storage (30GB paid) which is available for free Google online

Microsoft Office 365

- Office 365, one of the well-known cloud-based productivity software from Microsoft.
- With Microsoft Office, you will be able to keep everything organized and up-to-date.
- One will be able to keep polished documents through the use of Microsoft Office. The best part of Office remains that it is accessible across devices.
- You will also get to have free training with each application of Microsoft Office.
- It is easy, smart and very creative, compatibility with Windows 10.
- Microsoft Office is famous for its use because of its usage for teamwork.
- It is one of the most used Office online features which is built-in with the setup of Office blog.
- It is a perfectly able, compatible set of Docs.

Zoho Office Online Suite

- Zoho is right behind Microsoft and Google in cloud-based productivity software list.
- This name has evolved and achieved its standing in the market thanks to its many features which lets it do so.
- It helps you with your daily activities, boost sales records and increase employee productivity.
- It also has features of campaigns, connects and bug tracker available with it.

- It allows you to see show times, set up documents or survey forms for research or even personalize your contacts to the largest extent.

Apple iWork for iCloud

- Apple launched iWork, a cloud-based productivity software access with the iCloud account.
- Apple will stay back and not invent its inventory of Office online products under the tagline 'I'.
- Some of its features are Keynote for office users, Libre office project, iScript, and Automation features.
- One of the best features according to many experts is the iPhoto feature which is unique and sets it apart from the rest.
- The only downside is that Apple created its own echo system, which allows only their own device owners to enjoy the benefits of these suits.
- If you don't have free Apple / iCloud ID, you can get one free.

Part – A			
Q.No	Questions	Competence	BT Level
1	List the Economic characteristics of cloud	Remember	BT1
2	Show the Limitations of cloud infrastructure	Remember	BT1
3	Recognize the Advantages of cloud based software	Remember	BT1
4	Compare Public and private cloud	Understand	BT2
5	Identify the Factors for determining the total cost of an ownership	Understand	BT2
6	Define thread	Remember	BT1
7	Define Multithreading in Aneka	Remember	BT1
8	Define container	Remember	BT1
9	Explain the Anatomy of Aneka container	Understand	BT2

Part – B			
Q.No	Questions	Competence	BT Level
1	Describe about Aneka Architecture in detail	Understand	BT2
2	Demonstrate the Programming application with threads	Apply	BT3
3	Explain about Multithreading in Aneka	Understand	BT2
4	Illustrate about Cloud Infrastructure management interface	Understand	BT2
5	Outline about the Software productivity application in cloud with an example	Understand	BT2
6	Summarize the Parallelism for single machine computation	Understand	BT2



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 COMPUTING
DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING

UNIT- IV - CLOUD COMPUTING – SBSA1402

UNIT - IV

Multi-entity support - Multi-schema approach - Multi-tenancy using cloud data stores - Data access control for enterprise applications. Data in the cloud: Relational databases - Cloud file systems: GFS and HDFS - BigTable, HBase - Cloud data stores: Datastore and SimpleDB

Multi-entity support

A **multi-tenant cloud** is a **cloud computing** architecture that allows customers to share **computing** resources in a public or private **cloud**. Each **tenant's** data is isolated and remains invisible to other **tenants**. In a **multi-tenant cloud** system, users have individualized space for storing their projects and data.

Benefits of multitenancy

Better use of resources: One machine reserved for one tenant isn't efficient, as that one tenant is not likely to use all of the machine's computing power. By sharing machines among multiple tenants, use of available resources is maximized.

Lower costs: With multiple customers sharing resources, a cloud vendor can offer their services to many customers at a much lower cost than if each customer required their own dedicated infrastructure.

Drawbacks of multitenancy

Possible security risks and compliance issues: Some companies may not be able to store data within shared infrastructure, no matter how secure, due to regulatory requirements. Additionally, security problems or corrupted data from one tenant could spread to other tenants on the same machine, although this is extremely rare and shouldn't occur if the cloud vendor has configured their infrastructure correctly. These security risks are somewhat mitigated by the fact that cloud vendors typically are able to invest more in their security than individual businesses can.

The "noisy neighbor" effect: If one tenant is using an inordinate amount of computing power, this could slow down performance for the other tenants. Again, this should not occur if the cloud vendor has set up their infrastructure correctly.

Multitenancy in public cloud computing

This is similar to the way many public cloud providers implement multitenancy. Most cloud providers define multitenancy as a shared software instance. They store metadata* about each tenant and use this data to alter the software instance at runtime to fit each tenant's needs. The tenants are isolated from each other via

permissions. Even though they all share the same software instance, they each use and experience the software differently.

Multitenancy in private cloud computing

Private cloud computing uses multitenant architecture in much the same way that public cloud computing does. The difference is that the other tenants are not from external organizations. In public cloud computing, Company A shares infrastructure with Company B. In private cloud computing, different teams within Company A share infrastructure with each other.

Multitenancy is often used in [cloud computing](#), to offer shared tenancy on [public cloud](#) providers like [Amazon Web Services](#) and [Microsoft Azure](#). Additionally, multitenancy is a key part of another cloud model, [software as a service](#), and so is deployed by many [software as a service companies](#).

Development of Multi-Tenant Architecture

The concept of multitenancy actually dates back to the 1960s, when companies rented time on mainframes, which were rare and expensive. Back then it was called time sharing. Multiple customers could access the same apps at the same time, a feat only mainframes could do.

Starting in the 1990s, application service providers (ASPs) hosted applications on behalf of their customers and like mainframes, the same apps were made available to multiple customers.

In the modern era, multitenancy is part and parcel with software-as-a-service model like Salesforce, Office 365, Zoho, Box, Zendesk, Slack, and many more applications on demand. As mentioned above, cloud providers do offer multitenancy, in a technique of shared use of computing resources. However, this shared use of resources should not be confused with

virtualization, a closely related concept. In a multitenancy environment, multiple customers share the same application, in the same operating environment, on the same hardware, with the same storage mechanism. In virtualization, every application runs on a separate virtual machine with its own OS.

Single-Tenant vs. Multi-Tenant

Single-tenancy is largely seen as the "deluxe" option, in that a client operates in a solely dedicated environment. In a multi-tenant environment, each customer shares the software application along with a single database, so multiple people from the same company can access the database. Still, even in multi-tenant, each tenant is isolated from other tenants.

Advantages of Single-Tenant Hosting

Single-tenant hosting gives clients more control because there is no sharing of resources. This manifests in a number of ways:

- **Greater customization:** Since they have only client, single tenants can customize the software for their needs, whereas multi-tenant tends to be one-size-fits-all.
- **Greater isolation from security risks:** You control the environment and (hopefully) what goes in and out of it.
- **Faster recovery:** Restoring one client is faster and easier than many.
- **Better control:** Single-tenant can be choosier about accepting software changes and updates and decide what add-ons they want to use.
- **Avoiding “noisy neighbor” syndrome:** Since you are sharing resources in multi-tenant scenarios, someone else who is really heavily using the system might slow you down.

Disadvantages of Single-Tenant Hosting

Though some companies prefer it, there are downsides to single-tenancy.

- **Cost:** There is no cost sharing for things like balancing, services, system monitoring, and deployment.
- **Client responsibility:** Clients are responsible for software updates, patches, backup, restore, and disaster recovery.
- **Less efficient:** Single-tenant systems can be less efficient if they don't run at full

capacity or if they are over-provisioned.

Advantages of Multi-Tenant Hosting

The chief advantage of multi-tenant hosting is that it is less expensive. The resource pooling greatly reduces the cost since you only pay for what you need. And since multi-tenant is part of a SaaS provider, you are not paying for on- premises hardware. Functions like system monitoring and servicing the deployment become shared among all of the customers, which makes it less expensive as the cost is spread around.

There are other advantages as well.

- **Simplified hosting:** It's not your hardware to manage any more, reducing a lot of time and expense.
- **Better protection of systems:** With less interaction with the outside world, exposure to malicious software is reduced.
- **Upgrading software is no longer your problem:** You always get the latest version of software pushed out to you by the provider.

Disadvantages of Multi-Tenant Hosting

Despite its cost advantage, multi-tenant environments have some downsides.

- **They have their own security risks:** For starters, you need strict authentication and access controls to make sure the right people get access. Second, data corruption could possibly spread from one user to all, though precautions guard against this.
- **Downtime:** Outages can be nationwide, and often make the news when they happen. SaaS providers tend to build enough redundancy into the system to minimize this.
- **Noise neighbors:** As mentioned earlier, someone else on your CPU might be consuming cycles and slowing you down. Capacity is supposed to be elastic and expand as needed but that's not always the case.

Multi-Tenant Databases

As we said above, in a multi-tenant environment, multiple customers share the same application, in the same operating environment, on the same hardware, with the same storage mechanism and database. This is how Salesforce and every other SaaS operator

runs. Every tenant is a customer/user who has common access and specific privileges in the software instance.

The database, however, is another matter. There are three ways to architect your database in a multi-tenant system.

A single, shared database schema

A schema is a layout for database tables that relate to each other. In the first approach, one database is used with tenant tables all linked to the database. The tables handle relations and version control or updates, such as handling two people attempting to manipulate the same table or data entry. This is the fastest way to operate, since only one database is being used, assuming it scales.

Single database, multiple schemas

Multiple schemas inside a single database is a popular method to have sub-databases, so you can divide up your data without having to set up multiple databases. Each schema is isolated from the others and operates differently, which can be useful in situations where different data has different regulation, like international data.

Multiple databases

This takes the multi-schema approach one step further because now you have the data in multiple databases. Sales or customers can be divided up by region, for example. So the upside is you get the best isolation of data. Of course, that also adds to the complexity of management, maintenance and scalability that would come with deploying multiple databases.

Multi-Tenant Architecture Examples

In a virtual system, one system may have 20 instances running 20 operating systems, each with its own application and database. In a multi-tenant architecture, all of the instances share the OS, app, and most importantly, the database.

Does this sound familiar? That's not only how SaaS works but how Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) operate, since both can be used to write highly scalable

multitenant apps that are outward-facing, meaning customers and partners can use them.

That's how 50 people from the same company can work on Salesforce CRM. Similarly, a SAP system is composed of a database backend and Web application servers that host Web services in a highly scalable manner. The Web services that make up the SaaS app are exposed to different customers via different domain names. Scaling is achieved by starting more services.

Most of the examples conform to these three ways to structure a multi-tenant application:

- **URL-based SaaS.** This is the easiest one to do since it uses a single domain and database. You would have specific URLs, like `subdomain1.maindomain.com`, `subdomain2.maindomain.com`, etc. Data management and security is handled at the application level. Some SaaS operates this way, especially those that put a Web app interface between the user and the database.
- **Multi-tenant SaaS.** This is more complex because of multiple databases and/or schemas, and restrictions are done at the database level. This is how many SaaS apps operate and often allow more direct interaction with the database.
- **Virtualization-based SaaS, like containers.** This is the most complex setup because there is so much interaction between the containers as well as the apps and databases. That's why there is so much emphasis on orchestrators like Docker and Kubernetes.

Choosing Between Multi-Tenant and Single Tenant

Choosing between single- and multi-tenant comes down to a choice of on-premises vs. cloud. There is no single-tenant version of Salesforce and, in contrast, databases like Oracle are single-tenant so as to have full access to resources.

Security of data is clearly a concern, but that falls primarily on the shoulders of SaaS providers. They are the ones responsible for monitoring tenants and making sure there is no data bleed from one customer to another, and they do a good job of it.

The client's primary responsibility for securing the data falls to their client device. All of the major SaaS providers do offer two-factor authentication to secure logins. After that it's up to you to maintain the security of the endpoint device.

A major concern between single- and multi-tenant systems is data ownership. Larry Ellison

once called Salesforce a “roach motel of clouds,” meaning data goes in but doesn’t come out. Obviously, he was overstating the matter. It is certainly possible to get your CRM data out of Salesforce, but more likely in a CSV format than a SQL database.

Multi-tenancy is at the heart of cloud computing. It is designed to help scale up thousands of users both within an enterprise and externally as companies interact and do business. Whether it’s your Salesforce account or an app you built on AWS for customers, multi-tenancy can scale through public and private clouds and provide true economies of scale.

Multi-tenancy using cloud data stores

The cloud storage has both financial and security advantages. Financial advantages in the case that the virtual resources are cheaper than dedicated physical resources connected to a personal computer or network.

There are four main types of cloud storage:

Personal Cloud Storage: It is also called mobile cloud storage which is a subset of public cloud storage that applies to storing an individual's data in the cloud and provides access to data from anywhere. . Apple's iCloud is an example of personal cloud storage.

Public Cloud Storage: Public cloud storage is where the storage service providers and enterprises are separate and there aren't any cloud resources stored in the enterprise's data center. The enterprise data management is done by cloud storage provider.

Private Cloud Storage: A form of cloud storage where the data storage is done within the enterprise. This helps to resolve the potential risk for performance and security concerns while still providing the advantages of cloud storage.

Hybrid Cloud Storage: Hybrid cloud storage is a combination of private and public cloud storage where some data that is critical resides in the enterprise's private cloud while other data is stored and accessible from a public storage provider.

Tenant Data Management in Multi-Tenant Cloud Systems

Multi-tenancy enables the client service providers to serve multiple clients by a single application instance, isolating each tenant’s data. The benefits of this method include increased utilization of available hardware resources and improved ease of maintenance and deployment. These results are in lower overall application costs. In a multi-tenant architecture, a software application is designed virtually to partition data. Each tenant works in a virtual application

instance in which some tenants may be divided into multiple subtenants, each having its own users. As the number of tenants increases, a scalable architecture is needed. The business logic layer provides security by the access control component which should be isolated from the offered services. The database layer stores the application data. To enforce scalability of the database layer a logical hierarchical representation of the tenants and subtenants, and making a mapping to the physical storage is followed. The data access layer provides load balancing between the different data stores and stores the decision support for managing the data store. Tenants and subtenants in a multi-tenant organization can be structured hierarchically in the form of a tenant tree. At the top level of the tree is the SaaS provider, which is the root node of the tenant tree. Various tenants using the applications are located at the next level, and are the child nodes of SaaS provider. All tenants share the same parent or root node. Some tenants may be even divided into multiple subtenants.

The representation of tenants is done by logical hierarchical representation. The data could be split as in the logical representation, by creating a data store for tenants or merging smaller databases. For small applications with a less number of tenants, can share a single data store.

There are 3 different models for load balancing




Fig 1: Tenant Data Management

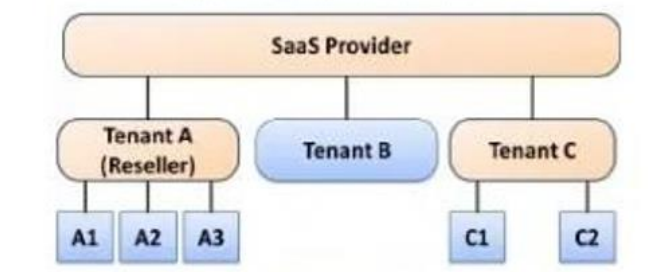


Fig 2: Interaction of SAAS with Tenants

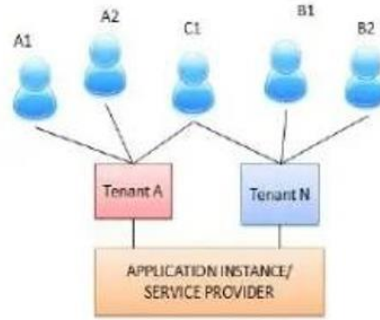


Fig 3: Application Instances

In the monolithic model, all data is stored in a single data store. In the fully distributed model every tenants and subtenants has its own data store. Whereas the hybrid model is a mixture of both monolithic and fully distributed models. For new SaaS applications with limited number of tenants, the service provider could use the monolithic model for storing data, and move to a fully distributed model or hybrid as the amount of data and the number of tenants increases. The monolithic model is easier to implement at a lower costs, as it requires only a single database instance. The fully distributed model is more expensive but it offers more scalability and flexibility.

MULTI-ENTITY SUPPORT

Long before ASPs and SaaS, large globally distributed organizations often needed their applications to support multiple organizational units, or ‘entities,’ in a segregated manner.




Fig 4: Data Storage Models

For example, consider a bank with many branches needing to transition from separate branch specific installations of its core banking software to a centralized deployment where

the same software would run on data from all branches.

The software designed to operate at the branch level clearly could not be used directly on data from all branches: For example branch-level users should see data related only to their branch and branch-wise accounting should consider transactions segregated by branch. If there was a need to enhance the system, say by introducing a new field, such a change would need to apply across all branches; at the same time, sometimes branch specific extensions would need to be supported as well. These requirements are almost exactly the same as for multitenancy!

In a multi-entity scenario there are also additional needs, such as where a subset of users needed to be given access to data from all branches, or a subset of branches, depending on their position in an organizational hierarchy. Similarly, some global processing would also need to be supported, such as inter-branch reconciliation or enterprise-level analytics, without which the benefits of centralization of data might not be realized.

Figure depicts the changes that need to be made in an application to support basic multi-entity features, so that users only access data belonging to their own units. Each database table is appended with a column (OU_ID) which marks the organizational unit each data record belongs to.

Each database query needs to be appended with a condition that ensures that data is filtered depending on the organizational unit of the currently logged-in user, which itself needs to be set in a variable, such as `current_user_OU`, during each transaction.

An exactly similar mechanism can be used to support multi-tenancy, with `OU_ID` now representing the customer to whom data records belong. Note that the application runs on a single schema containing data from all organizational units; we shall refer to this as the *single SCHEMA* model.

Many early implementations of SaaS products utilized the single schema model, especially those that built their SaaS applications from scratch. One advantage of the single schema structure is that upgrading functionality of the

Other	OU_ID
-------	-------

fields					
					North
					North
					North
					South
					South

SELECT ... FROM T WHERE OU_ID=:current_user_OU

Fig 5 : Multiple Entity Implementation

Application, say by adding a field, can be done at once for all customers. At the same time, there are **disadvantages**:

Re-engineering an existing application using the single schema approach entails significant restructuring of application code. For a complex software product, often having millions of lines of code, this cost can be prohibitive. Further, while modifications to the data model can be done for all customers at once, it becomes difficult to support customer specific extensions, such as custom fields, using a single schema structure. Meta-data describing such customizations, as well as the data in such extra fields has to be maintained separately. Further, it remains the responsibility of the application code to interpret such meta-data for, say, displaying and handling custom fields on the screen. Additionally, any queries that require, say, filtering or sorting on these custom fields become very complex to handle. Some of these issues can be seen more clearly through the example in Figure 4.2 that depicts a multi-tenant architecture using a single schema model which also supports custom fields:

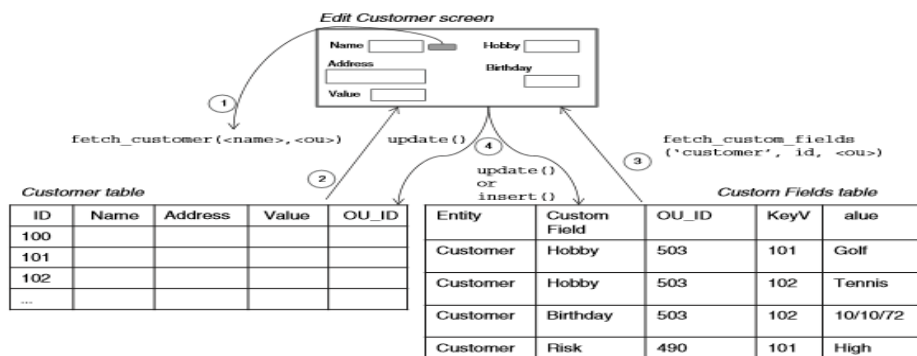


Fig 6: Custom Fields Table management

In the single schema model of above Figure Custom Fields table stores meta- Information and data values for *ALL* tables in the application. Mechanisms for handling custom fields in single schema architecture are usually variants of this scheme. Consider a screen that is used to retrieve and update records in the Customer table. First the record from the main table is retrieved by name, suitably filtered by the OU attribute of the logged in user. Next, custom fields along with their values are retrieved from the Custom Fields table, for *this* particular record *AND* the OU of the logged in user. For example, in OU503, there are two custom fields as displayed on the screen, but only one in OU

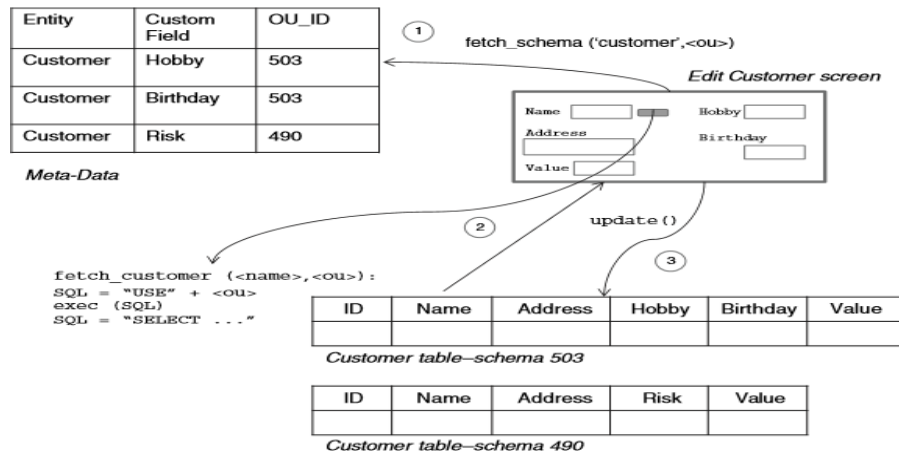


Fig 7: Multi-tenancy using Multiple Schemas

490, and none otherwise. Furthermore, some records may have missing values for these fields, so while saving the record care must be taken to appropriately either insert or update records in the Custom Fields table. The above example is a simple case; more complex requirements also need to be handled, for example where a list of records is to be displayed with the ability to sort and filter on custom fields. It should be clear from this example that the single schema approach to multi-tenancy, while seemingly having the advantage of being able to upgrade the data model in one shot for all customers, has many complicating disadvantages in addition to the fact that major re-engineering of legacy applications is needed to move to this model.

MULTI-SCHEMA APPROACH

Instead of insisting on a single schema, it is sometimes easier to modify even an existing application to use multiple schemas, as are supported by most relational databases. In this model, the application computes which OU the logged in user belongs to, and then connects to the appropriate database schema. Such architecture is

In the multiple schema approach a separate database schema is maintained for each customer, so each schema can implement customer-specific customizations directly. Meta-data describing customizations to the core schema is also maintained in a separate table, but unlike the Custom Fields table of Figure 4.2, this is pure meta-data and does not contain field values in individual records. As a result, the application design is simpler, and in case a legacy application needs to be re-engineered for multi-tenancy, it is likely that the modifications will be fewer and easier to accomplish.

Consider implementing the Edit Customer screen as discussed earlier using a multiple schema approach: The application renders the appropriate fields on the screen using information from the Meta-Data table. When making a database query, the application sets the database schema before issuing data manipulation (i.e. SQL) statements so as to access the appropriate schema. We have described a rather simple implementation to illustrate the concept of using multiple schemas for multi-tenancy. In practice, web-application servers need to have schema names configured during deployment so that they can maintain database connection pools to each schema. Therefore, another level of indirection is usually required, where customer name (i.e. OU) is mapped to the actual schema name, so that customers can be added or deleted online without bringing the system down.

In the case of a multi-entity scenario within a single organization, the number of users was relatively small, probably in the thousands at most. For a SaaS application, the number of users will be orders of magnitude larger. Thus additional factors need to be considered for a multi-tenant SaaS deployment, such as how many application server and database instances are needed, and how a large set of users are efficiently and dynamically mapped to OUs so as to be connected to the appropriate application server and database instance.

MULTI-TENANCY USING CLOUD DATA STORES

As discussed in the previous chapter, cloud data stores exhibit non-relational storage models. Furthermore, each of these data stores are built to be multi-tenant from scratch since effectively a single instance of such a large-scale distributed data store caters to multiple applications created by cloud users. For example, each user of the Google App Engine can create a fixed number of applications, and each of these *APPEARS* to have a separate data store;

Here we focus on a different problem: As a *user* (application developer) of a cloud platform, how does one create one's own multi-tenant application? In the case of Amazon EC2 the answer is straightforward; since this is an infrastructure cloud it gives users direct access to (virtual) servers where one can recreate exactly the same multi-tenant architectures discussed earlier using standard application servers and database systems.

However the situation is different using a PaaS platform such as Google's App Engine with its Datastore, Amazon's SimpleDB or even Azure's data services. For example, a single App Engine application has *one* data store name

DATA ACCESS CONTROL FOR ENTERPRISE APPLICATIONS

So far we have covered the typical strategies used to achieve multi-tenancy from the perspective of enabling a single application code base, running in a single instance, to work with data of multiple customers, thereby bringing down costs of management across a potentially large number of customers.

For the most part, multi-tenancy as discussed above appears to be of use primarily in a software as a service model. There are also certain cases where multi-tenancy can be useful within the enterprise as well. We have already seen that supporting multiple entities, such as bank branches, is essentially a multi-tenancy requirement. Similar needs can arise if a workgroup level application needs to be rolled out to many independent teams, who usually do not need to share data. Customizations of the application schema may also be needed in such scenarios, to support variations in business processes. Similar requirements also arise in supporting multiple *LEGAL* entities each of which could be operating in different regulatory environments.

As we mentioned earlier, in a multi-entity scenario a subset of users may need to be given access to data from all branches, or a subset of branches, depending on their position in an organizational unit hierarchy. More generally, access to data may need to be controlled based on the *VALUES* of any field of a table, such as high-value transactions being visible only to

some users, or special customer names being invisible without explicit permission. Such requirements are referred to as *DATA ACCESS control* needs, which while common, are less often handled in a reusable and generic manner. Data access control (or DAC) is a generalization of multi-tenancy in that the latter can often be implemented using DAC. In Figure 4.5 we illustrate how data access control can be implemented in a generic manner within a single schema to support fairly general rules for controlling access to records based on *field VALUES*.

Each application table, such as Customer, is augmented with an additional field DAC_ID. The DAC

Rules table lists patterns based on value ranges of arbitrary fields using which the values of the DAC_ID in each Customer record are filled through a batch process. Users are assigned privileges to access records satisfying one or more such DAC rules as specified in the User DAC Roles table.

This information is expanded, via a batch process, to data

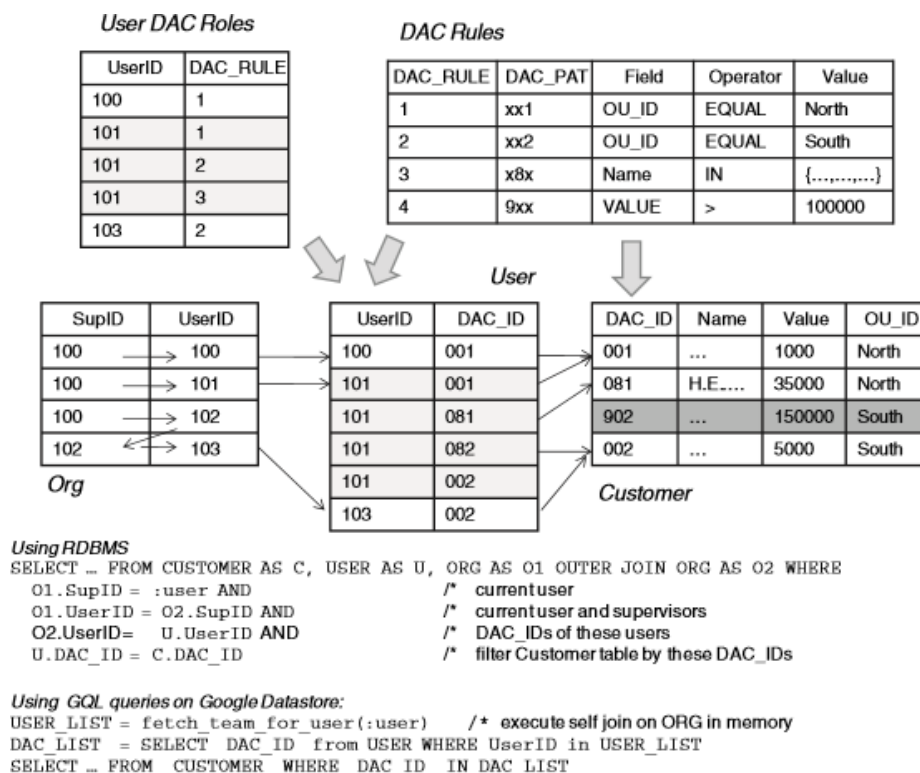


Fig 8: Data Access Control

Data in the cloud

Since the 80s relational database technology has been the ‘default’ data storage and retrieval mechanism used in the vast majority of enterprise applications. The origins of relational databases, beginning with System R and Ingres in the 70s, focused on introducing this new paradigm as a general purpose replacement for hierarchical and network databases, for the most common

business computing tasks at the time, viz. transaction processing.

In the process of creating a planetary scale web search service, Google in particular has developed a massively parallel and fault tolerant distributed file system (GFS) along with a data organization (BigTable) and programming paradigm (MapReduce) that is markedly different from the traditional relational model. Such ‘cloud data strategies’ are particularly well suited for large-volume massively parallel text processing, as well as possibly other tasks, such as enterprise analytics. The public cloud computing offerings from Google (i.e. App Engine) as well as those from other vendors have made similar data models (Google’s Datastore, Amazon’s SimpleDB) and programming paradigms (Hadoop on Amazon’s EC2) available to users as part of their cloud platforms.

At the same time there have been new advances in building specialized database organizations optimized for analytical data processing, in particular column-oriented databases such as Vertica. It is instructive to note that the BigTable-based data organization underlying cloud databases exhibits some similarities to column-oriented databases. These concurrent trends along with the ease of access to cloud platforms are witnessing a resurgence of interest in non-relational data organizations and an exploration of how these can best be leveraged for enterprise applications.

In this chapter we examine the structure of Google App Engine’s Data store and its underlying technologies, Google’s distributed file system (GFS) and BigTable abstraction, as well as the open source project Hadoop’s HBase and HDFS (clones of BigTable and GFS respectively).

RELATIONAL DATABASES

Before we delve into cloud data structures we first review traditional relational database systems and how they store data. Users (including application programs) interact with an RDBMS via SQL; the database ‘front-end’ or parser transforms queries into memory and disk level operations to optimize execution time. Data records are stored on pages of contiguous disk blocks, which are managed by the disk-space-management layer.

Pages are fetched from disk into memory buffers as they are requested, in many ways similar to the file and buffer management functions of the operating system, using pre-fetching and page

replacement policies. However, database systems usually do not rely on the file system layer of the OS and instead manage disk space themselves. This is primarily so that the database can have full control of when to retain a page in memory and when to release it. The database needs to be able to adjust page replacement policy when needed and pre-fetch pages from disk based on expected access patterns that can be very different from file operations. Finally, the operating system files used by databases need to span multiple disks so as to handle the large storage requirements of a database, by efficiently exploiting parallel I/O systems such as RAID disk arrays or multi-process or clusters.

The storage indexing layer of the database system is responsible for locating records and their organization on disk pages. Relational records (tabular rows) are stored on disk pages and accessed through indexes on specified columns, which can be B^+ -tree indexes, hash indexes, or bitmap indexes. Normally rows are stored on pages contiguously, also called a ‘row-store’, and indexed using B^+ -trees. An index can be *PRIMARY*, in which case rows of the table are physically stored in as close as possible to sorted order based on the column

specified by the index. Clearly only one such primary index is possible per table; the remaining indexes are *SECONDARY*, and maintain pointers to the actual row location on disk.

While B^+ -tree indexes on a row-store are optimal for write oriented workloads, such as the case in transaction processing applications, these are not the best for applications where reads dominate; in the latter case bitmap indexes, cross-table indexes and materialized views are used for efficient access to records and their attributes. Further, a row-oriented storage of records on disk may also not be optimal for read-dominated workloads, especially analytical applications. Recently column-oriented storage has been proposed as a more efficient mechanism suited for analytical workloads, where an aggregation of measures over columns (e.g. Sales) need to be performed based on values of dimension columns (e.g. Month). Figure 4.6 illustrates the difference between row-oriented and column-oriented storage. Notice that in a column store, projections of the table are stored sorted by dimension values, which are themselves compressed (as bitmaps, for example) for ease of comparison as well as reduced storage. Notice also that the column store needs additional ‘join indexes’ that map the sort orders of different projections so as to be able to recover the original row when required. When the cardinality of dimension values are small these join indexes can also be efficiently compressed

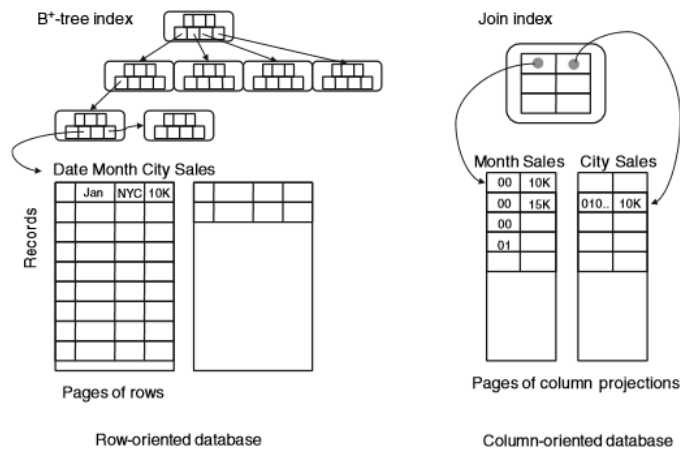


Fig 9: Row vs Columns




Fig 10: Parallel database architectures

CLOUD FILE SYSTEMS: GFS AND HDFS

The Google File System (GFS) is designed to manage relatively large files using a very large distributed cluster of commodity servers connected by a high-speed network. It is therefore designed to (a) expect and tolerate hardware failures, even during the reading or writing of an individual file (since files are expected to be very large) and (b) support parallel reads, writes and appends by multiple client programs. A common use case that is efficiently supported is that of many ‘producers’ appending to the same file in parallel, which is also being simultaneously read by many parallel ‘consumers’.

As a result they also do not scale as well as data organizations built on GFS-like platforms such as the Google Datastore. The Hadoop Distributed File System (HDFS) is an open source

implementation of the GFS architecture that is also available on the Amazon EC2 cloud platform; we refer to both GFS and HDFS as ‘cloud filesystems.’

The architecture of cloud file systems is illustrated in Figure 10.3. Large files are broken up into ‘chunks’ (GFS) or ‘blocks’ (HDFS), which are themselves large (64MB being typical). These chunks are stored on commodity (Linux) servers called Chunk Servers (GFS) or Data Nodes (HDFS); further each chunk is replicated at least three times, both on a different physical rack as well as a different network segment in anticipation of possible failures of these components apart from server failures.

When a client program (‘cloud application’) needs to read/write a file, it sends the full path and offset to the Master (GFS) which sends back meta-data




Fig 11: Cloud file systems

for one (in the case of read) or all (in the case of write) of the replicas of the chunk where this data is to be found. The client caches such meta-data so that it need not contact the Master each time. Thereafter the client directly reads data from the designated chunk server; this data is not cached since most reads are large and caching would only complicate writes.

In case of a write, in particular an append, the client sends only the data to be appended to all the chunk servers; when they all acknowledge receiving this data it informs a designated ‘primary’ chunk server, whose identity it receives (and also caches) from the Master. The primary chunk server appends its copy of data into the chunk at an offset of its choice; note that this may be beyond the EOF to account for multiple writers who may be appending to this file

simultaneously. The primary then forwards the request to all other replicas which in turn write the data at the same offset if possible or return a failure. In case of a failure the primary rewrites the data at possibly another offset and retries the process.

The Master maintains regular contact with each chunk server through heartbeat messages and in case it detects a failure its meta-data is updated to reflect this, and if required assigns a new primary for the chunks being served by a failed chunk server. Since clients cache meta-data, occasionally they will try to connect to failed chunk servers, in which case they update their meta-data from the master and retry.

It is shown that this architecture efficiently supports multiple parallel readers and writers. It also supports writing (appending) and reading the same file by parallel sets of writers and readers while maintaining a consistent view, i.e. each reader always sees the same data regardless of the replica it happens to read from. Finally, note that computational processes (the ‘client’ applications above) run on the same set of servers that files are stored on. As a result, distributed programming systems, such as MapReduce, can often schedule tasks so that their data is found locally as far as possible, as illustrated by the Clustera system [16].

BIGTABLE, HBASE

BigTable is a distributed *structured* storage system built on GFS; Hadoop’s HBase is a similar open source system that uses HDFS. A BigTable is essentially a sparse, distributed, persistent, multidimensional sorted ‘map.’¹ Data in a

In the programming language sense, i.e. a dictionary of key-value pairs.

	<i>location:city</i>	<i>location:region</i>	<i>sale: value</i>	<i>products: details</i>	<i>products: types</i>
Txn ID 0088997	NYC	US East Coast		ACME Detergent XYZ Soap KLLGS Cereal A	Cleaner Breakfast item
		US North East	\$ 80		

Fig 12: Data organization in a BigTable

BigTable is accessed by a row key, column key and a timestamp. Each column can store arbitrary name-value pairs of the form *COLUMN-FAMILY:LABEL, string*. The set of possible column-

families for a table is fixed when it is created whereas columns, i.e. labels within the column family, can be created dynamically at any time. Column families are stored close together in the distributed file system; thus the BigTable model shares elements of column-oriented databases. Further, each Bigtable cell (row, column) can contain multiple versions of the data that are stored in decreasing timestamp order. We illustrate these features below through an example.

Figure 12 illustrates the BigTable data structure: Each row stores information about a specific sale transaction and the row key is a transaction identifier. The ‘location’ column family stores columns relating to where the sale occurred, whereas the ‘product’ column family stores the actual products sold and their classification. Note that there are two values for region having different timestamps, possibly because of a reorganization of sales regions. Notice also that in this example the data happens to be stored in a de-normalized fashion, as compared to how one would possibly store it in a relational structure; for example the fact that XYZ Soap is a Cleaner is not maintained.

Since data in each column family is stored together, using this data organization results in efficient data access patterns depending on the nature of analysis: For example, only the *LOCATION* column family may be read for traditional data-cube based analysis of sales, whereas only the *product* column family is needed for say, market-basket analysis. Thus, the BigTable structure can be used in a manner similar to a column-oriented database.

CLOUD DATA STORES: DATASTORE AND SIMPLEDB

The Google and Amazon cloud services do not directly offer BigTable and Dynamo to cloud users. Using Hadoop’s HDFS and HBase, which are available as Amazon AMIs, users can set up their own BigTable-like stores on Amazon’s EC2. However, as we have seen in Chapter 5, Google and Amazon both offer simple key-value pair database stores, viz. Google App Engine’s Datastore and Amazon’s SimpleDB. Here we describe how the Google Datastore is implemented using an underlying BigTable infrastructure, as illustrated.

All entities (objects) in a Datastore reside in one BigTable table, the Entities table having *one* column family. In addition to the single Entities table there are index tables that are used to support efficient queries.

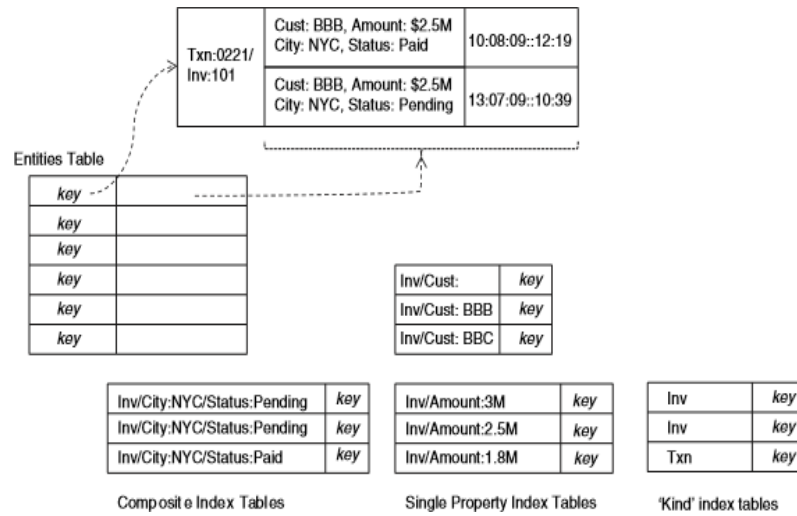


Fig 13: Google Datastore using BigTable

For the purposes of this discussion, it is useful to think of a BigTable as an array that has been horizontally partitioned (also called ‘sharded’) across disks, and *sorted* lexicographically by key values. In addition to single key lookup, this structure also enables highly efficient execution of *prefix* and *RANGE* queries on key values, e.g. all keys having prefix ‘Txn,’ or in the range ‘Amount:1M’

Part – A			
Q.No	Questions	Competence	BT Level
1	Define multitenancy	Remember	BT1
2	List the benefits of multitenancy	Remember	BT1
3	Show the types of cloud data storages	Remember	BT1
4	Compare GFS and HDFS	Understand	BT2
5	Differentiate Data store and simple database	Understand	BT2
6	Compare Single tenancy and multitenancy	Understand	BT2
7	Show the disadvantages of multitenancy	Remember	BT1

Part – B			
Q.No	Questions	Competence	BT Level
1	Explain about Multitenancy Architecture	Understand	BT2
2	Outline about Tenant Data management in multitenant cloud systems	Understand	BT2
3	Illustrate the Multi Entities support in multitenancy organization	Understand	BT2
4	Discuss more about Multi schema with multi tenancy	Understand	BT2
5	Describe about Multi tenancy using cloud data stores	Understand	BT2
6	Explain about Data access control for enterprise application	Understand	BT2
7	Describe about Parallel database architecture	Understand	BT2



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 COMPUTING
DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING

UNIT- V - CLOUD COMPUTING – SBSA1402

UNIT – V

Amazon web services: Compute services - Storage services - Communication services - Additional services. Google AppEngine: Architecture and core concepts - Application life cycle - Cost model – Observations. Microsoft azure: Azure core concepts - SQL azure - Windows azure platform appliance.

Amazon web services

Amazon Web Services (AWS) is a platform that allows the development of flexible applications by providing solutions for elastic infrastructure scalability, messaging, and data storage. The platform is accessible through SOAP or RESTful Web service interfaces and provides a Web-based console where users can handle administration and monitoring of the resources required, as well as their expenses computed on a pay-as-you-go basis.

Figure 5 shows all the services available in the AWS ecosystem. At the base of the solution stack are services that provide raw compute and raw storage: *Amazon Elastic Compute (EC2)* and *Amazon Simple Storage Service (S3)*. These are the two most popular services, which are generally complemented with other offerings for building a complete system. At the higher level, *Elastic MapReduce* and *AutoScaling* provide additional capabilities for building smarter and more elastic computing systems. On the data side, *Elastic Block Store (EBS)*, *Amazon SimpleDB*, *Amazon RDS*, and *Amazon ElastiCache* provide solutions for reliable data snapshots and the management of structured and semistructured data. Communication needs are covered at the networking level by *Amazon Virtual Private Cloud (VPC)*, *Elastic Load Balancing*, *Amazon Route 53*, and *Amazon Direct Connect*. More advanced services for connecting applications are *Amazon Simple Queue*

Service (SQS), *Amazon Simple Notification Service (SNS)*, and *Amazon Simple E-mail Service (SES)*. Other services include:

- *Amazon CloudFront* content delivery network solution
- *Amazon CloudWatch* monitoring solution for several Amazon services
- *Amazon Elastic Beanstalk* and *CloudFormation* flexible application packaging and deployment

As shown, AWS comprise a wide set of services. We discuss the most important services by examining the solutions proposed by AWS regarding compute, storage, communication, and complementary services.

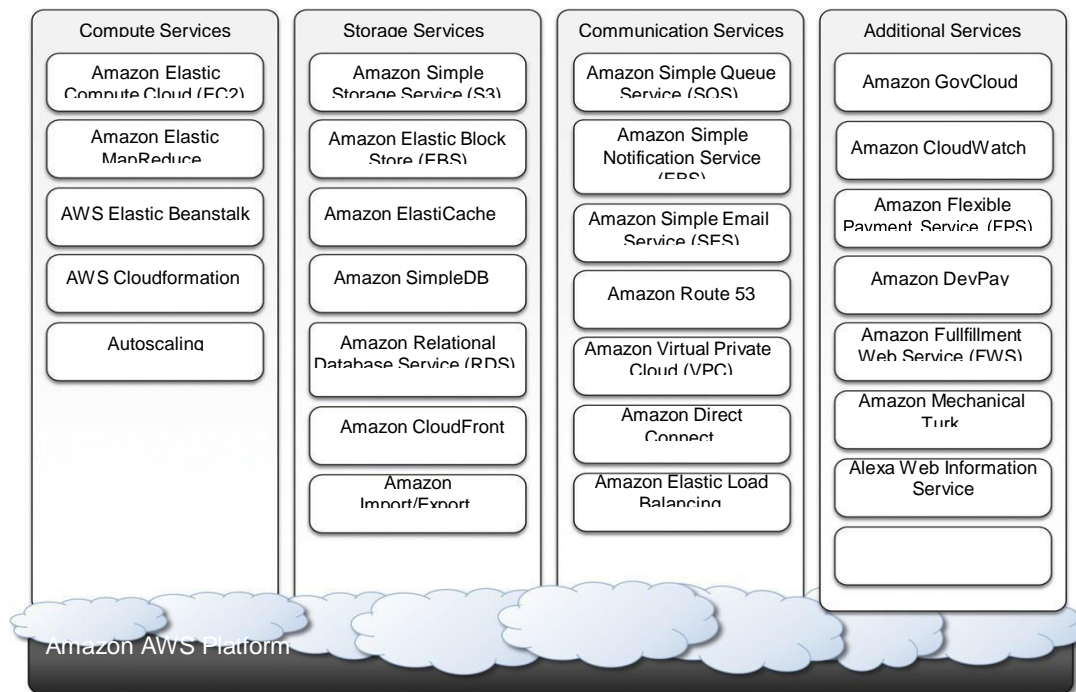


Fig 1 : Amazon Web service Ecosystem

Compute services

Compute services constitute the fundamental element of cloud computing systems. The fundamental service in this space is Amazon EC2, which delivers an IaaS solution that has served as a reference model for several offerings from other vendors in the same market segment. Amazon EC2 allows deploying servers in the form of virtual machines created as instances of a specific image. Images come with a preinstalled operating system and a software stack, and instances can be configured for memory, number of processors, and storage. Users are provided with credentials to remotely access the instance and further configure or install software if needed.

Amazon machine images

Amazon Machine Images (AMIs) are templates from which it is possible to create a virtual machine. They are stored in Amazon S3 and identified by a unique identifier in the form of *ami-xxxxxx* and a manifest XML file. An AMI contains a physical file system layout with a

predefined operating system installed. These are specified by the *Amazon Ramdisk Image* (ARI, id: *ari-yyyyyy*) and the *Amazon Kernel Image* (AKI, id: *aki-zzzzzz*), which are part of the configuration of the template. AMIs are either created from scratch or “bundled” from existing EC2 instances. A common practice is to prepare new AMIs to create an instance from a preexisting AMI, log into it once it is booted and running, and install all the software needed. Using the tools provided by Amazon, we can convert the instance into a new image. Once an AMI is created, it is stored in an S3 bucket and the user can decide whether to make it available to other users or keep it for personal use. Finally, it is also possible to associate a product code with a given AMI, thus allowing the owner of the AMI to get revenue every time this AMI is used to create EC2 instances.

EC2 instances

EC2 instances represent virtual machines. They are created using AMI as templates, which are specialized by selecting the number of cores, their computing power, and the installed memory. The processing power is expressed in terms of virtual cores and EC2 Compute Units (ECUs). The ECU is a measure of the computing power of a virtual core; it is used to express a predictable quantity of real CPU power that is allocated to an instance. By using compute units instead of real frequency values, Amazon can change over time the mapping of such units to the underlying real amount of computing power allocated, thus keeping the performance of EC2 instances consistent with standards set by the times. Over time, the hardware supporting the underlying infrastructure will be replaced by more powerful hardware, and the use of ECUs helps give users a consistent view of the performance offered by EC2 instances. Since users rent computing capacity rather than buying hardware, this approach is reasonable. One ECU is defined as giving the same performance as a 1.0 —1.2 GHz 2007 Opteron or 2007 Xeon processor.¹

Table 5.2 shows all the currently available configurations for EC2 instances. We can identify **six major categories**:

- ***Standard instances***. This class offers a set of configurations that are suitable for most applications. EC2 provides three different categories of increasing computing power, storage, and memory.

- ***Micro instances.*** This class is suitable for those applications that consume a limited amount of computing power and memory and occasionally need bursts in CPU cycles to process surges in the workload. Micro instances can be used for small Web applications with limited traffic.
- ***High-memory instances.*** This class targets applications that need to process huge workloads and require large amounts of memory. Three-tier Web applications characterized by high traffic are the target profile. Three categories of increasing memory and CPU are available, with memory proportionally larger than computing power.
- ***High-CPU instances.*** This class targets compute-intensive applications. Two configurations are available where computing power proportionally increases more than memory.
- ***Cluster Compute instances.*** This class is used to provide virtual cluster services. Instances in this category are characterized by high CPU compute power and large memory and an extremely high I/O and network performance, which makes it suitable for HPC applications.
- ***Cluster GPU instances.*** This class provides instances featuring graphic processing units (GPUs) and high compute power, large memory, and extremely high I/O and network performance. This class is particularly suited for cluster applications that perform heavy graphic computations, such as rendering clusters. Since GPU can be used for general-purpose computing, users of such instances can benefit from additional computing power, which makes this class suitable for HPC applications.

EC2 instances are priced hourly according to the category they belong to. At the beginning of every hour of usage, the user will be charged the cost of the entire hour. The hourly expense charged for one instance is constant. Instance owners are responsible for providing their own backup strategies, since there is no guarantee that the instance will run for the entire hour. Another alternative is represented by *spot instances*. These instances are much more dynamic in terms of pricing and lifetime since they are made available to the user according to the load of EC2 and the availability of resources. Users define an upper bound for a price they want to pay for these instances; as long as the current price (the spot price) remains under the given bound, the instance is kept running. The price is sampled at the beginning of each hour. Spot instances are more volatile than normal instances; whereas for normal instances EC2 will try as much as possible to keep them active, there is no such guarantee for spot instances. Therefore, implementing backup and checkpointing strategies is inevitable.

EC2 instances can be run either by using the command-line tools provided by Amazon, which connects the Amazon Web Service that provides remote access to the EC2 infrastructure, or via the AWS console, which allows the management of other services, such as S3. By default an EC2 instance is created with the kernel and the disk associated to the AMI. These define the architecture (32 bit or 64 bit) and the space of disk available to the instance. This is an ephemeral disk; once the instance is shut down, the content of the disk will be lost. Alternatively, it is possible to attach an EBS volume to the instance, the content of which will be stored in S3. If the default AKI and ARI are not suitable, EC2 provides capabilities to run EC2 instances by specifying a different AKI and ARI, thus giving flexibility in the creation of instances.

EC2 ENVIRONMENT

EC2 instances are executed within a virtual environment, which provides them with the services they require to host applications. The EC2 environment is in charge of allocating addresses, attaching storage volumes, and configuring security in terms of access control and network connectivity.

By default, instances are created with an internal IP address, which makes them capable of communicating within the EC2 network and accessing the Internet as clients. It is possible to associate an *Elastic IP* to each instance, which can then be remapped to a different instance over time. Elastic IPs allow instances running in EC2 to act as servers reachable from the Internet and, since they are not strictly bound to specific instances, to implement failover capabilities. Together with an external IP, EC2 instances are also given a domain name that generally is in the form *ec2-xxx-xxx-xxx.compute-x.amazonaws.com*, where *xxx-xxx-xxx* normally represents the four parts of the external IP address separated by a dash, and *compute-x* gives information about the availability zone where instances are deployed. Currently, there are five availability zones that are priced differently: two in the United States (Virginia and Northern California), one in Europe (Ireland), and two in Asia Pacific (Singapore and Tokyo). Instance owners can partially control where to deploy instances. Instead, they have a finer control over the security of the instances as well as their network accessibility. Instance owners can associate a key pair to one or more instances when these instances are created. A key pair allows the owner to remotely connect to the instance once this is running and gain

root access to it. Amazon EC2 controls the accessibility of a virtual instance with basic firewall configuration, allowing the specification of source address, port, and protocols (TCP, UDP, ICMP). Rules can also be attached to security groups, and instances can be made part of one or more groups before their deployment. Security groups and firewall rules constitute a flexible way of providing basic security for EC2 instances, which has to be complemented by appropriate security configuration within the instance itself.

ADVANCED compute SERVICES

EC2 instances and AMIs constitute the basic blocks for building an IaaS computing cloud. On top of these, Amazon Web Services provide more sophisticated services that allow the easy packaging and deploying of applications and a computing platform that supports the execution of MapReduce- based applications.

AWS CloudFormation

AWS CloudFormation constitutes an extension of the simple deployment model that characterizes EC2 instances. CloudFormation introduces the concepts of *templates*, which are JSON formatted text files that describe the resources needed to run an application or a service in EC2 together with the relations between them. CloudFormation allows easily and explicitly linking EC2 instances together and introducing dependencies among them. Templates provide a simple and declarative way to build complex systems and integrate EC2 instances with other AWS services such as S3, SimpleDB, SQS, SNS, Route 53, Elastic Beanstalk, and others.

AWS elastic beanstalk

AWS Elastic Beanstalk constitutes a simple and easy way to package applications and deploy them on the AWS Cloud. This service simplifies the process of provisioning instances and deploying application code and provides appropriate access to them. Currently, this service is available only for Web applications developed with the Java/Tomcat technology stack. Developers can conveniently package their Web application into a WAR file and use Beanstalk to automate its deployment on the AWS Cloud.

With respect to other solutions that automate cloud deployment, Beanstalk simplifies tedious tasks without removing the user's capability of accessing—and taking over control of—the

underlying EC2 instances that make up the virtual infrastructure on top of which the application is running. With respect to AWS CloudFormation, AWS Elastic Beanstalk provides a higher-level approach for application deployment on the cloud, which does not require the user to specify the infrastructure in terms of EC2 instances and their dependencies.

Amazon elastic MapReduce

Amazon Elastic MapReduce provides AWS users with a cloud computing platform for MapReduce applications. It utilizes Hadoop as the MapReduce engine, deployed on a virtual infrastructure composed of EC2 instances, and uses Amazon S3 for storage needs.

Apart from supporting all the application stack connected to Hadoop (Pig, Hive, etc.), Elastic MapReduce introduces elasticity and allows users to dynamically size the Hadoop cluster according to their needs, as well as select the appropriate configuration of EC2 instances to compose the cluster (Small, High-Memory, High-CPU, Cluster Compute, and Cluster GPU). On top of these services, basic Web applications allowing users to quickly run data-intensive applications without writing code are offered.

Storage services

AWS provides a collection of services for data storage and information management. The core service in this area is represented by Amazon *Simple Storage Service (S3)*. This is a distributed object store that allows users to store information in different formats. The core components of S3 are two: *buckets* and *objects*. Buckets represent virtual containers in which to store objects; objects represent the content that is actually stored. Objects can also be enriched with metadata that can be used to tag the stored content with additional information.

S3 key concepts

As the name suggests, S3 has been designed to provide a simple storage service that's accessible through a Representational State Transfer (REST) interface, which is quite similar to a distributed file system but which presents some important differences that allow the infrastructure to be highly efficient:

- ***The storage is organized in a two-level hierarchy.*** S3 organizes its storage space into buckets that cannot be further partitioned. This means that it is not possible to create directories or other kinds of physical groupings for objects stored in a bucket. Despite this fact, there are

few limitations in naming objects, and this allows users to simulate directories and create logical groupings.

- ***Stored objects cannot be manipulated like standard files.*** S3 has been designed to essentially provide storage for objects that will not change over time. Therefore, it does not allow renaming, modifying, or relocating an object. Once an object has been added to a bucket, its content and position is immutable, and the only way to change it is to remove the object from the store and add it again.
- ***Content is not immediately available to users.*** The main design goal of S3 is to provide an eventually consistent data store. As a result, because it is a large distributed storage facility, changes are not immediately reflected. For instance, S3 uses replication to provide redundancy and efficiently serve objects across the globe; this practice introduces latencies when adding objects to the store—especially large ones—which are not available instantly across the entire globe.
- ***Requests will occasionally fail.*** Due to the large distributed infrastructure being managed, requests for object may occasionally fail. Under certain conditions, S3 can decide to drop a request by returning an internal server error. Therefore, it is expected to have a small failure rate during day-to-day operations, which is generally not identified as a persistent failure.

Resource naming

Buckets, objects, and attached metadata are made accessible through a REST interface. Therefore, they are represented by *uniform resource identifiers (URIs)* under the s3.amazonaws.com domain. All the operations are then performed by expressing the entity they are directed to in the form of a request for a URI.

Amazon offers three different ways of addressing a bucket:

- ***Canonical form:*** http://s3.amazonaws.com/bucket_name/. The bucket name is expressed as a path component of the domain name s3.amazonaws.com. This is the naming convention that has less restriction in terms of allowed characters, since all the characters that are allowed for a path component can be used.
- ***Subdomain form:*** <http://bucketname.s3.amazonaws.com/>. Alternatively, it is also possible to reference a bucket as a subdomain of s3.amazonaws.com. To express a bucket name in this form, the name has to do all of the following:

- Be between 3 and 63 characters long
- Contain only letters, numbers, periods, and dashes
- Start with a letter or a number
- Contain at least one letter
- Have no fragments between periods that start with a dash or end with a dash or that are empty strings

This form is equivalent to the previous one when it can be used, but it is the one to be preferred since it works more effectively for all the geographical locations serving resources stored in S3.

- **Virtual hosting form:** <http://bucket-name.com/>. Amazon also allows referencing of its resources with custom URLs. This is accomplished by entering a CNAME record into the DNS that points to the subdomain form of the bucket URI.

Buckets

A *bucket* is a container of objects. It can be thought of as a virtual drive hosted on the S3 distributed storage, which provides users with a flat store to which they can add objects. Buckets are top-level elements of the S3 storage architecture and do not support nesting. That is, it is not possible to create “subbuckets” or other kinds of physical divisions.

A bucket is located in a specific geographic location and eventually replicated for fault tolerance and better content distribution. Users can select the location at which to create buckets, which by default are created in Amazon’s U.S. datacenters. Once a bucket is created, all the objects that belong to the bucket will be stored in the same availability zone of the bucket. Users create a bucket by sending a *PUT* request to <http://s3.amazonaws.com/> with the name of the bucket and, if they want to specify the availability zone, additional information about the preferred location. The content of a bucket can be listed by sending a *GET* request specifying the name of the bucket. Once created, the bucket cannot be renamed or relocated. If it is necessary to do so, the bucket needs to be deleted and recreated. The deletion of a bucket is performed by a *DELETE* request, which can be successful if and only if the bucket is empty.

Amazon elastic block store

The Amazon Elastic Block Store (EBS) allows AWS users to provide EC2 instances with persistent storage in the form of volumes that can be mounted at instance startup. They accommodate up to 1 TB of space and are accessed through a block device interface, thus

allowing users to format them according to the needs of the instance they are connected to (raw storage, file system, or other). The content of an EBS volume survives the instance life cycle and is persisted into S3. EBS volumes can be cloned, used as boot partitions, and constitute durable storage since they rely on S3 and it is possible to take incremental snapshots of their content.

EBS volumes normally reside within the same availability zone of the EC2 instances that will use them to maximize the I/O performance. It is also possible to connect volumes located in different availability zones. Once mounted as volumes, their content is lazily loaded in the background and according to the request made by the operating system. This reduces the number of I/O requests that go to the network. Volume images cannot be shared among instances, but multiple (separate) active volumes can be created from them. In addition, it is possible to attach multiple volumes to a single instance or create a volume from a given snapshot and modify its size, if the formatted file system allows such an operation.

The expense related to a volume comprises the cost generated by the amount of storage occupied in S3 and by the number of I/O requests performed against the volume. Currently, Amazon charges \$0.10/GB/month of allocated storage and \$0.10 per 1 million requests made to the volume.

Amazon ElastiCache

ElastiCache is an implementation of an elastic in-memory cache based on a cluster of EC2 instances. It provides fast data access from other EC2 instances through a Memcached-compatible protocol so that existing applications based on such technology do not need to be modified and can transparently migrate to ElastiCache. ElastiCache is based on a cluster of EC2 instances running the caching software, which is made available through Web services. An ElastiCache cluster can be dynamically resized according to the demand of the client applications. Furthermore, automatic patch management and failure detection and recovery of cache nodes allow the cache cluster to keep running without administrative intervention from AWS users, who have only to elastically size the cluster when needed.

Preconfigured EC2 AMIs

Preconfigured EC2 AMIs are predefined templates featuring an installation of a given database management system. EC2 instances created from these AMIs can be completed with an EBS

volume for storage persistence. Available AMIs include installations of IBM DB2, Microsoft SQL Server, MySQL, Oracle, PostgreSQL, Sybase, and Vertica. Instances are priced hourly according to the EC2 cost model. This solution poses most of the administrative burden on the EC2 user, who has to configure, maintain, and manage the relational database, but offers the greatest variety of products to choose from.

Amazon RDS

RDS is relational database service that relies on the EC2 infrastructure and is managed by Amazon. Developers do not have to worry about configuring the storage for high availability, designing failover strategies, or keeping the servers up-to-date with patches. Moreover, the service provides users with automatic backups, snapshots, point-in-time recoveries, and facilities for implementing replications. These and the common database management services are available through the AWS console or a specific Web service. Two relational engines are available: MySQL and Oracle.

Two key advanced features of RDS are *multi-AZ deployment* and *read replicas*. The first option provides users with a failover infrastructure for their RDBMS solutions. The high-availability solution is implemented by keeping in standby synchronized copies of the services in different availability zones that are activated if the primary service goes down. The second option provides users with increased performance for applications that are heavily based on database reads. In this case, Amazon deploys copies of the primary service that are only available for database reads, thus cutting down the response time of the service.

The available options and the relative pricing of the service during 2011—2012 are shown in [Table 5.4](#). The table shows the costing details of the on-demand instances. There is also the possibility of using reserved instances for long terms (one to three years) by paying up-front at discounted hourly rates.

With respect to the previous solution, users are not responsible for managing, configuring, and patching the database management software, but these operations are performed by the AWS. In addition, support for elastic management of servers is simplified. Therefore, this solution is optimal for applications based on the Oracle and MySQL engines, which are migrated on the AWS infrastructure and require a scalable database solution.

Amazon SimpleDB

Amazon SimpleDB is a lightweight, highly scalable, and flexible data storage solution for applications that do not require a fully relational model for their data. SimpleDB provides support for semistructured data, the model for which is based on the concept of *domains*, *items*, and *attributes*. With respect to the relational model, this model provides fewer constraints on the structure of data entries, thus obtaining improved performance in querying large quantities of data. As happens for Amazon RDS, this service frees AWS users from performing configuration, management, and high-availability design for their data stores.

SimpleDB uses *domains* as top-level elements to organize a data store. These domains are roughly comparable to tables in the relational model. Unlike tables, they allow items not to have all the same column structure; each item is therefore represented as a collection of attributes expressed in the form of a key-value pair. Each domain can grow up to 10 GB of data, and by default a single user can allocate a maximum of 250 domains. Clients can create, delete, modify, and make snapshots of domains. They can insert, modify, delete, and query items and attributes. Batch insertion and deletion are also supported. The capability of querying data is one of the most relevant functions of the model, and the *select* clause supports the following test operators: *=*, *!=*, *<*, *>*, *<=*, *>=*, *like*, *not like*, *between*, *is null*, *is not null*, and *every()*. Here is a simple example on how to query data:

```
select * from domain_name where every(attribute_name) = 'value'
```

Moreover, the *select* operator can extend its query beyond the boundaries of a single domain, thus allowing users to query effectively a large amount of data.

To efficiently provide AWS users with a scalable and fault-tolerant service, SimpleDB implements a relaxed constraint model, which leads to *eventually consistent* data. The adverb *eventually* denotes the fact that multiple accesses on the same data might not read the same value in the very short term, but they will eventually converge over time. This is because SimpleDB does not lock all the copies of the data during an update, which is propagated in the background. Therefore, there is a transient period of time in which different clients can access different copies of the same data that have different values. This approach is very scalable with minor drawbacks, and it is also reasonable, since the application scenario for SimpleDB is

mostly characterized by querying and indexing operations on data. Alternatively, it is possible to change the default behavior and ensure that all the readers are blocked during an update. Even though SimpleDB is not a transactional model, it allows clients to express conditional insertions or deletions, which are useful to prevent lost updates in multiple-writer scenarios. In this case, the operation is executed if and only if the condition is verified. This condition can be used to check preexisting values of attributes for an item.

The service charges either for data transfer or stored data. Data transfer within the AWS network is not charged. In addition, SimpleDB also charges users for machine usage. The first 25 SimpleDB instances per month are free; after this threshold there is an hourly charge (\$0.140 hour in the U.S. East region). If we compare this cost model with the one characterizing S3, it becomes evident that S3 is a cheaper option for storing large objects. This is useful information for clarifying the different nature of SimpleDB with respect to S3: The former has been designed to provide fast access to semistructured collections of small objects and not for being a long-term storage option for large objects.

AWS provides users with simple Web service APIs to manage CloudFront. To make available content through CloudFront, it is necessary to create a distribution. This identifies an origin server, which contains the original version of the content being distributed, and it is referenced by a DNS domain under the *Cloudfront.net* domain name (i.e., my-distribution.Cloudfront.net). It is also possible to map a given domain name to a distribution. Once the distribution is created, it is sufficient to reference the distribution name, and the CloudFront engine will redirect the request to the closest replica and eventually download the original version from the origin server if the content is not found or expired on the selected edge server.

The content that can be delivered through CloudFront is static (HTTP and HTTPS) or streaming (Real Time Messaging Protocol, or RTMP). The origin server hosting the original copy of the distributed content can be an S3 bucket, an EC2 instance, or a server external to the Amazon network. Users can restrict access to the distribution to only one or a few of the available protocols, or they can set up access rules for finer control. It is also possible to invalidate content to remove it from the distribution or force its update before expiration.

Note that CloudFront is cheaper than S3. This reflects its different purpose: CloudFront is designed to optimize the distribution of very popular content that is frequently downloaded, potentially from the entire globe and not only the Amazon network.

Communication services

Amazon provides facilities to structure and facilitate the communication among existing applications and services residing within the AWS infrastructure. These facilities can be organized into two major categories: *virtual networking* and *messaging*.

Virtual networking

Virtual networking comprises a collection of services that allow AWS users to control the connectivity to and between compute and storage services. *Amazon Virtual Private Cloud (VPC)* and *Amazon Direct Connect* provide connectivity solutions in terms of infrastructure; *Route 53* facilitates connectivity in terms of naming.

Amazon VPC provides a great degree of flexibility in creating virtual private networks within the Amazon infrastructure and beyond. The service providers prepare either templates covering most of the usual scenarios or a fully customizable network service for advanced configurations. Prepared templates include public subnets, isolated networks, private networks accessing Internet through network address translation (NAT), and hybrid networks including AWS resources and private resources. Also, it is possible to control connectivity between different services (EC2 instances and S3 buckets) by using the *Identity Access Management (IAM)* service. During 2011, the cost of Amazon VPC was \$0.50 per connection hour.

Amazon Direct Connect allows AWS users to create dedicated networks between the user private network and Amazon Direct Connect locations, called *ports*. This connection can be further partitioned in multiple logical connections and give access to the public resources hosted on the Amazon infrastructure. The advantage of using Direct Connect versus other solutions is the consistent performance of the connection between the users' premises and the Direct Connect locations. This service is compatible with other services such as EC2, S3, and Amazon VPC and can be used in scenarios requiring high bandwidth between the Amazon network and the outside world. There are only two available ports located in the United States, but users can leverage external providers that offer guaranteed high

bandwidth to these ports. Two different bandwidths can be chosen: 1 Gbps, priced at \$0.30 per hour, and 10 Gbps, priced at \$2.25 per hour. Inbound traffic is free; outbound traffic is priced at \$0.02 per GB.

Amazon Route 53 implements dynamic DNS services that allow AWS resources to be reached through domain names different from the amazon.com domain. By leveraging the large and globally distributed network of Amazon DNS servers, AWS users can expose EC2 instances or S3 buckets as resources under a domain of their property, for which Amazon DNS servers become authoritative.³ EC2 instances are likely to be more dynamic than the physical machines, and S3 buckets might also exist for a limited time. To cope with such a volatile nature, the service provides AWS users with the capability of dynamically mapping names to resources as instances are launched on EC2 or as new buckets are created in S3. By interacting with the Route 53 Web service, users can manage a set of *hosted zones*, which represent the user domains controlled by the service, and edit the resources made available through it. Currently, a single user can have up to 100 zones. The costing model includes a fixed amount (\$1 per zone per month) and a dynamic component that depends on the number of queries resolved by the service for the hosted zones (\$0.50 per million queries for the first billion of queries a month, \$0.25 per million queries over 1 billion of queries a month).

Messaging

Messaging services constitute the next step in connecting applications by leveraging AWS capabilities.

The three different types of messaging services offered are

- ***Amazon Simple Queue Service (SQS),***
- ***Amazon Simple Notification Service (SNS), and***
- ***Amazon Simple Email Service (SES).***

Amazon SQS constitutes disconnected model for exchanging messages between applications by means of message queues, hosted within the AWS infrastructure. Using the AWS console or directly the underlying Web service AWS, users can create an unlimited number of message queues and configure them to control their access. Applications can send messages to any queue they have access to. These messages are securely and redundantly stored within the

AWS infrastructure for a limited period of time, and they can be accessed by other (authorized) applications. While a message is being read, it is kept locked to avoid spurious processing from other applications. Such a lock will expire after a given period.

Amazon SNS provides a publish-subscribe method for connecting heterogeneous applications. With respect to Amazon SQS, where it is necessary to continuously poll a given queue for a new message to process, Amazon SNS allows applications to be notified when new content of interest is available. This feature is accessible through a Web service whereby AWS users can create a topic, which other applications can subscribe to. At any time, applications can publish content on a given topic and subscribers can be automatically notified. The service provides subscribers with different notification models (HTTP/HTTPS, email/email JSON, and SQS).

Amazon SES provides AWS users with a scalable email service that leverages the AWS infrastructure. Once users are signed up for the service, they have to provide an email that SES will use to send emails on their behalf. To activate the service, SES will send an email to verify the given address and provide the users with the necessary information for the activation. Upon verification, the user is given an SES sandbox to test the service, and he can request access to the production version. Using SES, it is possible to send either SMTP-compliant emails or raw emails by specifying email headers and Multipurpose Internet Mail Extension (MIME) types. Emails are queued for delivery, and the users are notified of any failed delivery. SES also provides a wide range of statistics that help users to improve their email campaigns for effective communication with customers.

With regard to the costing, all three services do not require a minimum commitment but are based on a pay-as-you go model. Currently, users are not charged until they reach a minimum threshold. In addition, data transfer-in is not charged, but data transfer-out is charged by ranges.

Google AppEngine

Google AppEngine is a PaaS implementation that provides services for developing and hosting scalable Web applications. AppEngine is essentially a distributed and scalable runtime environment that leverages Google's distributed infrastructure to scale out applications facing a large number of requests by allocating more computing resources to them and balancing the load among them. The runtime is completed by a collection of services that allow developers

to design and implement applications that naturally scale on AppEngine. Developers can develop applications in Java, Python, and Go, a new programming language developed by Google to simplify the development of Web applications. Application usage of Google resources and services is metered by AppEngine, which bills users when their applications finish their free quotas.

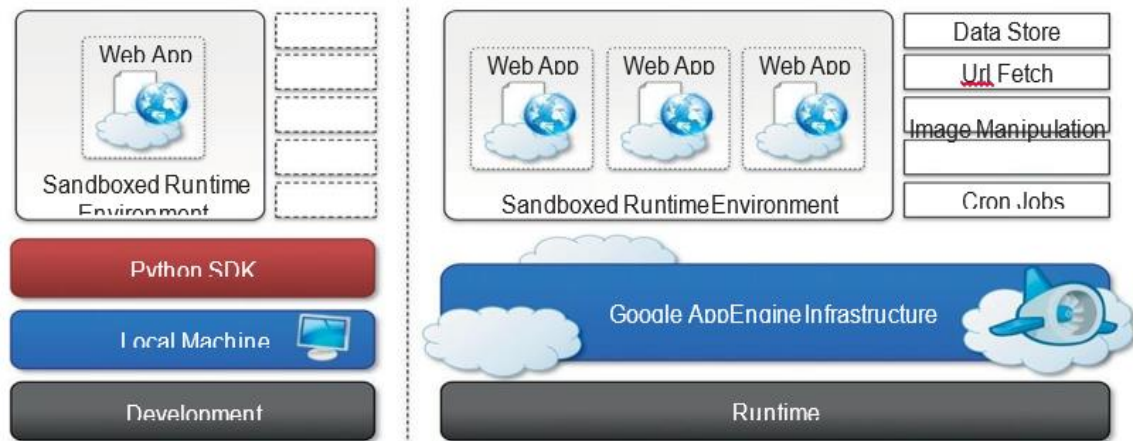


Fig 2: Google App Engine Platform Architecture

5.2.1 Architecture and core concepts

AppEngine is a platform for developing scalable applications accessible through the Web. The platform is logically divided into four major components: infrastructure, the run-time environment, the underlying storage, and the set of scalable services that can be used to develop applications.

5.2.1.1 *Infrastructure*

AppEngine hosts Web applications, and its primary function is to serve users requests efficiently. To do so, AppEngine's infrastructure takes advantage of many servers available within Google datacenters. For each HTTP request, AppEngine locates the servers hosting the application that processes the request, evaluates their load, and, if necessary, allocates additional resources (i.e., servers) or redirects the request to an existing server. The particular design of applications, which does not expect any state information to be implicitly maintained between requests to the same application, simplifies the work of the infrastructure, which can redirect each of the requests to any of the servers hosting the target application or

even allocate a new one. The infrastructure is also responsible for monitoring application performance and collecting statistics on which the billing is calculated

Application SERVICES

Applications hosted on AppEngine take the most from the services made available through the run-time environment. These services simplify most of the common operations that are performed in Web applications: access to data, account management, integration of external resources, messaging and communication, image manipulation, and asynchronous computation.

UrlFetch

Web 2.0 has introduced the concept of composite Web applications. Different resources are put together and organized as meshes within a single Web page. Meshes are fragments of HTML generated in different ways. They can be directly obtained from a remote server or rendered from an XML document retrieved from a Web service, or they can be rendered by the browser as the result of an embedded and remote component. A common characteristic of all these examples is the fact that the resource is not local to the server and often not even in the same administrative domain. Therefore, it is fundamental for Web applications to be able to retrieve remote resources. The sandbox environment does not allow applications to open arbitrary connections through sockets, but it does provide developers with the capability of retrieving a remote resource through HTTP/HTTPS by means of the *UrlFetch* service. Applications can make synchronous and asynchronous Web requests and integrate the resources obtained in this way into the normal request-handling cycle of the application. One of the interesting features of *UrlFetch* is the ability to set deadlines for requests so that they can be completed (or aborted) within a given time. Moreover, the ability to perform such requests asynchronously allows the applications to continue with their logic while the resource is retrieved in the background. *UrlFetch* is not only used to integrate meshes into a Web page but also to leverage remote Web services in accordance with the SOA reference model for distributed applications

MemCache

AppEngine provides developers with access to fast and reliable storage, which is DataStore. Despite this, the main objective of the service is to serve as a scalable and long-term storage, where data are persisted to disk redundantly in order to ensure reliability and availability of data against failures. This design poses a limit on how much faster the store can be compared to other solutions, especially for objects that are frequently accessed—for example, at each Web request.

AppEngine provides caching services by means of *MemCache*. This is a distributed in-memory cache that is optimized for fast access and provides developers with a volatile store for the objects that are frequently accessed. The caching algorithm implemented by MemCache will automatically remove the objects that are rarely accessed. The use of MemCache can significantly reduce the access time to data; developers can structure their applications so that each object is first looked up into MemCache and if there is a miss, it will be retrieved from DataStore and put into the cache for future lookups.

Mail and instant messaging

Communication is another important aspect of Web applications. It is common to use email for following up with users about operations performed by the application. Email can also be used to trigger activities in Web applications. To facilitate the implementation of such tasks, AppEngine provides developers with the ability to send and receive mails through *Mail*. The service allows sending email on behalf of the application to specific user accounts. It is also possible to include several types of attachments and to target multiple recipients. Mail operates asynchronously, and in case of failed delivery the sending address is notified through an email detailing the error.

AppEngine provides also another way to communicate with the external world: the Extensible Messaging and Presence Protocol (XMPP). Any chat service that supports XMPP, such as Google Talk, can send and receive chat messages to and from the Web application, which is identified by its own address. Even though the chat is a communication medium mostly used for human interactions, XMPP can be conveniently used to connect the Web application with chat bots or to implement a small administrative console.

Account management

Web applications often keep various data that customize their interaction with users. These data normally go under the user profile and are attached to an account. AppEngine simplifies account management by allowing developers to leverage Google account management by means of *Google Accounts*. The integration with the service also allows Web applications to offload the implementation of authentication capabilities to Google's authentication system. Using Google Accounts, Web applications can conveniently store profile settings in the form of key-value pairs, attach them to a given Google account, and quickly retrieve them once the user authenticates. With respect to a custom solution, the use of Google Accounts requires users to have a Google account, but it does not require any further implementation. The use of Google Accounts is particularly advantageous for developing Web applications within a corporate environment using Google Apps. In this case, the applications can be easily integrated with all the other services (and profile settings) included in Google Apps.

Image manipulation

Web applications render pages with graphics. Often simple operations, such as adding watermarks or applying simple filters, are required. AppEngine allows applications to perform image resizing, rotation, mirroring, and enhancement by means of *Image Manipulation*, a service that is also used in other Google products. Image Manipulation is mostly designed for lightweight image processing and is optimized for speed

Compute SERVICES

Web applications are mostly designed to interface applications with users by means of a ubiquitous channel, that is, the Web. Most of the interaction is performed synchronously: Users navigate the Web pages and get instantaneous feedback in response to their actions. This feedback is often the result of some computation happening on the Web application, which implements the intended logic to serve the user request. Sometimes this approach is not applicable—for example, in long computations or when some operations need to be triggered at a given point in time. A good design for these scenarios provides the user with immediate feedback and a notification once the required operation is completed. AppEngine offers additional services such as *Task Queues* and *Cron Jobs* that simplify the execution of

computations that are off-bandwidth or those that cannot be performed within the timeframe of the Web request.

Task queues

Task Queues allow applications to submit a task for a later execution. This service is particularly useful for long computations that cannot be completed within the maximum response time of a request handler. The service allows users to have up to 10 queues that can execute tasks at a configurable rate.

In fact, a task is defined by a Web request to a given URL, and the queue invokes the request handler by passing the payload as part of the Web request to the handler. It is the responsibility of the request handler to perform the “task execution,” which is seen from the queue as a simple Web request. The queue is designed to reexecute the task in case of failure in order to avoid transient failures preventing the task from a successful completion.

Cron jobs

Sometimes the length of computation might not be the primary reason that an operation is not performed within the scope of the Web request. It might be possible that the required operation needs to be performed at a specific time of the day, which does not coincide with the time of the Web request. In this case, it is possible to schedule the required operation at the desired time by using the *Cron Jobs* service. This service operates similarly to Task Queues but invokes the request handler specified in the task at a given time and does not reexecute the task in case of failure. This behavior can be useful to implement maintenance operations or send periodic notifications.

Application life cycle

AppEngine provides support for almost all the phases characterizing the life cycle of an application: testing and development, deployment, and monitoring. The SDKs released by Google provide developers with most of the functionalities required by these tasks. Currently there are two SDKs available for development: Java SDK and Python SDK.

Application DEVELOPMENT and testing

Developers can start building their Web applications on a local development server. This is a self-contained environment that helps developers tune applications without uploading them to AppEngine. The development server simulates the AppEngine runtime environment by providing a mock implementation of DataStore, MemCache, UrlFetch, and the other services leveraged by Web applications. Besides hosting Web applications, the development server contains a complete set of monitoring features that are helpful to profile the behavior of applications, especially regarding access to the DataStore service and the queries performed against it. This is a particularly important feature that will be of relevance in deploying the application to AppEngine. As discussed earlier, AppEngine builds indexes for each of the queries performed by a given application in order to speed up access to the relevant data. This capability is enabled by *a priori* knowledge about all the possible queries made by the application; such knowledge is made available to AppEngine by the developer while uploading the application. The development server analyzes application behavior while running and traces all the queries made during testing and development, thus providing the required information about the indexes to be built.

Java SDK

The Java SDK provides developers with the facility for building applications with the Java 5 and Java 6 runtime environments. Alternatively, it is possible to develop applications within the Eclipse development environment by using the Google AppEngine plug-in, which integrates the features of the SDK within the powerful Eclipse environment. Using the Eclipse software installer, it is possible to download and install Java SDK, Google Web Toolkit, and Google AppEngine plug-ins into Eclipse. These three components allow developers to program powerful and rich Java applications for AppEngine.

The SDK supports the development of applications by using the *servlet* abstraction, which is a common development model. Together with servlets, many other features are available to build applications. Moreover, developers can easily create Web applications by using the *Eclipse Web Platform*, which provides a set of tools and components.

The plug-in allows developing, testing, and deploying applications on AppEngine. Other tasks, such as retrieving the log of applications, are available by means of command-line tools that are part of the SDK.

Python SDK

The Python SDK allows developing Web applications for AppEngine with Python 2.5. It provides a standalone tool, called *GoogleAppEngineLauncher*, for managing Web applications locally and deploying them to AppEngine. The tool provides a convenient user interface that lists all the available Web applications, controls their execution, and integrates them with the default code editor for editing application files. In addition, the launcher provides access to some important services for application monitoring and analysis, such as the logs, the SDK console, and the dashboard. The log console captures all the information that is logged by the application while it is running. The console SDK provides developers with a Web interface via which they can see the application profile in terms of utilized resource. This feature is particularly useful because it allows developers to preview the behavior of the applications once they are deployed on AppEngine, and it can be used to tune applications made available through the runtime.

The Python implementation of the SDK also comes with an integrated Web application framework called *webapp* that includes a set of models, components, and tools that simplify the development of Web applications and enforce a set of coherent practices. This is not the only Web framework that can be used to develop Web applications. There are dozens of available Python Web frameworks that can be used. However, due to the restrictions enforced by the sandboxed environment, all of them cannot be used seamlessly. The *webapp* framework has been reimplemented and made available in the Python SDK so that it can be used with AppEngine. Another Web framework that is known to work well is *Django*.⁴

The SDK is completed by a set of command-line tools that allows developers to perform all the operations available through the launcher and more from the command shell.

Application deployment and management

Once the application has been developed and tested, it can be deployed on AppEngine with a simple click or command-line tool. Before performing such task, it is necessary to create an application identifier, which will be used to locate the application from the Web browser by typing the address *http://, application-id..appspot.com*. Alternatively, it is also possible to map the application with a registered DNS domain name. This is particularly useful for commercial

development, where users want to make the application available through a more appropriate name.

An application identifier is mandatory because it allows unique identification of the application while it's interacting with AppEngine. Developers use an app identifier to upload and update applications. Besides being unique, it also needs to be compliant to the rules that are enforced for domain names. It is possible to register an application identifier by logging into AppEngine and selecting the "Create application" option. It is also possible to provide an application title that is descriptive of the application; the title can be changed over time.

Once an application identifier has been created, it is possible to deploy an application on AppEngine. This task can be done using either the respective development environment (*GoogleAppEngineLauncher* and *Google AppEngine* plug-in) or the command-line tools. Once the application is uploaded, nothing else needs to be done to make it available. AppEngine will take care of everything. Developers can then manage the application by using the administrative console. This is the primary tool used for application monitoring and provides users with insight into resource usage (CPU, bandwidth) and services and other useful counters. It is also possible to manage multiple versions of a single application, select the one available for the release, and manage its billing-related issues.

Microsoft Azure

Microsoft Windows Azure is a cloud operating system built on top of Microsoft datacenters' infrastructure and provides developers with a collection of services for building applications with cloud technology. Services range from compute, storage, and networking to application connectivity,

access control, and business intelligence. Any application that is built on the Microsoft technology can be scaled using the Azure platform, which integrates the scalability features into the common Microsoft technologies such as Microsoft Windows Server 2008, SQL Server, and ASP.NET.

Figure provides an overview of services provided by Azure. These services can be managed and controlled through the *Windows Azure Management Portal*, which acts as an administrative console for all the services offered by the Azure platform. In this section, we present the core features of the major services available with Azure.

Azure core concepts

The Windows Azure platform is made up of a foundation layer and a set of developer services that can be used to build scalable applications. These services cover compute, storage, networking, and identity management, which are tied together by middleware called *AppFabric*. This scalable computing environment is hosted within Microsoft datacenters and accessible through the Windows Azure Management Portal. Alternatively, developers can recreate a Windows Azure environment (with limited capabilities) on their own machines for development and testing purposes. In this section, we provide an overview of the Azure middleware and its services.

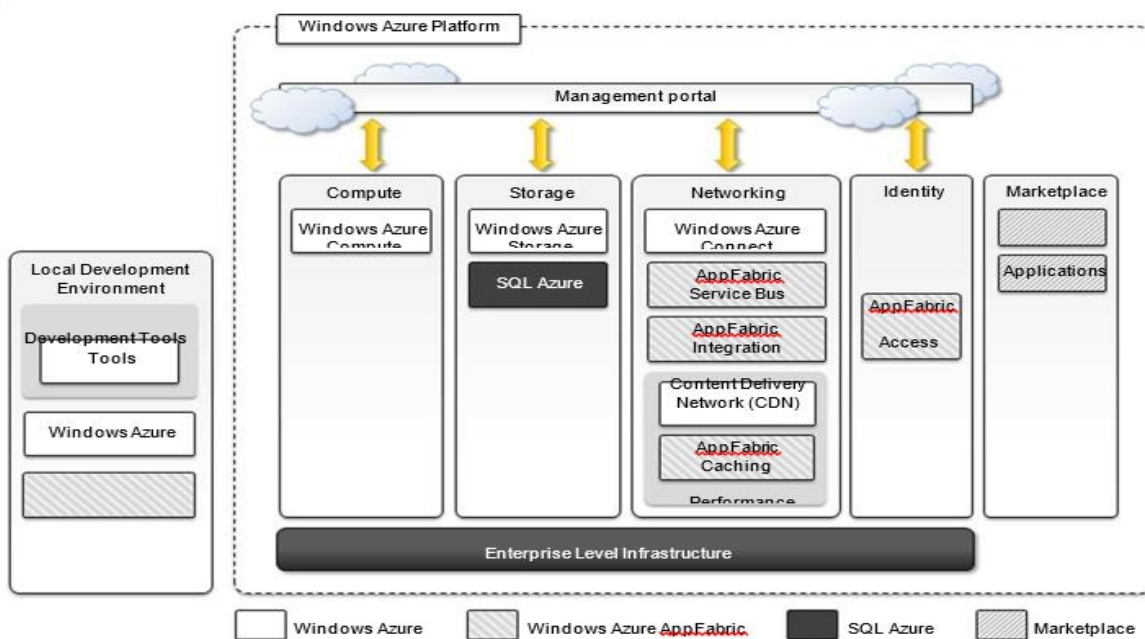


Fig 3: Microsoft Azure Platform Architecture

Compute SERVICES

Compute services are the core components of Microsoft Windows Azure, and they are delivered by means of the abstraction of *roles*. A role is a runtime environment that is customized for a specific compute task. Roles are managed by the Azure operating system and instantiated on demand in order to address surges in application demand. Currently, there are three different roles: *Web role*, *Worker role*, and *Virtual Machine (VM) role*.

Web role

The *Web role* is designed to implement scalable Web applications. Web roles represent the units of deployment of Web applications within the Azure infrastructure. They are hosted on the IIS 7 Web Server, which is a component of the infrastructure that supports Azure. When Azure detects peak loads in the request made to a given application, it instantiates multiple Web roles for that application and distributes the load among them by means of a load balancer.

Since version 3.5, the .NET technology natively supports Web roles; developers can directly develop their applications in Visual Studio, test them locally, and upload to Azure. It is possible to develop ASP.NET (*ASP.NET Web Role* and *ASP.NET MVC 2 Web Role*) and WCF (*WCF Service Web Role*) applications. Since IIS 7 also supports the PHP runtime environment by means of the FastCGI module, Web roles can be used to run and scale PHP Web applications on Azure (*CGI Web Role*). Other Web technologies that are not integrated with IIS can still be hosted on Azure (i.e., Java Server Pages on Apache Tomcat), but there is no advantage to using a Web role over a Worker role.

Windows Azure platform appliance

The Windows Azure platform can also be deployed as an appliance on third-party data centers and constitutes the cloud infrastructure governing the physical servers of the datacenter. The Windows Azure Platform Appliance includes Windows Azure, SQL Azure, and Microsoft-specified configuration of network, storage, and server hardware. The appliance is a solution that targets governments and service providers who want to have their own cloud computing infrastructure.

As introduced earlier, Azure already provides a development environment that allows building applications for Azure in their own premises. The local development environment is not intended to be production middleware, but it is designed for developing and testing the functionalities of applications that will eventually be deployed on Azure. The Azure appliance is instead a full-featured implementation of Windows Azure. Its goal is to replicate Azure on a third-party infrastructure and make available its services beyond the boundaries of the Microsoft Cloud. The appliance addresses two major scenarios: institutions that have very large computing needs (such as government agencies) and institutions that cannot afford to transfer their data outside their premises.

Observations

Windows Azure is Microsoft's solution for developing cloud computing applications. Azure is an implementation of the PaaS layer and provides the developer with a collection of services and scalable middleware hosted on Microsoft datacenters that address compute, storage, networking, and identity management needs of applications. The services Azure offers can be used either individually or all together for building both applications that integrate cloud features and elastic computing systems completely hosted in the cloud.

The core components of the platform are composed of compute services, storage services, and middleware. Compute services are based on the abstraction of roles, which identify a sandboxed environment where developers can build their distributed and scalable components. These roles are useful for Web applications, back-end processing, and virtual computing. Storage services include solutions for static and dynamic content, which is organized in the form of tables with fewer constraints than those imposed by the relational model. These and other services are implemented and made available through App Fabric, which constitutes the distributed and scalable middleware of Azure.

SQL Azure is another important element of Windows Azure and provides support for relational data in the cloud. SQL Azure is an extension of the capabilities of SQL Server adapted for the cloud environment and designed for dynamic scaling.

The platform is mostly based on the .NET technology and Windows systems, even though other technologies and systems can be supported.

Part – A			
Q.No	Questions	Competence	BT Level
1	List few examples for IAAS services	Remember	BT1
2	Define Bucketing in Cloud computing	Understand	BT2
3	Show the Permissions for access control	Remember	BT1
4	Define Virtual Networking	Remember	BT1
5	List the Messaging services in Amazon	Remember	BT1
6	Define Amazon Cloud watch	Remember	BT1
7	Define URL Fetching	Remember	BT1
8	Define Cron Jobs in AWS	Remember	BT1
9	Compare Block blobs and page blobs	Understand	BT2

Part – B			
Q.No	Questions	Competence	BT Level
1	Describe about the Cloud computing services	Understand	BT2
2	Discuss about amazon EC2 with its working	Understand	BT2
3	Explain about the Amazon S3	Understand	BT2
4	Illustrate about Elastic Block Storages with diagrams	Understand	BT2
5	Explain the Architecture of google app engine	Understand	BT2
6	Describe about the Microsoft Azure architecture	Understand	BT2
7	Summarize the working of SQL Azure Architecture	Understand	BT2

TEXT/ REFERENCE BOOKS

1. Rajkumar Buyya, Christian Vecchiola and S. Thamarai Selvi, “Mastering Cloud Computing”
- Foundations and Applications Programming , MK publications, 2013.
2. Gautam Shroff, “Enterprise Cloud Computing: Technology, Architecture, Applications” by
Cambridge University Press, 2010.