CMR
INSTITUTE OF
TECHNOLOGY

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

**Internal Assessment Test 3 – Jan. 2022**

| Sub: | CLOUD COMPUTING | | | | | | | Sub Code: | 20MCA342 |
|---|---|---|---|---|---|---|---|---|---|
| Date: | 25-01-2022 | Duration: | 90 min's | Max Marks: | 50 | Sem: | III | Branch: | MCA |

**Note : Answer FIVE FULL Questions, choosing ONE full question from each Module**

| | Marks | OBE | |
|---|---|---|---|
| | | CO | RBT |
| **PART-I** | | | |
| 1. Explain about SaaS. | [10] | CO4 | L2 |
| **(OR)** | | | |
| 2. Explain about the private cloud with necessary diagram. | [10] | CO4 | L2 |
| **PART-II** | | | |
| 3. Explain about the community cloud with necessary diagram. | [10] | CO4 | L2 |
| **(OR)** | | | |
| 4. Explain about the economics of the cloud. | [10] | CO4 | L2 |

CMR
INSTITUTE OF
TECHNOLOGY

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

**Internal Assessment Test 3 – Jan. 2022**

| Sub: | CLOUD COMPUTING | | | | | | | Sub Code: | 20MCA342 |
|---|---|---|---|---|---|---|---|---|---|
| Date: | 25-01-2022 | Duration: | 90 min's | Max Marks: | 50 | Sem: | III | Branch: | MCA |

**Note : Answer FIVE FULL Questions, choosing ONE full question from each Module**

| | Marks | OBE | |
|---|---|---|---|
| | | CO | RBT |
| **PART-I** | | | |
| 1. Explain about SaaS. | [10] | CO4 | L2 |
| **(OR)** | | | |
| 2. Explain about the private cloud with necessary diagram. | [10] | CO4 | L2 |
| **PART-II** | | | |
| 3. Explain about the community cloud with necessary diagram. | [10] | CO4 | L2 |
| **(OR)** | | | |
| 4. Explain about the economics of the cloud. | [10] | CO4 | L2 |

**PART-III**

| | | | |
|---|---|---|---|
| 5. Explain the list of open challenges in cloud computing. | [10] | CO4 | L2 |
| **(OR)** | | | |
| 6. Explain about the AWS with necessary diagram. | [10] | CO5 | L2 |
| **PART-IV** | | | |
| 7. Describe in detail about openstack. | [10] | CO5 | L2 |
| **(OR)** | | | |
| 8. Explain the architecture model of GAE with necessary diagram. | [10] | CO5 | L2 |
| **PART-V** | | | |
| 9. Explain how cloud computing can be used for ECG analysis. | [10] | CO5 | L2 |
| **(OR)** | | | |
| 10. Explain how cloud computing can be used for social networking applications. | [10] | CO5 | L2 |


**PART-III**

| | | | |
|---|---|---|---|
| 5. Explain the list of open challenges in cloud computing. | [10] | CO4 | L2 |
| **(OR)** | | | |
| 6. Explain about the AWS with necessary diagram. | [10] | CO5 | L2 |
| **PART-IV** | | | |
| 7. Describe in detail about openstack. | [10] | CO5 | L2 |
| **(OR)** | | | |
| 8. Explain the architecture model of GAE with necessary diagram. | [10] | CO5 | L2 |
| **PART-V** | | | |
| 9. Explain how cloud computing can be used for ECG analysis. | [10] | CO5 | L2 |
| **(OR)** | | | |
| 10. Explain how cloud computing can be used for social networking applications. | [10] | CO5 | L2 |

**1. Explain about SaaS. (10)**

- a software delivery model that provides access to applications through the Internet as a Web-based service.
- It provides a means to free users from complex hardware and software management by offloading such tasks to third parties, which build applications accessible to multiple users through a Web browser
- customers neither need install anything on their premises nor have to pay considerable up-front costs to purchase the software and the required licenses
- simply access the application website, enter their credentials and billing details, and can instantly use the application
- On the provider side, the specific details and features of each customer's application are maintained in the infrastructure and made available on demand
- applications serving a wide range of users and that can be adapted to specific needs with little further customization
- "one-to-many" software delivery model, whereby an application is shared across multiple users
- Example: ERP
- SaaS applications are naturally multi-tenant - allows providers to centralize and sustain the effort of managing large hardware infrastructures, maintaining and upgrading applications transparently to the users, and optimizing resources by sharing the costs among the large user base
- On the customer side, such costs constitute a minimal fraction of the usage fee paid for the software
- ASPs already had some of the core characteristics of SaaS:
  - The product sold to customer is application access.
  - The application is centrally managed.
  - The service delivered is one-to-many.
  - The service delivered is an integrated solution delivered on the contract, which means *provided as promised*.
- ASPs provided access to packaged software solutions that addressed the needs of a variety of customers
- Web 2.0 technologies allowed turning the Web browser into a full-featured interface, able even to support application composition and development
- SaaS applications can serve different needs.
- CRM, ERP, and social networking applications are definitely the most popular ones.
- Example for CRM: *Salesforce.com*, builds on top of the Force.com platform, which provides a fully featured environment for building applications
- *RightNow* is customer experience-centered SaaS application that integrates together different features, from chat to Web communities, to support the common activity of an enterprise.
- Another important class of popular SaaS applications comprises social networking applications such as *Facebook* and professional networking sites such as *LinkedIn*

**2. Explain about the private cloud with necessary diagram. (10)**

- Private clouds: implemented within the private premises of an institution and generally made accessible to the members of the institution or a subset of them.
- Private clouds: implemented within the private premises of an institution and generally made accessible to the members of the institution or a subset of them.
- Key Advantages:
  - Customer Information Protection: security concerns are less critical

- Infrastructure ensuring SLAs: QoS by means of system monitoring and maintenance, and disaster recovery, and other uptime services
- Compliance with standard procedures and operations: for third party specific procedures when deploying and executing applications

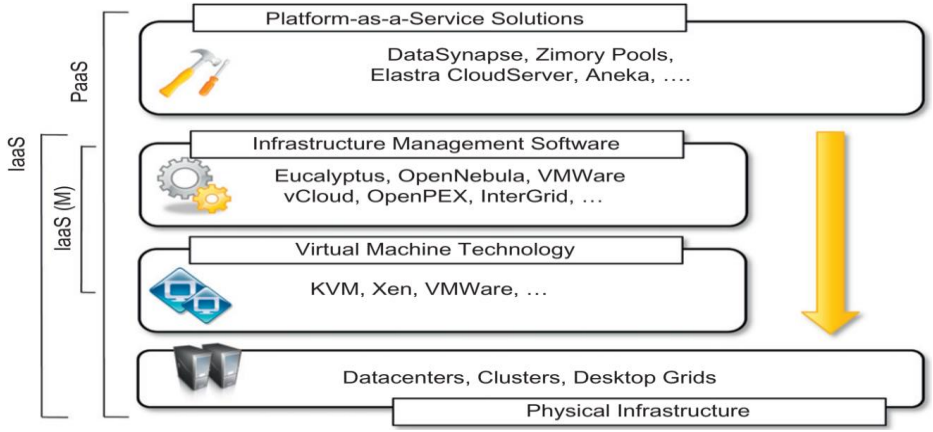


*Private clouds hardware and software stack*

- datacenter, a cluster, an enterprise desktop grid: existing IT infrastructure already deployed on the private premises
- The physical layer is complemented with infrastructure management software or a PaaS solution, according to the service delivered to the users of the cloud
- bottom layer of the software stack: serve as the foundations of the cloud
- VMware vCloud, Eucalyptus, and OpenNebula: used to control the virtual infrastructure and provide an IaaS solution
  - OpenNebula is an open-source solution for virtual infrastructure management
- Its modular architecture allows extending the software with additional features such as the capability of reserving virtual machine instances
- OpenPEX: Web-based system that allows the reservation of virtual machine instances and is designed to support different back ends
- InterGrid: provides added value on top of OpenNebula and Amazon EC2 by allowing the reservation of virtual machine instances and managing multi-administrative domain clouds
- PaaS solutions can provide an additional layer and deliver a high-level service for private clouds
  - DataSynapse: global provider of application virtualization software; provides a flexible environment for building private clouds on top of datacenters
  - Elastra Cloud Server: a platform for easily configuring and deploying distributed application infrastructures on clouds
  - Zimory: provides a software infrastructure layer that automates the use of resource pools based on Xen, KVM, and VMware virtualization technologies.
  - Aneka: a software development platform that can be used to deploy a cloud infrastructure on top of heterogeneous hardware: data-centers, clusters, and desktop grids.
  - provides a pluggable service-oriented architecture that's mainly devoted to supporting the execution of distributed applications with different programming models: bag of tasks, MapReduce, and others.
  - Private clouds can provide in-house solutions for cloud computing, but if compared to public clouds they exhibit more limited capability to scale elastically on demand

**3. Explain about the community cloud with necessary diagram. (10)**

- distributed systems created by integrating the services of different clouds to address the specific needs of an industry, a community, or a business sector.
- Community clouds are also different from private clouds, where the services are generally delivered within the institution that owns the cloud

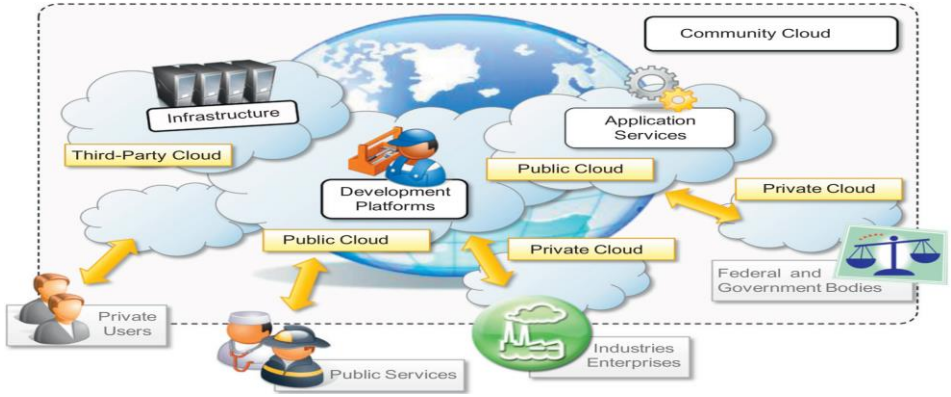


*A community cloud*

- Candidate sectors for community clouds are:
- Media industry: Community clouds can provide a shared environment where services can facilitate business-to-business collaboration and offer the horsepower in terms of aggregate bandwidth, CPU, and storage required to efficiently support media production
- Healthcare industry: community clouds can provide a global platform on which to share information and knowledge without revealing sensitive data maintained within the private infrastructure
- support the storing of patient-related data in a private cloud while using the shared infrastructure for noncritical services and automating processes within hospitals.
- Energy and other core industries: community clouds can bundle the comprehensive set of solutions that together vertically address management, deployment, and orchestration of services and operations
- Public sector: A community cloud can constitute the optimal venue to provide a distributed environment in which to create a communication platform for performing activities include invoice approval, infrastructure planning, and public hearings aimed at providing strategic solutions at local, national, and international administrative levels
- Scientific research: Science clouds are an interesting example of community clouds
- the common interest driving different organizations sharing a large distributed infrastructure is scientific computing
- The benefits of these community clouds are the following:
- Openness: community clouds are open systems in which fair competition between different solutions can happen
- Community: the infrastructure turns out to be more scalable because the system can grow simply by expanding its user base
- Graceful failures: Since there is no single provider or vendor in control of the infrastructure, there is no single point of failure
- Convenience and control: the cloud is shared and owned by the community, which makes all the decisions through a collective democratic process
- Environmental sustainability: community clouds have a smaller carbon footprint because it harnesses underutilized resources

**4. Explain about the economics of the cloud. (10)**

- cloud computing allows:
  - Reducing the capital costs associated to the IT infrastructure
  - Eliminating the depreciation or lifetime costs associated with IT capital assets
  - Replacing software licensing with subscriptions
  - Cutting the maintenance and administrative costs of IT resources
- In the case of a small startup, it is possible to completely leverage the cloud for many aspects, such as:
  - IT infrastructure
  - Software development
  - CRM and ERP
  - Another important aspect is the elimination of some indirect costs that are generated by IT assets, such as software licensing and support and carbon footprint emissions. (carbon footprints are taxable in Australia)
- In terms of the pricing models introduced by cloud computing, we can distinguish three different strategies that are adopted by the providers:

i. Tiered pricing: cloud services are offered in several tiers, each of which offers a fixed computing specification (Eg.: EC2 in AWS)

ii. Per-unit pricing: more suitable to cases where the principal source of revenue for the cloud provider is determined in terms of units of specific services, such as data transfer and memory allocation. (Eg.: GoGrid charges based on RAM/hr. units)

iii. Subscription-based pricing: used mostly by SaaS providers; users pay a periodic subscription fee for use of the software or the specific component services that are integrated in their applications.

**5. Explain the list of open challenges in cloud computing. (10)**

(i) Cloud interoperability and standards

- introducing standards and allowing interoperability between solutions offered by different vendors are objectives of fundamental importance
- Vendor lock-in constitutes one of the major strategic barriers against the seamless adoption of cloud computing at all stages
- Yet the first steps toward a standardization process have been made, and a few organizations:
  - the Cloud Computing Interoperability Forum,
  - the Open Cloud Consortium,
  - the DMTF Cloud Standards Incubator,
  - Open Cloud Manifesto - which embodies the point of view of various stakeholders on the benefits of open standards in the field
  - The Open Virtualization Format (OVF) - provide a common format for storing the information and metadata describing a virtual machine image
- the lack of a common set of APIs make the interaction with cloud-based solutions vendor specific

(ii) Scalability and fault tolerance

- Clouds allow scaling beyond the limits of the existing in-house IT resources, whether they are infrastructure (compute and storage) or applications services
- the ability to tolerate failure becomes fundamental, sometimes even more important than providing an extremely efficient and optimized system

(iii) Security, trust, and privacy

- The traditional cryptographic technologies are used to prevent data tampering and access to sensitive information.
- The massive use of virtualization technologies exposes the existing system to new threats, which previously were not considered applicable

- a lack of control over the environment in which the application is executed, which is made possible by leveraging the cloud

(iv) Organizational aspects

- Cloud computing introduces a significant change in the way IT services are consumed and managed
- storage, compute power, network infrastructure, and applications are delivered as metered services over the Internet
- a billing model that is new within typical enterprise IT departments, which requires a certain level of cultural and organizational process maturity
- a wide acceptance of cloud computing will require a significant change to business processes and organizational boundaries

**6. Explain about the AWS with necessary diagram. (10)**

- a platform that allows the development of flexible applications by providing solutions for elastic infrastructure scalability, messaging, and data storage
- Solutions stack:

(i) Compute Services

- constitute the fundamental element of cloud computing systems
- 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 pre-installed 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
  - EC2 instances: represent virtual machines.
  - provides virtual environments with the services they require to host applications
  - Six major categories:

i. Standard instances: computing power, storage, and memory
ii. Micro instances: used for small Web applications with limited traffic
iii. High-memory instances: need to process huge workloads and require large amounts of memory
iv. High-CPU instances: higher computing power than memory
v. Cluster Compute instances: extremely high I/O and network performance (HPC)
vi. Cluster GPU instances: heavy graphic computations, such as rendering clusters

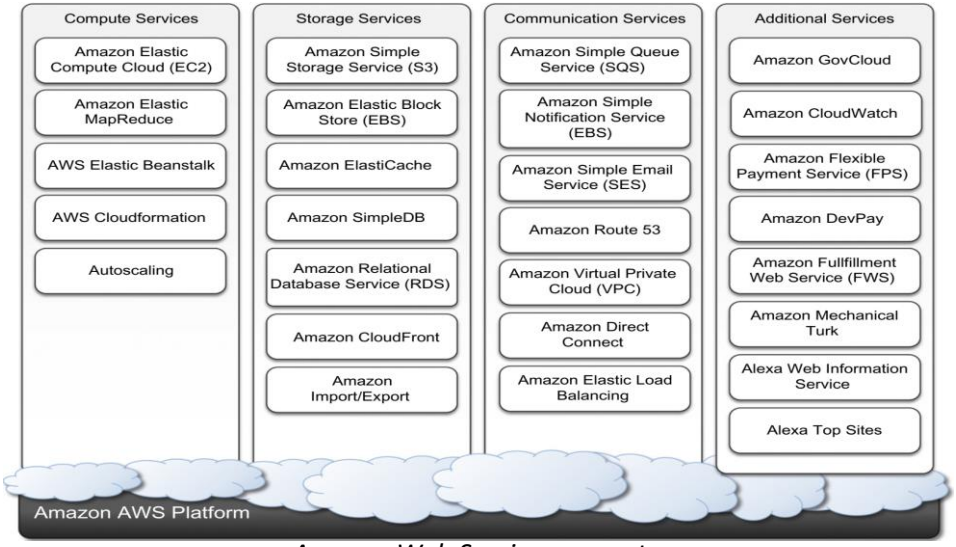


*Amazon Web Services ecosystem*

(ii) Storage Services

- AWS provides a collection of services for data storage and information management
- Two core components of S3 are: buckets and objects.
  i. Buckets: virtual containers in which to store objects
  ii. Objects represent the content that is actually stored. It can also be enriched with metadata that can be used to tag the stored content with additional information.
- 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 start-up.
  - 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)
- Amazon ElastiCache
  - an implementation of an elastic in-memory cache based on a cluster of EC2 instances.
  - It provides fast data access from other EC2 instances
  - It 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
- Amazon Relational Data Storage (RDS)
  - It 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
  - Two key advanced features of RDS: multi-AZ deployment and read replicas
    i. multi-AZ deployment: provides users with a failover infrastructure for their RDBMS solutions
    ii. read replicas: provides users with increased performance for applications that are heavily based on database reads
- Amazon SimpleDB
  - It is a lightweight, highly scalable, and flexible data storage solution for applications that do not require a fully relational model for their data
  - provides support for semi-structured data, the model for which is based on the concept of domains, items, and attributes
- Amazon CloudFront
  - It is an implementation of a Content Delivery Network (CDN) on top of the Amazon distributed storage infrastructure.
  - It leverages a collection of edge servers strategically located around the globe to better serve requests for static and streaming Web content so that

(iii) 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
  - 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

- 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)

(iv) Additional Services

- collection of services that allow users to utilize services in aggregation.
- The two relevant services are Amazon CloudWatch and Amazon Flexible Payment Service (FPS)
- Amazon CloudWatch:
  - provides a comprehensive set of statistics that help developers understand and optimize the behavior of their application hosted on AWS.
  - collects information from several other AWS services: EC2, S3, SimpleDB, CloudFront, and others.
  - Using CloudWatch, developers can see a detailed breakdown of their usage of the service they are renting on AWS and can devise more efficient and cost-saving applications

**7. Describe in detail about openstack. (10)**

- OpenStack is a collection of opensource software projects that can be collectively utilized to operate a cloud network infrastructure in order to provide IaaS. The OpenStack project began as a collaboration of Rackspace Hosting and NASA as an opensource project.
- Since Amazon's AWS was the first majorly used cloud service, OpenStack also makes its services available via Amazon EC2 and S3 compatible APIs. This ensures that all the existing tools that work with Amazon's cloud offerings, can work with deployments of OpenStack as well.
- The OpenStack project is combination of three main components:

a. OpenStack Compute (Nova) – is used to orchestrate, manage and offer virtual machines upon many hypervisors, including QEMU and KVM. This is analogous to the Amazon Elastic Compute Cloud (EC2).

b. OpenStack Object Store (Swift) – provides redundant storage for static objects. This service is scalable to massive data sizes and theoretically can provide infinite storage. It is analogous to the Amazon Simple Storage Service (S3).

c. OpenStack Image Service (Glance) – provides storage for virtual disk, kernel and images. Glance is also used to provide image registration and querying services. It is able to accept images in many formats, including the popular Amazon Machine Image (AMI), Amazon Kernel Image (AKI) and Amazon Ramdisk Image (ARI).

i. OpenStack Compute (Nova)

- Nova takes up the role of providing computing services within the OpenStack cloud. As such, any activity needed to support the life cycle of a virtual machine instance within the cloud is handled by Nova. This includes things like managing block storage, networking, scheduling, computing resources, authorization and hypervisors. However, Nova does not provide any virtualization capabilities by itself. It is designed to use libvirt APIs to interact with any supported hypervisors. This means that Nova is hypervisor agnostic and provides support for Xen, XenServer/XCP, KVM, UML, VMware vSphere and Hyper-V amongst others.

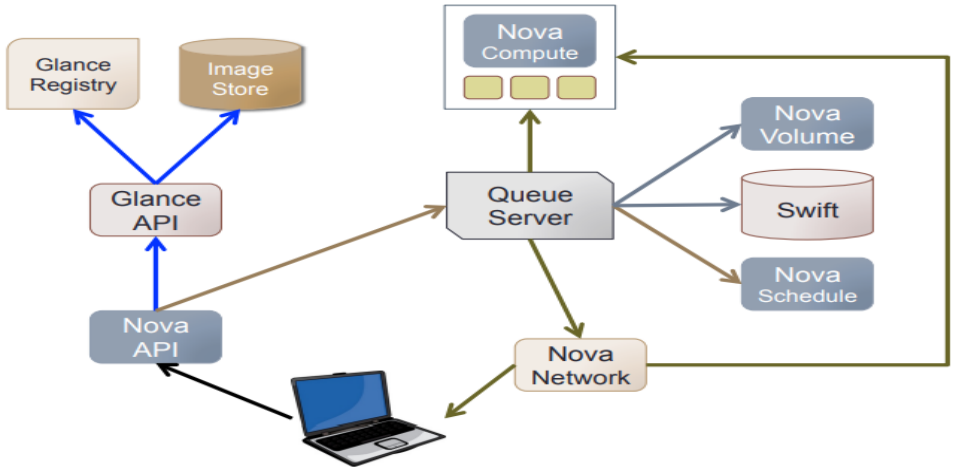


An overview of the OpenStack architecture

ii. API Services (nova-api)

- The nova-api service provides an interface to the outside world to interact with the cloud infrastructure. The API server is the only component that the outside world uses to manage the infrastructure. Management is done through RESTful calls using the EC2 API. The API server then, in turn, communicates with the relevant components of the cloud infrastructure by using the Message Queue.

iii. OpenStack Image Service (Glance)

- OpenStack Imaging Service is a lookup and retrieval system for virtual machine images. While it can be configured to use Swift or S3 storage to store the images, it normally uses a regular filesystem on the host for the glance service. The information regarding registered images is stored in an SQL database, which can be either MySQL, PostgreSQL, SQLite or many other varieties as well

**8. Explain the architecture model of GAE with necessary diagram. (10)**

- It 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 platform is logically divided into four major components:
  i. infrastructure,
  ii. the runtime environment,
  iii. the underlying storage, and
  iv. the set of scalable services that can be used to develop applications

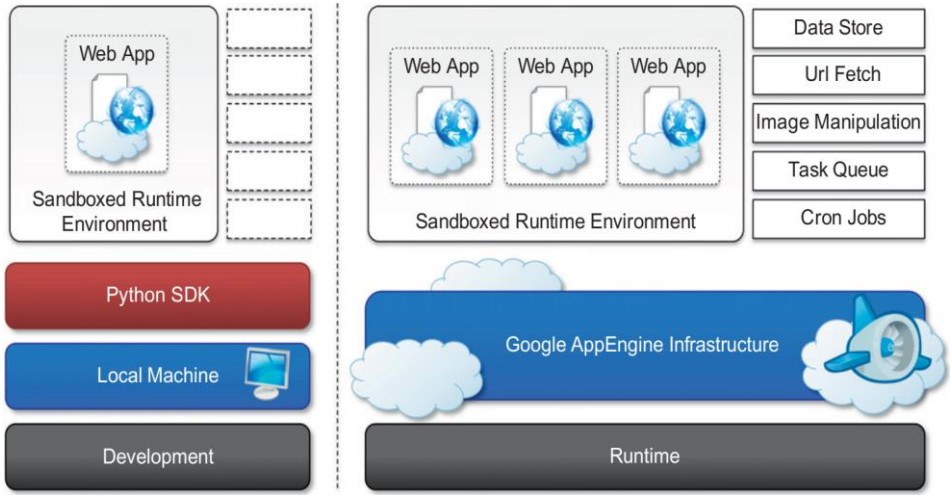


*Google AppEngine platform architecture*

(i) Infrastructure

- AppEngine hosts Web applications, and its primary function is to serve users requests efficiently
- AppEngine's infrastructure takes advantage of many servers available within Google data-centers.
- 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

(ii) Runtime environment

- represents the execution context of applications hosted on AppEngine
- the runtime comes into existence when the request handler starts executing and terminates once the handler has completed

- One of the major responsibilities of the runtime environment is to provide the application environment with an isolated and protected context (execute without causing a threat to the server)

(iii) Storage

- There are three different levels of storage: in memory-cache, storage for semi-structured data, and long-term storage for static data
- MemCache: 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

▪ Sandboxing

- sandboxing is achieved by means of modified runtimes for applications that disable some of the common features normally available with their default implementations.

▪ 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

▪ 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

▪ Task queues

- 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
- The queue is designed to re-execute the task in case of failure in order to avoid transient failures preventing the task from a successful completion
- allow applications to submit a task for a later execution

▪ 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.

**9. Explain how cloud computing can be used for ECG analysis. (10)**

- Healthcare is a domain in which computer technology has found several and diverse applications: from supporting the business functions to assisting scientists in developing solutions to cure diseases.
- An important application is the use of cloud technologies to support doctors in providing more effective diagnostic processes.
- Cloud computing technologies allow the remote monitoring of a patient's heartbeat data, data analysis in minimal time, and the notification of first-aid personnel and doctors should these data reveal potentially dangerous conditions.
  - The Web service forms the front-end of a platform that is entirely hosted in the cloud and that leverages the three layers of the cloud computing stack: SaaS, PaaS, and IaaS.
  - The Web service constitute the SaaS application that will store ECG data in the Amazon S3 service and issue a processing request to the scalable cloud platform.
  - The runtime platform is composed of a dynamically sizable number of instances running the workflow engine and Aneka.

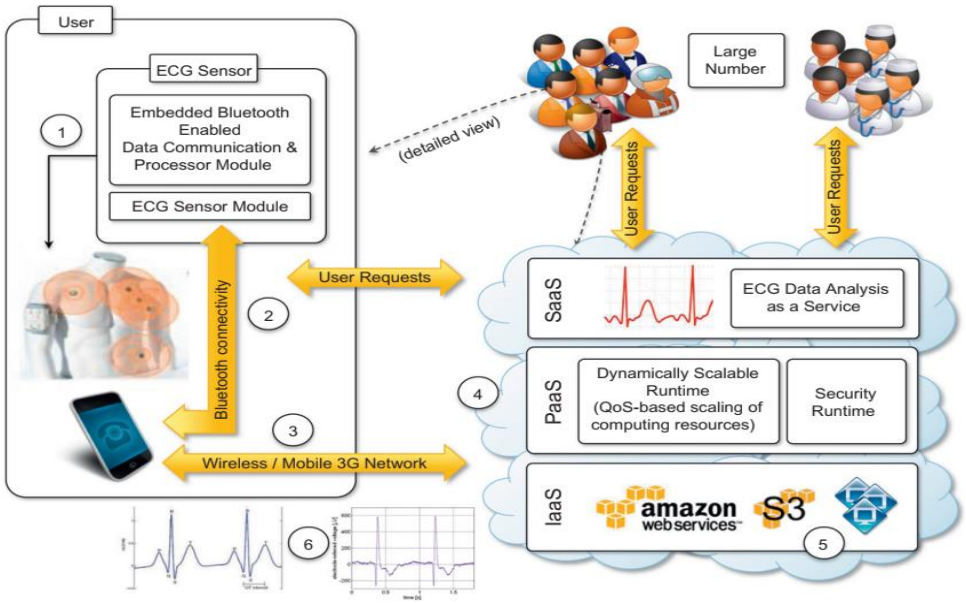


*An online health monitoring system hosted in the cloud*

- Advantages:
  i. the elasticity of the cloud infrastructure that can grow and shrink according to the requests served
  ii. ubiquity - promise to deliver systems with minimum or no downtime
  iii. cost savings constitute another reason for the use of cloud technology in healthcare
- Cloud services are priced on a pay-per-use basis and with volume prices for large numbers of service requests.
- These two models provide a set of flexible options that can be used to price the service, thus actually charging costs based on effective use rather than capital costs.

**10. Explain how cloud computing can be used for social networking applications. (10)**

- Facebook
  - Facebook is probably the most evident and interesting environment in social networking.
  - the social network is backed by two data centers that have been built and optimized to reduce costs and impact on the environment
- a completely customized stack of opportunely modified and refined open-source technologies constitutes the back-end of the largest social network
- these technologies constitute a powerful platform for developing cloud applications.
- This platform primarily supports Facebook itself and offers APIs to integrate third-party applications with Facebook's core infrastructure to deliver additional services such as social games and quizzes created by others.
- The reference stack serving Facebook is based on LAMP (Linux, Apache, MySQL, and PHP).
- Most of the user data are served by querying a distributed cluster of MySQL instances, which mostly contain key-value pairs.
- The development of services is facilitated by a set of internally developed tools.
- Thrift: a collection of abstractions (and language bindings) that allow cross-language development.
- It allows services developed in different languages to communicate and exchange data.
- Other relevant services and tools are Scribe, which aggregates streaming log feeds, and applications for alerting and monitoring.