CLOUD COMPUTING WITH AWS

M.Sc.(IT & CA)

Geetajnali College Of Computer Science & Commerce (BBA)
Rajkot.

Prepared By : Er. Harsh Joshi

- Introduction of cloud computing
  - Cloud Computing is the delivery of computing services such as servers, storage, databases, networking, software, analytics, intelligence, and more, over the Cloud (Internet).

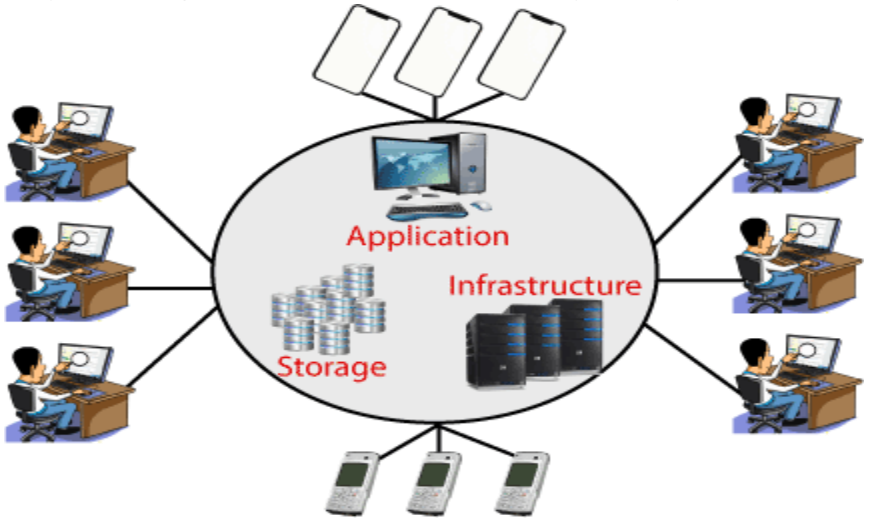

  - Cloud Computing provides an alternative to the on-premises datacenter. With an on-premises datacenter, we have to manage everything, such as purchasing and installing hardware, virtualization, installing the operating system, and any other required applications, setting up the network, configuring the firewall, and setting up storage for data. After doing all the set-up, we become responsible for maintaining it through its entire lifecycle.
  - But if we choose Cloud Computing, a cloud vendor is responsible for the hardware purchase and maintenance. They also provide a wide variety of software and platform as a service. We can take any required services on rent. The cloud computing services will be charged based on usage.

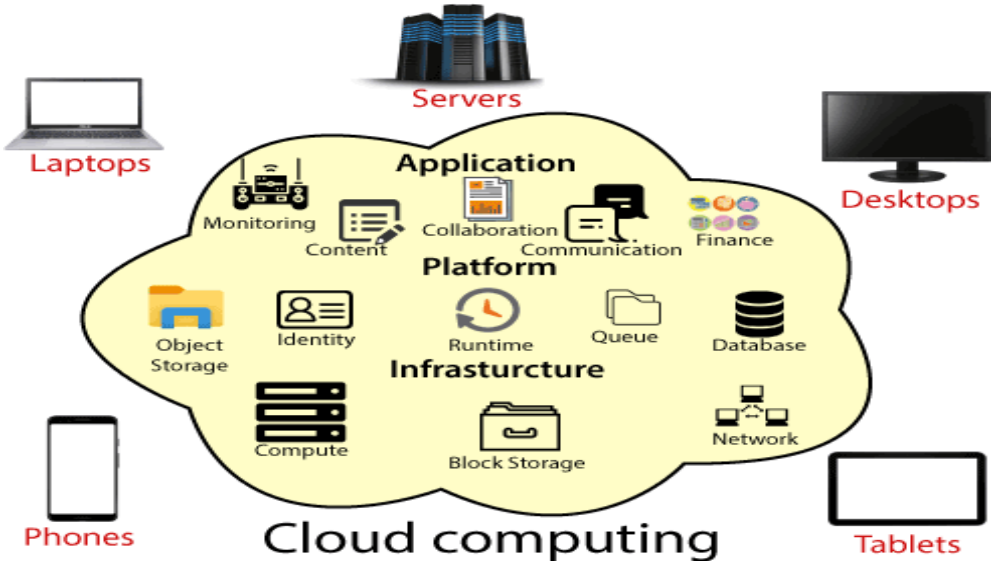

  - The cloud environment provides an easily accessible online portal that makes handy for the user to manage the compute, storage, network, and application resources. Some cloud service providers are in the following figure.
  - Advantages of cloud computing
    - Cost: It reduces the huge capital costs of buying hardware and software.
    - Speed: Resources can be accessed in minutes, typically within a few clicks.
    - Scalability: We can increase or decrease the requirement of resources according to the business requirements.

- Productivity: While using cloud computing, we put less operational effort. We do not need to apply patching, as well as no need to maintain hardware and software. So, in this way, the IT team can be more productive and focus on achieving business goals.
- Reliability: Backup and recovery of data are less expensive and very fast for business continuity.
- Security: Many cloud vendors offer a broad set of policies, technologies, and controls that strengthen our data security.

- How Cloud Computing Works
  - Assume that you are an executive at a very big corporation. Your particular responsibilities include to make sure that all of your employees have the right hardware and software they need to do their jobs. To buy computers for everyone is not enough. You also have to purchase software as well as software licenses and then provide these software's to your employees as they require. Whenever you hire a new employee, you need to buy more software or make sure your current software license allows another user. It is so stressful that you have to spend lots of money.
  - But, there may be an alternative for executives like you. So, instead of installing a suite of software for each computer, you just need to load one application. That application will allow the employees to log-in into a Web-based service which hosts all the programs for the user that is required for his/her job. Remote servers owned by another company and that will run everything from e-mail to word processing to complex data analysis programs. It is called cloud computing, and it could change the entire computer industry.

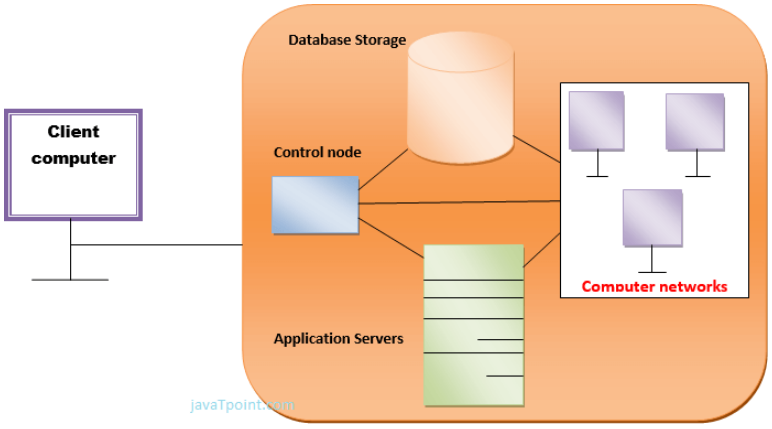

  - In a cloud computing system, there is a significant workload shift. Local computers have no longer to do all the heavy lifting when it comes to run applications. But cloud computing can handle that much heavy load easily and automatically. Hardware and software demands on the user's side decrease. The only thing the user's computer requires to be able to run is the cloud computing interface software of the system, which can be as simple as a Web browser and the cloud's network takes care of the rest.

- Types Of Cloud

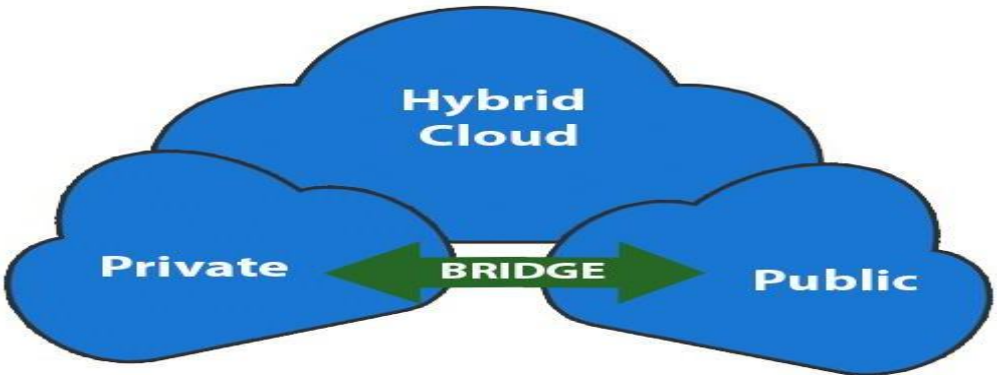

- Public Cloud: The cloud resources that are owned and operated by a third-party cloud service provider are termed as public clouds. It delivers computing resources such as servers, software, and storage over the internet
- Private Cloud: The cloud computing resources that are exclusively used inside a single business or organization are termed as a private cloud. A private cloud may physically be located on the company's on-site datacenter or hosted by a third-party service provider.
- Hybrid Cloud: It is the combination of public and private clouds, which is bounded together by technology that allows data applications to be shared between them. Hybrid cloud provides flexibility and more deployment options to the business.

- What is Virtualization
  - Virtualization is the "creation of a virtual (rather than actual) version of something, such as a server, a desktop, a storage device, an operating system or network resources".

  - In other words, Virtualization is a technique, which allows to share a single physical instance of a resource or an application among multiple customers and organizations. It does by assigning a logical name to a physical storage and providing a pointer to that physical resource when demanded.

- Advantages of Cloud

  As we all know that Cloud computing is trending technology. Almost every company switched their services on the cloud to rise the company growth.

  Here, we are going to discuss some important advantages of Cloud Computing:

  - Back-up and restore data
    Once the data is stored in the cloud, it is easier to get back-up and restore that data using the cloud.
  - Improved collaboration
    Cloud applications improve collaboration by allowing groups of people to quickly and easily share information in the cloud via shared storage.
  - Excellent accessibility
    Cloud allows us to quickly and easily access store information anywhere, anytime in the whole world, using an internet connection. An internet cloud infrastructure increases organization productivity and efficiency by ensuring that our data is always accessible.
  - Low maintenance cost
    Cloud computing reduces both hardware and software maintenance costs for organizations.
  - Mobility
    Cloud computing allows us to easily access all cloud data via mobile.
  - Services in the pay-per-use model
    Cloud computing offers Application Programming Interfaces (APIs) to the users for access services on the cloud and pays the charges as per the usage of service.
  - Unlimited storage capacity
    Cloud offers us a huge amount of storing capacity for storing our important data such as documents, images, audio, video, etc. in one place.
  - Data security
    Data security is one of the biggest advantages of cloud computing. Cloud offers many advanced features related to security and ensures that data is securely stored and handled.

- AWS History
  - 2003: In 2003, Chris Pinkham and Benjamin Black presented a paper on how Amazon's own internal infrastructure should look like. They suggested to sell it as a service and prepared a business case on it.

They prepared a six-page document and had a look over it to proceed with it or not. They decided to proceed with the documentation.

- 2004: SQS stands for "Simple Queue Service" was officially launched in 2004. A team launched this service in Cape Town, South Africa.
- 2006: AWS (Amazon Web Services) was officially launched.
- 2007: In 2007, over 180,000 developers had signed up for the AWS.
- 2010: In 2010, amazon.com retail web services were moved to the AWS, i.e., amazon.com is now running on AWS.
- 2011: AWS suffered from some major problems. Some parts of volume of EBS (Elastic Block Store) was stuck and were unable to read and write requests. It took two days for the problem to get resolved.
- 2012: AWS hosted a first customer event known as re:Invent conference. First re:invent conference occurred in which new products were launched. In AWS, another major problem occurred that affects many popular sites such as Pinterest, Reddit, and Foursquare.
- 2013: In 2013, certifications were launched. AWS started a certifications program for software engineers who had expertise in cloud computing.
- 2014: AWS committed to achieve 100% renewable energy usage for its global footprint.
- 2015: AWS breaks its revenue and reaches to $6 Billion USD per annum. The revenue was growing 90% every year.
- 2016: By 2016, revenue doubled and reached $13Billion USD per annum.
- 2017: In 2017, AWS re: invent releases a host of Artificial Intelligence Services due to which revenue of AWS doubled and reached $27 Billion USD per annum.
- 2018: In 2018, AWS launched a Machine Learning Speciality Certs. It heavily focussed on automating Artificial Intelligence and Machine learning.

- AWS Dashboard / Console

  Let's look at it this way, imagine you want to access all the features of your Facebook account; You will need an interface where you can input your username and password to gain access. You can then navigate your updates, messages, and other account information. Similarly, the AWS console works the same way. AWS Console is a web application that allows users to access Amazon Web Services. The console can be considered the backbone or basic web infrastructure through which Amazon Web Services can be accessed. Without the console presented so that users can easily navigate to every Amazon web service, it won't be easy to have centralized access to all the Amazon web services.

  The console provides an inbuilt interface for users to perform tasks like provisioning resources, lets you launch instances, and works with Amazon S3 buckets. We'll learn more about features and discover significant other configurations as we move forward.

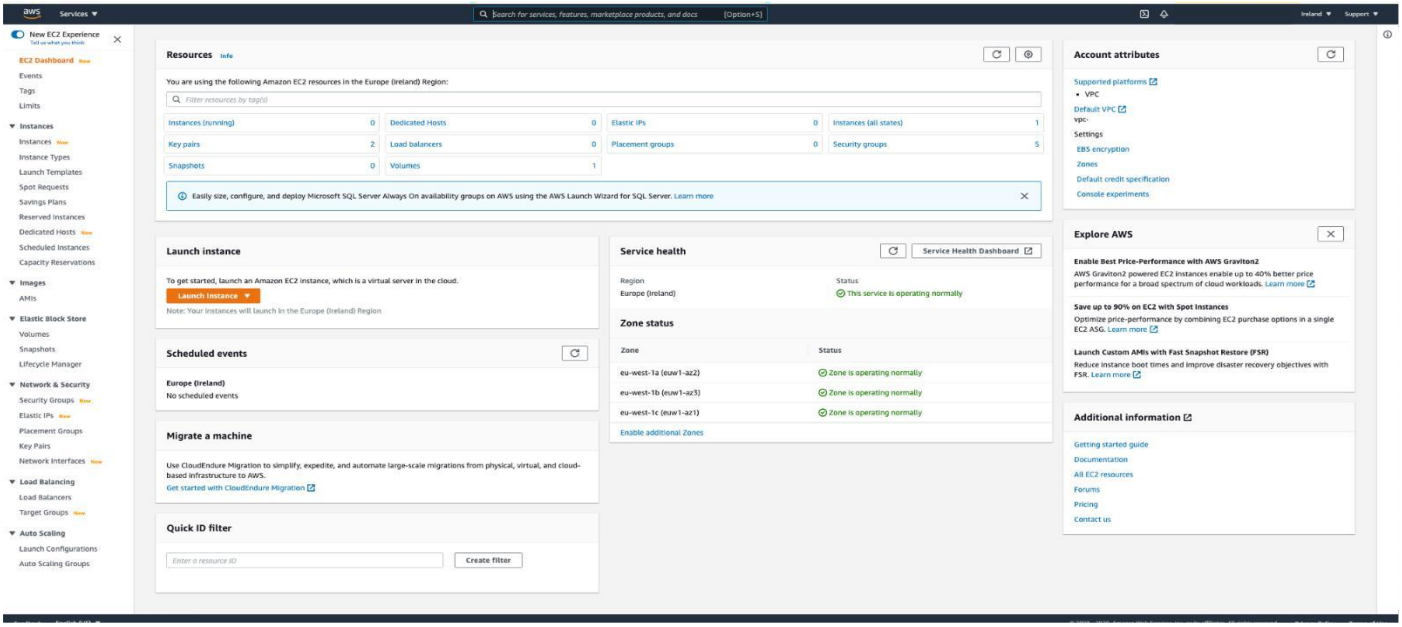

- AWS Overview
  - AWS stands for Amazon Web Services.
  - The AWS service is provided by the Amazon that uses distributed IT infrastructure to provide different IT resources available on demand. It provides different services such as infrastructure as a service (IaaS), platform as a service (PaaS) and packaged software as a service (SaaS).
  - Amazon launched AWS, a cloud computing platform to allow the different organizations to take advantage of reliable IT infrastructure.

- AWS Architecture

  This is the basic structure of AWS EC2, where EC2 stands for Elastic Compute Cloud. EC2 allow users to use virtual machines of different configurations as per their requirement. It allows various configuration options, mapping of individual server, various pricing options, etc. We will discuss these in detail in AWS Products section. Following is the diagrammatic representation of the architecture.
  - Architecture

    Note – In the above diagram S3 stands for Simple Storage Service. It allows the users to store and retrieve various types of data using API calls. It doesn't contain any computing element. We will discuss this topic in detail in AWS products section.
  - Load Balancing

    Load balancing simply means to hardware or software load over web servers, that improver's the efficiency of the server as well as the application. Following is the diagrammatic representation of AWS architecture with load balancing.Hardware load balancer is a very common network appliance used in traditional web application architectures.

    AWS provides the Elastic Load Balancing service, it distributes the traffic to EC2 instances across multiple available sources, and dynamic addition and removal of Amazon EC2 hosts from the load-balancing rotation.

    Elastic Load Balancing can dynamically grow and shrink the load-balancing capacity to adjust to traffic demands and also support sticky sessions to address more advanced routing needs.
  - Amazon Cloud-front

    It is responsible for content delivery, i.e. used to deliver website. It may contain dynamic, static, and streaming content using a global network of edge locations. Requests for content at the user's end are automatically routed to the nearest edge location, which improves the performance.

    Amazon Cloud-front is optimized to work with other Amazon Web Services, like Amazon S3 and Amazon EC2. It also works fine with any non-AWS origin server and stores the original files in a similar manner.

In Amazon Web Services, there are no contracts or monthly commitments. We pay only for as much or as little content as we deliver through the service.

- Elastic Load Balancer
  It is used to spread the traffic to web servers, which improves performance. AWS provides the Elastic Load Balancing service, in which traffic is distributed to EC2 instances over multiple available zones, and dynamic addition and removal of Amazon EC2 hosts from the load-balancing rotation.
  Elastic Load Balancing can dynamically grow and shrink the load-balancing capacity as per the traffic conditions.
- Security Management
  Amazon's Elastic Compute Cloud (EC2) provides a feature called security groups, which is similar to an inbound network firewall, in which we have to specify the protocols, ports, and source IP ranges that are allowed to reach your EC2 instances.
  Each EC2 instance can be assigned one or more security groups, each of which routes the appropriate traffic to each instance. Security groups can be configured using specific subnets or IP addresses which limits access to EC2 instances.
- Elastic Caches
  Amazon Elastic Cache is a web service that manages the memory cache in the cloud. In memory management, cache has a very important role and helps to reduce the load on the services, improves the performance and scalability on the database tier by caching frequently used information.

- Amazon RDS
  Amazon RDS (Relational Database Service) provides a similar access as that of MySQL, Oracle, or Microsoft SQL Server database engine. The same queries, applications, and tools can be used with Amazon RDS.
  It automatically patches the database software and manages backups as per the user's instruction. It also supports point-in-time recovery. There are no up-front investments required, and we pay only for the resources we use.
- Hosting RDMS on EC2 Instances
  Amazon RDS allows users to install RDBMS (Relational Database Management System) of your choice like MySQL, Oracle, SQL Server, DB2, etc. on an EC2 instance and can manage as required.
  Amazon EC2 uses Amazon EBS (Elastic Block Storage) similar to network-attached storage. All data and logs running on EC2 instances should be placed on Amazon EBS volumes, which will be available even if the database host fails.
  Amazon EBS volumes automatically provide redundancy within the availability zone, which increases the availability of simple disks. Further if the volume is not sufficient for our databases needs, volume can be added to increase the performance for our database.
  Using Amazon RDS, the service provider manages the storage and we only focus on managing the data.
- Storage & Backups
  AWS cloud provides various options for storing, accessing, and backing up web application data and assets. The Amazon S3 (Simple Storage Service) provides a simple web-services interface that can be used to store and retrieve any amount of data, at any time, from anywhere on the web.
  Amazon S3 stores data as objects within resources called buckets. The user can store as many objects as per requirement within the bucket, and can read, write and delete objects from the bucket.
  Amazon EBS is effective for data that needs to be accessed as block storage and requires persistence beyond the life of the running instance, such as database partitions and application logs.
  Amazon EBS volumes can be maximized up to 1 TB, and these volumes can be striped for larger volumes and increased performance. Provisioned IOPS volumes are designed to meet the needs of database workloads that are sensitive to storage performance and consistency.
  Amazon EBS currently supports up to 1,000 IOPS per volume. We can stripe multiple volumes together to deliver thousands of IOPS per instance to an application.
- Auto Scaling

The difference between AWS cloud architecture and the traditional hosting model is that AWS can dynamically scale the web application fleet on demand to handle changes in traffic.
In the traditional hosting model, traffic forecasting models are generally used to provision hosts ahead of projected traffic. In AWS, instances can be provisioned on the fly according to a set of triggers for scaling the fleet out and back in. Amazon Auto Scaling can create capacity groups of servers that can grow or shrink on demand.

- Software as a Service (SaaS) introduction
  SaaS is also known as "On-Demand Software". It is a software distribution model in which services are hosted by a cloud service provider. These services are available to end-users over the internet so, the end-users do not need to install any software on their devices to access these services.

- SaaS Model

There are the following services provided by SaaS providers

- Business Services - SaaS Provider provides various business services to start-up the business. The SaaS business services include ERP (Enterprise Resource Planning), CRM (Customer Relationship Management), billing, and sales.
- Document Management - SaaS document management is a software application offered by a third party (SaaS providers) to create, manage, and track electronic documents.
- Social Networks - As we all know, social networking sites are used by the general public, so social networking service providers use SaaS for their convenience and handle the general public's information.
- Mail Services - To handle the unpredictable number of users and load on e-mail services, many e-mail providers offering their services using SaaS.

- SaaS Integration Services
  Almost every company in the world uses software as a service (SaaS) applications — Salesforce is currently the largest provider of SaaS apps, followed by Microsoft, Adobe, Box, and Amazon. You probably use dozens of them every single day, and you're not alone. Even small non-digital businesses, such as your local coffee shop, use them. They might use Shopify to take online orders, Square to process payments, and Google Docs to share their employee's schedules.

  SaaS Integration Benefits
  - It saves you time
    This benefit is fairly obvious. When your business tools can transfer information automatically, you don't need to spend the time and resources doing it yourself. Set up the integration and that's it. No more manual entry.
  - It reduces or eliminates human error
    Dealing with mistakes that happen during manual data entry can be one of the costliest aspects of running a business. It takes resources to track down the error and correct it, and the error itself may have caused delays and unhappy customers for your business.

By automatically sending the correct information from one application to another, SaaS integration removes the middle-person. The information from one tool goes straight to the other, eliminating the chance of human error.

- It gives your teams visibility
  Integration allows for the right information to be at the right place at the right time. The information can be automatically kept up to date in the tool of your choice. This gives your team the visibility and data they need to make decisions and do their job.
  In the Shopify/FreshBooks example above, your accounting team can find what they need in FreshBooks, the tool they likely use the most, without having to check Shopify.
- It improves customer service
  People expect things fast these days, without error. When someone books a hotel room online, they expect a confirmation email and the front desk clerk to have their reservation. If they get a refund on a purchase, the money should be transferred to them immediately, along with another email.
  When your business tools are connected, these processes are ready to go, allowing you to provide the best customer service possible.
- It makes processes scalable
  When companies grow, some of the aforementioned functions can get unwieldy. But when your apps are integrated, they can automatically pass information back and forth, without needing humans to do it — and without copy errors.
  For any company that is trying to scale, integration is a key component of any automation strategy. It's an absolute must for large enterprises, which often use 100+ business apps across dozens of departments.

- Advantages & Disadvantages Of SaaS
  - Advantages
    - SaaS is easy to buy
      SaaS pricing is based on a monthly fee or annual fee subscription, so it allows organizations to access business functionality at a low cost, which is less than licensed applications.
      Unlike traditional software, which is sold as a licensed based with an up-front cost (and often an optional ongoing support fee), SaaS providers are generally pricing the applications using a subscription fee, most commonly a monthly or annually fee.
    - One to Many
      SaaS services are offered as a one-to-many model means a single instance of the application is shared by multiple users.
    - Less hardware required for SaaS
      The software is hosted remotely, so organizations do not need to invest in additional hardware.
    - Low maintenance required for SaaS
      Software as a service removes the need for installation, set-up, and daily maintenance for the organizations. The initial set-up cost for SaaS is typically less than the enterprise software. SaaS vendors are pricing their applications based on some usage parameters, such as a number of users using the application. So SaaS does easy to monitor and automatic updates.
    - No special software or hardware versions required
      All users will have the same version of the software and typically access it through the web browser. SaaS reduces IT support costs by outsourcing hardware and software maintenance and support to the IaaS provider.
    - Multidevice support
      SaaS services can be accessed from any device such as desktops, laptops, tablets, phones, and thin clients.
    - API Integration
      SaaS services easily integrate with other software or services through standard APIs.

- No client-side installation
  SaaS services are accessed directly from the service provider using the internet connection, so do not need to require any software installation.

- Disadvantages
  - Security
    Actually, data is stored in the cloud, so security may be an issue for some users. However, cloud computing is not more secure than in-house deployment.
  - Latency issue
    Since data and applications are stored in the cloud at a variable distance from the end-user, there is a possibility that there may be greater latency when interacting with the application compared to local deployment. Therefore, the SaaS model is not suitable for applications whose demand response time is in milliseconds.
  - Total Dependency on Internet
    Without an internet connection, most SaaS applications are not usable.
  - Switching between SaaS vendors is difficult
    Switching SaaS vendors involves the difficult and slow task of transferring the very large data files over the internet and then converting and importing them into another SaaS also.

- Infrastructure As A Service
  - Introduction
    Infrastructure as a Service (IaaS) is a business model that delivers IT infrastructure like compute, storage, and network resources on a pay-as-you-go basis over the internet. You can use IaaS to request and configure the resources you require to run your applications and IT systems.

  - Virtual Machines
    In computing, a "virtual machine" is the virtualization or emulation of a computer system. Virtual machines are based on computer architectures and provide the functionality of a physical computer. Their implementations may involve specialized hardware, software, or a combination of the two.
  - VM Migration Services
    A virtual machine (VM) is a software-defines computer that is the same as the physical computer. The critical feature of VM is that it is independent of the physical computer and also portable. These characteristics make VM popular among over the globe for both on premise and cloud environments.
    Today we will explore virtual machine migration to the cloud and much more. Let's dive in!
    Types of VM Migration Services
    VM migration is divided into three parts mentioned below:
    - VM migration to the cloud with new environment setup
    - Lift and Shift VM migration to cloud with a small downtime window
    - Lift and Shift VM migration to cloud with a large downtime window.
  - Advantages Of IAAS
    IaaS is advantageous to companies in scenarios where scalability and quick provisioning are key. In other words, organizations experiencing rapid growth but lacking the capital to invest in hardware are great candidates for IaaS models. IaaS can also be beneficial to companies with steady application workloads that simply want to offload some of the routine operations and maintenance involved in managing infrastructure.
    Other advantages may include the following:
    - Pay for What You Use: Fees are computed via usage-based metrics
    - Reduce Capital Expenditures: IaaS is typically a monthly operational expense
    - Dynamically Scale: Rapidly add capacity in peak times and scale down as needed

- Increase Security: IaaS providers invest heavily in security technology and expertise
- Future-Proof: Access to state-of-the-art data center, hardware and operating systems
- Self-Service Provisioning: Access via simple internet connection
- Reallocate IT Resources: Free up IT staff for higher value projects
- Reduce Downtime: IaaS enables instant recovery from outages
- Boost Speed: Developers can begin projects once IaaS machines are provisioned
- Enable Innovation: Add new capabilities and leverage APIs
- Level the Playing Field: SMBs can compete with much larger firms

- Disadvantages Of IAAS

  There are many benefits to using IaaS in an organization, but there are also challenges. Some of these hurdles can be overcome with advanced preparation, but others present risks that a customer should weigh in on before deployment.

  Challenges may include the following:
  - Unexpected Costs: Monthly fees can add up, or peak usage may be more than expected
  - Process Changes: IaaS may require changes to processes and workflows
  - Runaway Inventory: Instances may be deployed, but not taken down
  - Security Risks: While IaaS providers secure the infrastructure, businesses are responsible for anything they host
  - Lack of Support: Live help is sometimes hard to come by
  - Complex Integration: Challenges with interaction with existing systems
  - Security Risks: New vulnerabilities may emerge around the loss of direct control
  - Limited Customization: Public cloud users may have limited control and ability to customize
  - Vendor Lock-In: Moving from one IaaS provider to another can be challenging
  - Broadband Dependency: Only as good as the reliability of the internet connection
  - Providers Not Created Equally: Vendor vetting and selection can be challenging
  - Managing Availability: Even the largest service providers experience downtime
  - Confusing SLAs: Service level agreements (SLAs) can be difficult to understand
  - Regulatory Uncertainty: Evolving federal and state laws can impact some industries' use of IaaS, especially across country borders
  - Vendor Consolidation: Providers may be acquired or go out of business
  - Third-Party Expertise: Lack of mature service providers, guidance or ecosystem support.

- Platform As A Service
  - Introduction

    Platform as a service. Platform as a service (PaaS) is a complete development and deployment environment in the cloud, with resources that enable you to deliver everything from simple cloud-based apps to sophisticated, cloud-enabled enterprise applications.
  - Integration Of Private & Public Cloud

    The term cloud computing spans a range of classifications, types, and architecture models. This networked computing model has transformed how we work—you're likely already using the cloud. But the cloud isn't one thing—cloud computing can be categorized into three general types:
    - Public cloud is cloud computing that's delivered via the internet and shared across organizations.
    - Private cloud is cloud computing that is dedicated solely to your organization.
    - Hybrid cloud is any environment that uses both public and private clouds.
  - Advantages Of PAAS

    PaaS works well for small businesses and startup companies for two very basic reasons. First, it's cost effective, allowing smaller organizations access to state-of-the-art resources without

the big price tag. Most small firms have never been able to build robust development environments on premises, so PaaS provides a path for accelerating software development. Second, it allows companies to focus on what they specialize in without worrying about maintaining basic infrastructure.

Other advantages include the following:

- Cost Effective: No need to purchase hardware or pay expenses during downtime
- Time Savings: No need to spend time setting up/maintaining the core stack
- Speed to Market: Speed up the creation of apps
- Future-Proof: Access to state-of-the-art data center, hardware and operating systems
- Increase Security: PaaS providers invest heavily in security technology and expertise
- Dynamically Scale: Rapidly add capacity in peak times and scale down as needed
- Custom Solutions: Operational tools in place so developers can create custom software
- Flexibility: Allows employees to log in and work on applications from anywhere

- Disadvantages Of PAAS

  There are always two sides to every story. While it's easy to make the case for PaaS, there's bound to be some challenges as well. Some of these hurdles are simply the flip side of the positives and the nature of the beast. Others can be overcome with advanced planning and preparation.

  Challenges may include the following:

  - Vendor Dependency: Very dependent upon the vendor's capabilities
  - Risk of Lock-In: Customers may get locked into a language, interface or program they no longer need
  - Compatibility: Difficulties may arise if PaaS is used in conjunction with existing development platforms
  - Security Risks: While PaaS providers secure the infrastructure and platform, businesses are responsible for security of the applications they build.

- IAM Overview and Policies
  You manage access in AWS by creating policies and attaching them to IAM identities (users, groups of users, or roles) or AWS resources. A policy is an object in AWS that, when associated with an identity or resource, defines their permissions. AWS evaluates these policies when an IAM principal (user or role) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in AWS as JSON documents. AWS supports six types of policies: identity-based policies, resource-based policies, permissions boundaries, Organizations SCPs, ACLs, and session policies.

  - Policies
    The following policy types, listed in order from most frequently used to less frequently used, are available for use in AWS. For more details, see the sections below for each policy type.
    - Identity-based policies – Attach managed and inline policies to IAM identities (users, groups to which users belong, or roles). Identity-based policies grant permissions to an identity.
    - Resource-based policies – Attach inline policies to resources. The most common examples of resource-based policies are Amazon S3 bucket policies and IAM role trust policies. Resource-based policies grant permissions to the principal that is specified in the policy. Principals can be in the same account as the resource or in other accounts.
    - Permissions boundaries – Use a managed policy as the permissions boundary for an IAM entity (user or role). That policy defines the maximum permissions that the identity-based policies can grant to an entity, but does not grant permissions. Permissions boundaries do not define the maximum permissions that a resource-based policy can grant to an entity.
    - Organizations SCPs – Use an AWS Organizations service control policy (SCP) to define the maximum permissions for account members of an organization or organizational unit (OU). SCPs limit permissions that identity-based policies or resource-based policies grant to entities (users or roles) within the account, but do not grant permissions.
    - Access control lists (ACLs) – Use ACLs to control which principals in other accounts can access the resource to which the ACL is attached. ACLs are similar to resource-based policies, although they are the only policy type that does not use the JSON policy document structure. ACLs are cross-account permissions policies that grant permissions to the specified principal. ACLs cannot grant permissions to entities within the same account.
    - Session policies – Pass advanced session policies when you use the AWS CLI or AWS API to assume a role or a federated user. Session policies limit the permissions that the role or user's identity-based policies grant to the session. Session policies limit permissions for a created session, but do not grant permissions. For more information, see Session Policies.

- IAM Users
  An IAM user is a resource in IAM that has associated credentials and permissions. An IAM user can represent a person or an application that uses its credentials to make AWS requests. This is typically referred to as a service account.

- IAM Groups
  An IAM group is an identity that specifies a collection of IAM users. You can't use a group to sign-in. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier to manage for large sets of users.

- Access Keys And Secret Access Keys
  Access keys consist of an access key ID and secret access key, which are used to sign programmatic requests that you make to AWS. If you don't have access keys, you can create them by using the IAM console at https://console.aws.amazon.com/iam/

- MFA

  Use the IAM console to check whether an AWS account root user or IAM user has a valid MFA device enabled.

  To check the MFA status of a root user
  - Sign in to the AWS Management Console with your root user credentials and then open the IAM console at https://console.aws.amazon.com/iam/
  - In the navigation bar on the upper right, choose your user name, and then choose Security credentials.
  - Check under Multi-factor Authentication (MFA) to see whether MFA is enabled or disabled. If MFA has not been activated, an alert symbol (Alert icon) is displayed.

- Reports

  The AWS Cost and Usage Reports (AWS CUR) contains the most comprehensive set of cost and usage data available. You can use Cost and Usage Reports to publish your AWS billing reports to an Amazon Simple Storage Service (Amazon S3) bucket that you own. You can receive reports that break down your costs by the hour, day, or month, by product or product resource, or by tags that you define yourself. AWS updates the report in your bucket once a day in comma-separated value (CSV) format. You can view the reports using spreadsheet software such as Microsoft Excel or Apache OpenOffice Calc, or access them from an application using the Amazon S3 API.

  AWS Cost and Usage Reports tracks your AWS usage and provides estimated charges associated with your account. Each report contains line items for each unique combination of AWS products, usage type, and operation that you use in your AWS account. You can customize the AWS Cost and Usage Reports to aggregate the information either by the hour, day, or month.

  AWS Cost and Usage Reports can do the following:
  - Deliver report files to your Amazon S3 bucket
  - Update the report up to three times a day
  - Create, retrieve, and delete your reports using the AWS CUR API Reference.

- Amazon EC2 Overview

  Amazon Elastic Compute Cloud (Amazon EC2) provides scalable computing capacity in the Amazon Web Services (AWS) Cloud. Using Amazon EC2 eliminates your need to invest in hardware up front, so you can develop and deploy applications faster. You can use Amazon EC2 to launch as many or as few virtual servers as you need, configure security and networking, and manage storage. Amazon EC2 enables you to scale up or down to handle changes in requirements or spikes in popularity, reducing your need to forecast traffic.

  Features of Amazon EC2

  Amazon EC2 provides the following features:
  - Virtual computing environments, known as instances
  - Preconfigured templates for your instances, known as Amazon Machine Images (AMIs), that package the bits you need for your server (including the operating system and additional software)
  - Various configurations of CPU, memory, storage, and networking capacity for your instances, known as instance types
  - Secure login information for your instances using key pairs (AWS stores the public key, and you store the private key in a secure place)
  - Storage volumes for temporary data that's deleted when you stop, hibernate, or terminate your instance, known as instance store volumes
  - Persistent storage volumes for your data using Amazon Elastic Block Store (Amazon EBS), known as Amazon EBS volumes
  - Multiple physical locations for your resources, such as instances and Amazon EBS volumes, known as Regions and Availability Zones

- A firewall that enables you to specify the protocols, ports, and source IP ranges that can reach your instances using security groups
- Static IPv4 addresses for dynamic cloud computing, known as Elastic IP addresses
- Metadata, known as tags that you can create and assign to your Amazon EC2 resources.
- Virtual networks you can create that are logically isolated from the rest of the AWS Cloud, and that you can optionally connect to your own network, known as virtual private clouds (VPCs)

- Elastic Block Storage (EBS)
  Amazon Elastic Block Store (Amazon EBS) provides block level storage volumes for use with EC2 instances. EBS volumes behave like raw, unformatted block devices. You can mount these volumes as devices on your instances. EBS volumes that are attached to an instance are exposed as storage volumes that persist independently from the life of the instance. You can create a file system on top of these volumes, or use them in any way you would use a block device (such as a hard drive). You can dynamically change the configuration of a volume attached to an instance.

  We recommend Amazon EBS for data that must be quickly accessible and requires long-term persistence. EBS volumes are particularly well-suited for use as the primary storage for file systems, databases, or for any applications that require fine granular updates and access to raw, unformatted, block-level storage. Amazon EBS is well suited to both database-style applications that rely on random reads and writes, and to throughput-intensive applications that perform long, continuous reads and writes.

- Amazon Machine Image (AMI)
  An Amazon Machine Image (AMI) is a supported and maintained image provided by AWS that provides the information required to launch an instance. You must specify an AMI when you launch an instance. You can launch multiple instances from a single AMI when you require multiple instances with the same configuration. You can use different AMIs to launch instances when you require instances with different configurations.
  An AMI includes the following:
  - One or more Amazon Elastic Block Store (Amazon EBS) snapshots, or, for instance-store-backed AMIs, a template for the root volume of the instance (for example, an operating system, an application server, and applications).
  - Launch permissions that control which AWS accounts can use the AMI to launch instances.
  - A block device mapping that specifies the volumes to attach to the instance when it's launched.

- Instance Purchasing Options & Introduction To EC2 Instance Types
  Amazon EC2 provides the following purchasing options to enable you to optimize your costs based on your needs:
  - On-Demand Instances – Pay, by the second, for the instances that you launch.
  - Savings Plans – Reduce your Amazon EC2 costs by making a commitment to a consistent amount of usage, in USD per hour, for a term of 1 or 3 years.
  - Reserved Instances – Reduce your Amazon EC2 costs by making a commitment to a consistent instance configuration, including instance type and Region, for a term of 1 or 3 years.
  - Spot Instances – Request unused EC2 instances, which can reduce your Amazon EC2 costs significantly.
  - Dedicated Hosts – Pay for a physical host that is fully dedicated to running your instances, and bring your existing per-socket, per-core, or per-VM software licenses to reduce costs.
  - Dedicated Instances – Pay, by the hour, for instances that run on single-tenant hardware.
  - Capacity Reservations – Reserve capacity for your EC2 instances in a specific Availability Zone for any duration.

  If you require a capacity reservation, purchase Reserved Instances or Capacity Reservations for a specific Availability Zone. Spot Instances are a cost-effective choice if you can be flexible about when your applications run and if they can be interrupted. Dedicated Hosts or Dedicated Instances can help you address compliance requirements and reduce costs by using your existing server-bound software licenses.

- Security Group Elastic
  Security groups are virtual firewalls that control traffic to our EC2 instances. We cannot block individual IP addresses using security groups and we cannot block an individual port. In SG, everything is blocked by default. We have to allow explicitly.

- Public And Private IP Overview
  A private IPv4 address is an IP address that's not reachable over the Internet. You can use private IPv4 addresses for communication between instances in the same VPC. For more information about the standards and specifications of private IPv4 addresses, see RFC 1918. We allocate private IPv4 addresses to instances using DHCP.

  A public IP address is an IPv4 address that's reachable from the Internet. You can use public addresses for communication between your instances and the Internet.
  When you launch an instance in a default VPC, we assign it a public IP address by default. When you launch an instance into a no default VPC, the subnet has an attribute that determines whether instances launched into that subnet receive a public IP address from the public IPv4 address pool. By default, we don't assign a public IP address to instances launched in a no default subnet.

- Amazon EBS Snapshot
  You can back up the data on your Amazon EBS volumes to Amazon S3 by taking point-in-time snapshots. Snapshots are incremental backups, which means that only the blocks on the device that have changed after your most recent snapshot are saved. This minimizes the time required to create the snapshot and saves on storage costs by not duplicating data. Each snapshot contains all of the information that is needed to restore your data (from the moment when the snapshot was taken) to a new EBS volume.
  When you create an EBS volume based on a snapshot, the new volume begins as an exact replica of the original volume that was used to create the snapshot. The replicated volume loads data in the background so that you can begin using it immediately. If you access data that hasn't been loaded yet, the volume immediately downloads the requested data from Amazon S3, and then continues loading the rest of the volume's data in the background.

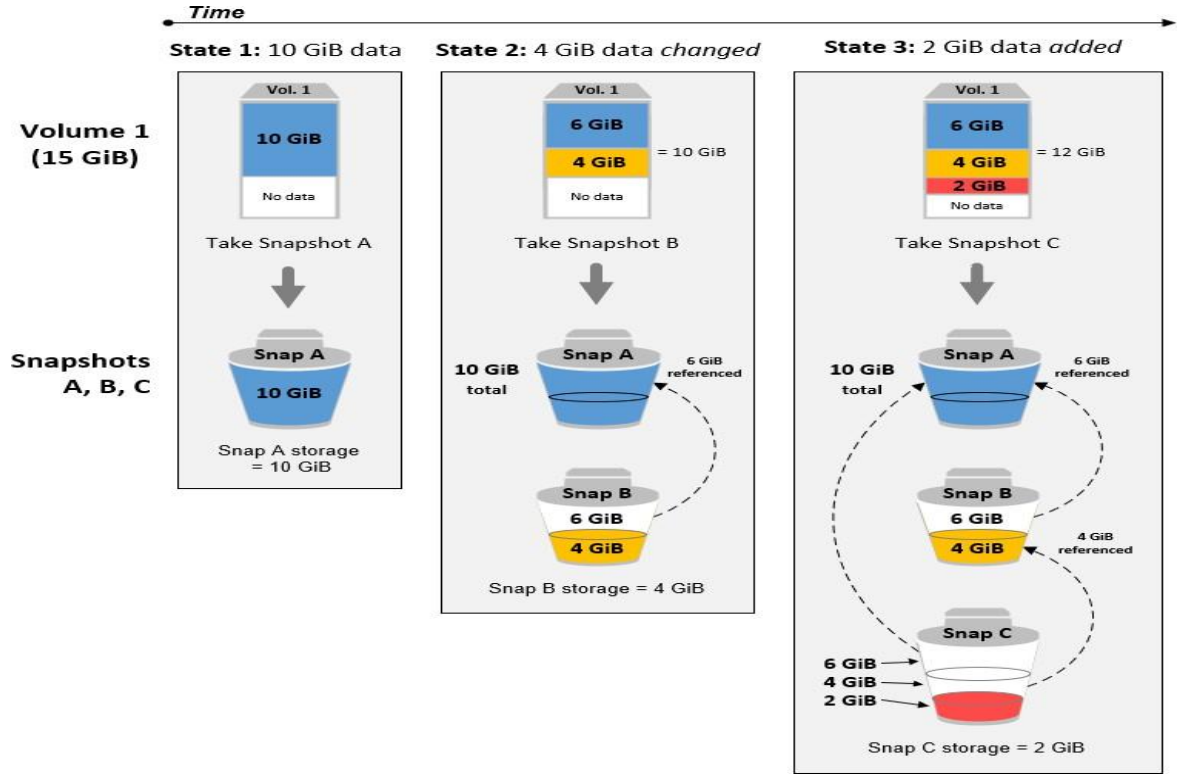

- AWS CLI
  The AWS Command Line Interface (AWS CLI) is a unified tool to manage your AWS services. With just one tool to download and configure, you can control multiple AWS services from the command line and automate them through scripts.
  The AWS CLI v2 offers several new features including improved installers, new configuration options such as AWS IAM Identity Center (successor to AWS SSO), and various interactive features.

- Bootstrap Script
  You can use a bootstrap action to install additional software or customize the configuration of cluster instances. Bootstrap actions are scripts that run on cluster after Amazon EMR launches the instance using the Amazon Linux Amazon Machine Image (AMI). Bootstrap actions run before Amazon EMR installs the applications that you specify when you create the cluster and before cluster nodes begin processing data. If you add nodes to a running cluster, bootstrap actions also run on those nodes in the same way. You can create custom bootstrap actions and specify them when you create your cluster.
  Most predefined bootstrap actions for Amazon EMR AMI versions 2.x and 3.x are not supported in Amazon EMR releases 4.x. For example, configure-Hadoop and configure-daemons are not supported in Amazon EMR release 4.x. Instead, Amazon EMR release 4.x natively provides this functionality. For more information about how to migrate bootstrap actions from Amazon EMR AMI versions 2.x and 3.x to Amazon EMR release 4.x, go to Customizing cluster and application configuration with earlier AMI versions of Amazon EMR in the Amazon EMR Release Guide.

- Elastic Load Balancing (EBS)
  In simplest terms, cloud computing means storing and accessing the data and programs on remote servers that are hosted on the internet instead of the computer's hard drive or local server. It is also referred to as Internet-based computing.
  Features of cloud
    - No up-front investment
    - Lowering operating cost
    - Highly scalable and efficient
    - Easy access
    - Reducing business risks and maintenance expenses

  Advantages of Elastic Load Balancer
    - ELB automatically distributes incoming application traffic across multiple targets, such as EC2 instances, containers, and IP addresses, to achieve high availability.
    - It can automatically scale to handle changes in traffic demand, allowing you to maintain consistent application performance.
    - It can monitor the health of its registered targets and route traffic only to the healthy targets.
    - It evenly distributes traffic across all availability zones in a region, improving fault tolerance.

- Auto Scaling
  Traditional IT environments are limited, using a specific number of servers to handle loads for any given application. When the amount of requests increases, so does the load on the server. Eventually, the demand on the load causes degraded performance and failure. Amazon Elastic Compute Cloud (EC2) provides an Auto Scaling service that overcomes this challenge.
  Auto Scaling makes sure there are enough EC2 instances to run applications. Before the service can run, you define auto scaling groups. For each group, you specify a minimal or maximum number of EC2 instances. Auto Scaling then detects if there is an error or failure on an instance, and immediately launches another instance to maintain the required capacity.
  Amazon EC2 also offers dynamic auto scaling policies, based on load metrics, Cloud Watch alarms, events from other Amazon services such as SQS, or a fixed schedule.

- Amazon Virtual Private Cloud (VPC)
  - VPC stands for Virtual Private Cloud.
  - Amazon Virtual Private Cloud (Amazon VPC) provides a logically isolated area of the AWS cloud where you can launch AWS resources in a virtual network that you define.
  - You have complete control over your virtual networking environment, including a selection of your IP address range, the creation of subnets, and configuration of route tables and network gateways.
  - You can easily customize the network configuration for your Amazon Virtual Private Cloud. For example, you can create a public-facing subnet for web servers that can access to the internet and can also place your backend system such as databases or application servers to a private-facing subnet.
  - You can provide multiple layers of security, including security groups and network access control lists, to help control access to Amazon EC2 instances in each subnet.

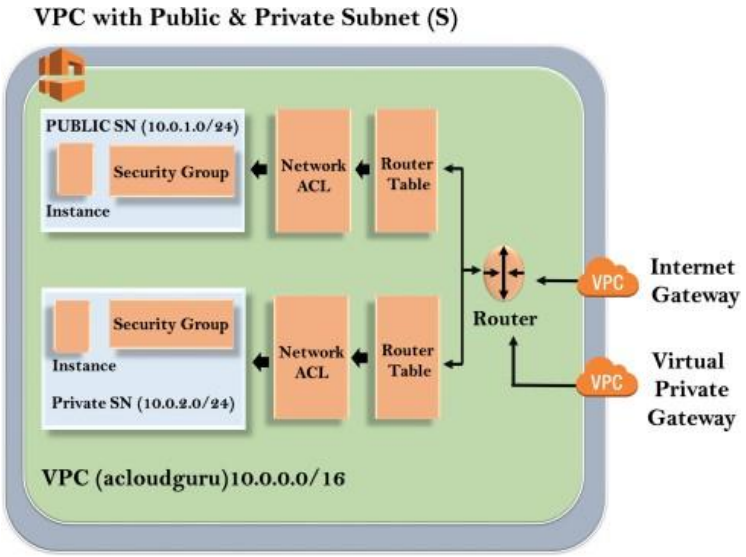

- Amazon VPC And Subnets

  The outer line represents the region, and the region is us-east-1. Inside the region, we have VPC, and outside the VPC, we have internet gateway and virtual private gateway. Internet Gateway and Virtual Private Gateway are the ways of connecting to the VPC. Both these connections go to the router in a VPC and then router directs the traffic to the route table. Route table will then direct the traffic to Network ACL. Network ACL is the firewall or much like security groups. Network ACL are stateliest which allows as well as deny the roles. You can also block the IP address on your Network ACL. Now, move over to the security group that accesses another line against the EC2 instance. It has two subnets, i.e., Public and Private Subnet. In public subnet, the internet is accessible by an EC2 instance, but in private subnet, an EC2 instance cannot access the internet on their own. We can connect the instances. To connect an instance, move over to the public subnet and then it SSH to the private subnet. This is known as jump boxes. In this way, we can connect an instance in public subnet to an instance in private subnet.
  Some ranges are reserved for private subnet:
  10.0.0.0 - 10.255.255.255 (10/8 prefix)
  172.16.0.0 - 172.31.255.255 (172.16/12 prefix)
  192.168.0.0 - 192.168.255.255 (192.108/16 prefix)

- Route Table

  AWS Route Table Concepts and How Route Tables Work. For instance, the following are the key concepts for route tables.

  Main route table—the route table that automatically comes with your VPC. It controls the routing for all subnets that are not explicitly associated with any other route table.

Custom route table—a route table that you create for your VPC.
Edge association – For example, you use to route inbound VPC traffic to an appliance. You associate a route table with the internet gateway or virtual private gateway, and specify the network interface of your appliance as the target for VPC traffic.
Route table association—the association between a route table and a subnet, internet gateway, or virtual private gateway.
Subnet route table—a route table that's associated with a subnet.
Gateway route table— this is, for instance, associated with an internet gateway or virtual private gateway.
Local gateway route table—it is mainly associated with an Outposts local gateway. For information about local gateways, see Local Gateways in the AWS Outposts User Guide.
Destination—the destination CIDR where you want traffic to go. For example, an external corporate network with a 172.16.0.0/12 CIDR.
Target—the target through which to send the destination traffic; for example, an internet gateway.
Local route—a default route for communication within the VPC

- Internet Gateway
    - Internet Gateway (IGW) is a horizontally scaled, redundant, and highly available VPC component that allows communication between your VPC and the internet.
    - Internet Gateway enables resources (like EC2 instances) in public subnets to connect to the internet. Similarly, resources on the internet can initiate a connection to resources in your subnet using the public.
    - If a VPC does not have an Internet Gateway, then the resources in the VPC cannot be accessed from the Internet (unless the traffic flows via a Corporate Network and VPN/Direct Connect).
    - Internet Gateway supports IPv4 and IPv6 traffic.
    - Internet Gateway does not cause availability risks or bandwidth constraints on your network traffic.
    - In order to make subnet public, add a route to your subnet's route table that directs internet-bound traffic to the internet gateway.
    - You can associate exactly one Internet Gateway with a VPC.
    - Internet Gateway is not Availability Zone specific.
    - There's no additional charge for having an internet gateway in your account.
- Simple Storage Services (S3)

  Amazon Simple Storage Service (Amazon S3) is an object storage service that offers industry-leading scalability, data availability, security, and performance. Customers of all sizes and industries can use Amazon S3 to store and protect any amount of data for a range of use cases, such as data lakes, websites, mobile applications, backup and restore, archive, enterprise applications, IoT devices, and big data analytics. Amazon S3 provides management features so that you can optimize, organize, and configure access to your data to meet your specific business, organizational, and compliance requirements.
  Features of Amazon S3

  Storage classes
  Amazon S3 offers a range of storage classes designed for different use cases. For example, you can store mission-critical production data in S3 Standard for frequent access, save costs by storing infrequently accessed data in S3 Standard-IA or S3 One Zone-IA, and archive data at the lowest costs in S3 Glacier Instant Retrieval, S3 Glacier Flexible Retrieval, and S3 Glacier Deep Archive.
  You can store data with changing or unknown access patterns in S3 Intelligent-Tiering, which optimizes storage costs by automatically moving your data between four access tiers when your access patterns change. These four access tiers include two low-latency access tiers optimized for frequent and infrequent access, and two opt-in archive access tiers designed for asynchronous access for rarely accessed data.

Storage management

Amazon S3 has storage management features that you can use to manage costs, meet regulatory requirements, reduce latency, and save multiple distinct copies of your data for compliance requirements.

S3 Lifecycle – Configure a lifecycle configuration to manage your objects and store them cost effectively throughout their lifecycle. You can transition objects to other S3 storage classes or expire objects that reach the end of their lifetimes.

S3 Object Lock – Prevent Amazon S3 objects from being deleted or overwritten for a fixed amount of time or indefinitely. You can use Object Lock to help meet regulatory requirements that require write-once-read-many (WORM) storage or to simply add another layer of protection against object changes and deletions.

S3 Replication – Replicate objects and their respective metadata and object tags to one or more destination buckets in the same or different AWS Regions for reduced latency, compliance, security, and other use cases.

S3 Batch Operations – Manage billions of objects at scale with a single S3 API request or a few clicks in the Amazon S3 console. You can use Batch Operations to perform operations such as Copy, Invoke AWS Lambda function, and Restore on millions or billions of objects.

Access management

Amazon S3 provides features for auditing and managing access to your buckets and objects. By default, S3 buckets and the objects in them are private. You have access only to the S3 resources that you create. To grant granular resource permissions that support your specific use case or to audit the permissions of your Amazon S3 resources, you can use the following features.

S3 Block Public Access – Block public access to S3 buckets and objects. By default, Block Public Access settings are turned on at the account and bucket level.

AWS Identity and Access Management (IAM) – IAM is a web service that helps you securely control access to AWS resources, including your Amazon S3 resources. With IAM, you can centrally manage permissions that control which AWS resources users can access. You use IAM to control who is authenticated (signed in) and authorized (has permissions) to use resources.

Bucket policies – Use IAM-based policy language to configure resource-based permissions for your S3 buckets and the objects in them.

Amazon S3 access points – Configure named network endpoints with dedicated access policies to manage data access at scale for shared datasets in Amazon S3.

Access control lists (ACLs) – Grant read and write permissions for individual buckets and objects to authorized users. As a general rule, we recommend using S3 resource-based policies (bucket policies and access point policies) or IAM policies for access control instead of ACLs. ACLs are an access control mechanism that predates resource-based policies and IAM. For more information about when you'd use ACLs instead of resource-based policies or IAM policies, see Access policy guidelines.

S3 Object Ownership – Disable ACLs and take ownership of every object in your bucket, simplifying access management for data stored in Amazon S3. You, as the bucket owner, automatically own and have full control over every object in your bucket, and access control for your data is based on policies.

Access Analyzer for S3 – Evaluate and monitor your S3 bucket access policies, ensuring that the policies provide only the intended access to your S3 resources.

Data processing

To transform data and trigger workflows to automate a variety of other processing activities at scale, you can use the following features.

S3 Object Lambda – Add your own code to S3 GET, HEAD, and LIST requests to modify and process data as it is returned to an application. Filter rows, dynamically resize images, redact confidential data, and much more.

Event notifications – Trigger workflows that use Amazon Simple Notification Service (Amazon SNS), Amazon Simple Queue Service (Amazon SQS), and AWS Lambda when a change is made to your S3 resources.

- S3 Object Storage and Buckets

  Object storage is a technology that stores and manages data in an unstructured format called objects. Modern organizations create and analyze large volumes of unstructured data such as photos, videos, email, web pages,

sensor data, and audio files. Cloud object storage systems distribute this data across multiple physical devices but allow users to access the content efficiently from a single, virtual storage repository. Object storage solutions are ideal for building cloud native applications that require scale and flexibility, and can also be used to import existing data stores for analytics, backup, or archive.

Metadata is critical to object storage technology. With object storage, objects are kept in a single bucket and are not files inside of folders. Instead, object storage combines the pieces of data that make up a file, adds all the user-created metadata to that file, and attaches a custom identifier. This creates a flat structure, called a bucket, as opposed to hierarchical or tiered storage. This lets you retrieve and analyze any object in the bucket, no matter the file type, based on its function and characteristics.

Object storage is the ideal storage for data lakes because it delivers an architecture for large amounts of data, with each piece of data stored as an object, and the object metadata provides a unique identifier for easier access. This architecture removes the scaling limitations of traditional storage, and is why object storage is the storage of the cloud.

The major benefits of object storage are the virtually unlimited scalability and the lower cost of storing large volumes of data for use cases such as data lakes, cloud native applications, analytics, log files, and machine learning (ML). Object storage also delivers greater data durability and resiliency because it stores objects on multiple devices, across multiple systems, and even across multiple data centers and regions. This allows for virtually unlimited scale and also improves resilience and availability of the data.

- Security On Buckets

  To make sure your files and Amazon S3 buckets are secure, follow these best practices:

  Restrict access to your S3 resources: When using AWS, restrict access to your resources to the people that absolutely need it. Follow the principle of least privilege.

  Monitor your S3 resources: Monitor your resources using AWS CloudTrail logs, S3 server access logging, AWS Config, AWS Identity and Access Management (IAM) Access Analyzer, Amazon Macie, Amazon CloudWatch, or AWS Trusted Advisor's S3 bucket permissions check.

  Use encryption to protect your data: Amazon S3 supports encryption during transmission, server-side encryption (SSE), and client-side encryption.

  Resolution

  Restrict access to your S3 resources

  By default, all S3 buckets are private and can be accessed only by users who are explicitly granted access.

  Restrict access to your S3 buckets or objects by doing the following:

  Writing IAM user policies that specify the users that can access specific buckets and objects. IAM policies provide a programmatic way to manage Amazon S3 permissions for multiple users. For more information about creating and testing user policies, see the AWS Policy Generator and IAM Policy Simulator.

  Writing bucket policies that define access to specific buckets and objects. You can use a bucket policy to grant access across AWS accounts, grant public or anonymous permissions, and allow or block access based on conditions. For more information about creating and testing bucket policies, see the AWS Policy Generator.

  Note: You can use a deny statement in a bucket policy to restrict access to specific IAM users. You can restrict access even if the users are granted access in an IAM policy.

  Using Amazon S3 Block Public Access as a centralized way to limit public access. Block Public Access settings override bucket policies and object permissions. Be sure to enable Block Public Access for all accounts and buckets that you don't want publicly accessible.

  Setting access control lists (ACLs) on your buckets and objects.

  Note: If you need a programmatic way to manage permissions, use IAM policies or bucket policies instead of ACLs. However, you can use ACLs when your bucket policy exceeds the 20 KB maximum file size. Or, you can use ACLs to grant access for Amazon S3 server access logs or Amazon CloudFront logs.

  Consider these best practices when you use ACLs to secure your resources:

  Be sure to review ACL permissions that allow Amazon S3 actions on a bucket or an object. For the list of ACL permissions and the actions that they allow, see What permissions can I grant?

Be stringent about who gets Read and Write access to your buckets.
Carefully consider your use case before granting Read access to the Everyone group because this allows anyone to access the bucket or object.
Never allow Write access to the Everyone group. This setting allows anyone to add objects to your bucket, which you will then be billed for. This setting also allows anyone to delete objects in the bucket.
Never allow Write access to the Any authenticated AWS user group. This group includes anyone with an active AWS account, not just IAM users in your account. To control access for IAM users on your account, use an IAM policy instead. For more information on how Amazon S3 evaluates IAM policies, see How Amazon S3 authorizes a request.

In addition to using policies, Block Public Access, and ACLs, you can also restrict access to specific actions in these ways:
Enable MFA delete, which requires a user to authenticate using a multi-factor authentication (MFA) device before deleting an object or disabling bucket versioning.
Set up MFA-protected API access, which requires that users authenticate with an AWS MFA device before they call certain Amazon S3 API operations.
If you temporarily share an S3 object with another user, create a presigned URL to grant time-limited access to the object. For more information, see Share an object with others.

- Web Hosting
  You can use Amazon S3 to host a static website. On a static website, individual webpages include static content. They might also contain client-side scripts.
  By contrast, a dynamic website relies on server-side processing, including server-side scripts, such as PHP, JSP, or ASP.NET. Amazon S3 does not support server-side scripting, but AWS has other resources for hosting dynamic websites. To learn more about website hosting on AWS

- Logging & Event
  To log data events for an S3 bucket to AWS CloudTrail and CloudWatch Events, create a trail. A trail captures API calls and related events in your account and delivers the log files to an S3 bucket that you specify. You can update an existing trail or create a new one.
  Open the CloudTrail console at https://console.aws.amazon.com/cloudtrail/.

  - In the navigation pane, choose Trails, Create trail.
  - For Trail name, type a name for the trail.
  - For Data events, type the bucket name and prefix (optional). For each trail, you can add up to 250 Amazon S3 objects.
  - To log data events for all Amazon S3 objects in a bucket, specify an S3 bucket and an empty prefix. When an event occurs on an object in that bucket, the trail processes and logs the event.
  - To log data events for specific Amazon S3 objects, choose Add S3 bucket, then specify an S3 bucket and optionally the object prefix. When an event occurs on an object in that bucket and the object starts with the specified prefix, the trail processes and logs the event.
  - For each resource, specify whether to log Read events, Write events, or both.
  - For Storage location, create or choose an existing S3 bucket to designate for log file storage.
  - Choose Create.

- Glacier
  The Amazon S3 Glacier storage classes are purpose-built for data archiving, providing you with the highest performance, most retrieval flexibility, and the lowest cost archive storage in the cloud. All S3 Glacier storage classes provide virtually unlimited scalability and are designed for 99.999999999% (11 nines) of data durability.

The S3 Glacier storage classes deliver options for the fastest access to your archive data and the lowest-cost archive storage in the cloud.

You can choose from three archive storage classes optimized for different access patterns and storage duration. For archive data that needs immediate access, such as medical images, news media assets, or genomics data, choose the S3 Glacier Instant Retrieval storage class, an archive storage class that delivers the lowest cost storage with milliseconds retrieval. For archive data that does not require immediate access but needs the flexibility to retrieve large sets of data at no cost, such as backup or disaster recovery use cases, choose S3 Glacier Flexible Retrieval (formerly S3 Glacier), with retrieval in minutes or free bulk retrievals in 5-12 hours. To save even more on long-lived archive storage such as compliance archives and digital media preservation, choose S3 Glacier Deep Archive, the lowest cost storage in the cloud with data retrieval within twelve hours.

- Versioning & Lifecycle Policy
  Amazon S3 now supports lifecycle rules for versioning. This means that you can now use lifecycle rules for S3 buckets regardless of whether they are enabled for versioning or not.
  Versioning provides protection against overwrites and deletes by enabling you to preserve, retrieve, and restore every version of every object in an Amazon S3 bucket. Lifecycle rules provide you the ability to configure rules that define what automatically happens to objects stored in your buckets after a specific date or period of time. Now that lifecycle rules are supported for versioning, you can combine versioning and lifecycle rules to enable scenarios such as easily setting up a rule that stores all your previous object versions in the lower cost Glacier storage class and deletes them from Glacier storage after 100 days. This example would provide a 100 day window to roll back any changes made to your data and automatically lower your storage costs.

- Cross Region Replication
  Replication enables automatic, asynchronous copying of objects across Amazon S3 buckets. Buckets that are configured for object replication can be owned by the same AWS account or by different accounts. You can replicate objects to a single destination bucket or to multiple destination buckets. The destination buckets can be in different AWS Regions or within the same Region as the source bucket.
  To automatically replicate new objects as they are written to the bucket, use live replication, such as Cross-Region Replication (CRR). To replicate existing objects to a different bucket on demand, use S3 Batch Replication. For more information about replicating existing objects, see When to use S3 Batch Replication.
  To enable CRR, you add a replication configuration to your source bucket. The minimum configuration must provide the following:
  The destination bucket or buckets where you want Amazon S3 to replicate objects
  An AWS Identity and Access Management (IAM) role that Amazon S3 can assume to replicate objects on your behalf
  Additional configuration options are available. For more information, see Additional replication configurations.
  To get detailed metrics for S3 Replication, including replication rule count metrics, you can use Amazon S3 Storage Lens. S3 Storage Lens is a cloud-storage analytics feature that you can use to gain organization-wide visibility into object-storage usage and activity. For more information, see Using S3 Storage Lens to protect your data. For a complete list of metrics.

- DNS Records
  Let's understand what is Amazon Route 53 in technical terms. AWS Route 53 lets developers and organizations route end users to their web applications in a very reliable and cost-effective manner. It is a Domain Name System (DNS) that translates domain names into IP addresses to direct traffic to your website. In simple terms, it converts World Wide Web addresses like www.example.com to IP addresses like 10.20.30.40.
  Basically, domain queries are automatically routed to the nearest DNS server to provide the quickest response possible. If you use a web hosting company like GoDaddy, it takes 30 minutes to 24 hours to remap a domain to a different IP, but by using Route 53 in AWS it takes only a few minutes.

- Website Hosting
  Suppose that you want to host a static website on Amazon S3. You've registered a domain with Amazon Route 53 (for example, example.com), and you want requests for http://www.example.com and http://example.com to be served from your Amazon S3 content. You can use this walkthrough to learn how to host a static website and create redirects on Amazon S3 for a website with a custom domain name that is registered with Route 53. You can work with an existing website that you want to host on Amazon S3, or use this walkthrough to start from scratch.
  Automating static website setup with an AWS CloudFormation template
  You can use an AWS CloudFormation template to automate your static website setup. The AWS CloudFormation template sets up the components that you need to host a secure static website so that you can focus more on your website's content and less on configuring components.
  The AWS CloudFormation template includes the following components:
  Amazon S3 – Creates an Amazon S3 bucket to host your static website.
  CloudFront – Creates a CloudFront distribution to speed up your static website.
  Lambda@Edge – Uses Lambda@Edge to add security headers to every server response. Security headers are a group of headers in the web server response that tell web browsers to take extra security precautions.

- Routing Policy
  When you create a record, you choose a routing policy, which determines how Amazon Route 53 responds to queries:
  - Simple routing policy – Use for a single resource that performs a given function for your domain, for example, a web server that serves content for the example.com website. You can use simple routing to create records in a private hosted zone.
  - Failover routing policy – Use when you want to configure active-passive failover. You can use failover routing to create records in a private hosted zone.
  - Geolocation routing policy – Use when you want to route traffic based on the location of your users. You can use geolocation routing to create records in a private hosted zone.
  - Geoproximity routing policy – Use when you want to route traffic based on the location of your resources and, optionally, shift traffic from resources in one location to resources in another.
  - Latency routing policy – Use when you have resources in multiple AWS Regions and you want to route traffic to the region that provides the best latency. You can use latency routing to create records in a private hosted zone.
  - IP-based routing policy – Use when you want to route traffic based on the location of your users, and have the IP addresses that the traffic originates from.
  - Multivalue answer routing policy – Use when you want Route 53 to respond to DNS queries with up to eight healthy records selected at random. You can use multivalue answer routing to create records in a private hosted zone.
  - Weighted routing policy – Use to route traffic to multiple resources in proportions that you specify. You can use weighted routing to create records in a private hosted zone.

- Health Check
  Amazon Route 53 health checks monitor the health and performance of your web applications, web servers, and other resources. Each health check that you create can monitor one of the following:
  The health of a specified resource, such as a web server.
  The status of other health checks.
  The status of an Amazon CloudWatch alarm.
  After you create a health check, you can get the status of the health check, get notifications when the status changes, and configure DNS failover:
  Getting health check status and notifications
  You can view the current and recent status of your health checks on the Route 53 console. You can also work with health checks programmatically through one of the AWS SDKs, the AWS Command Line Interface, AWS Tools for Windows PowerShell, or the Route 53 API.
  If you want to receive a notification when the status of a health check changes, you can configure an Amazon CloudWatch alarm for each health check.
  For information about viewing health check status and receiving notifications, see Monitoring health check status and getting notifications.
  Configuring DNS failover
  If you have multiple resources that perform the same function, you can configure DNS failover so that Route 53 will route your traffic from an unhealthy resource to a healthy resource. For example, if you have two web servers and one web server becomes unhealthy, Route 53 can route traffic to the other web server.

- Relation Database System
  RDS is a service on separate VPS, optimized for working with the databases.
  The following database management systems are available on Amazon RDS:

  - MySQL Community Edition
  - Oracle Database Standard Edition One
  - Oracle Database Standard Edition
  - Oracle Database Enterprise Edition

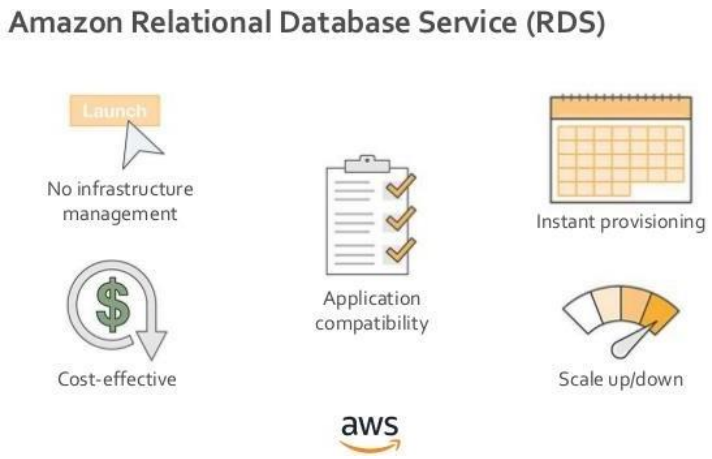

  RDS instance disk space can be ordered by the client. The minimum size of a storage is 5 GB.
  There is a possibility of flexible settings for server access. Only one address is available. This can be useful, for example, for autoscaling.
  You can also configure replication between servers.
  RDS supports instant blind (snapshot) and auto-backup, which allows you quickly and accurately recover a data.
  RDS will automatically transfer your host to a healthy node.

If there are any updates, DBMSs can be automatically patched and rebooted. Customers are notified in advance.
But there is no root access to the DBMS. The storage capabilities of the built-in procedures and fine-tuning are implemented through the API and command line utilities.
All RDS instances work on a 64 bit platform.

- DB Engine & Instance Details
  The DB instance class determines the computation and memory capacity of an Amazon RDS DB instance. The DB instance class that you need depends on your processing power and memory requirements.
  A DB instance class consists of both the DB instance type and the size. For example, db.m6g is a general-purpose DB instance type powered by AWS Graviton2 processors. Within the db.m6g instance type, db.m6g.2xlarge is a DB instance class.
  DB instance class types
  Amazon RDS supports three types of DB instance classes: general purpose, memory-optimized, and burstable performance. For more information about Amazon EC2 instance types, see Instance types in the Amazon EC2 documentation.

  The following are the general-purpose DB instance types available:
  - db.m6g – General-purpose DB instance classes powered by AWS Graviton2 processors. These instance classes deliver balanced compute, memory, and networking for a broad range of general-purpose workloads.
  - You can modify a DB instance to use one of the DB instance classes powered by AWS Graviton2 processors. To do so, complete the same steps as with any other DB instance modification.
  - db.m6gd – General-purpose DB instance classes powered by AWS Graviton2 processors. These instance classes deliver balanced compute, memory, and networking for a broad range of general-purpose workloads. They have local NVMe-based SSD block-level storage for applications that need high-speed, low latency local storage.
  - db.m6i – General-purpose DB instance classes that are well suited for a broad range of general-purpose workloads.
  - db.m5d – General-purpose DB instance classes that are optimized for low latency, very high random I/O performance, and high sequential read throughput.
  - db.m5 –General-purpose DB instance classes that provide a balance of compute, memory, and network resources, and are a good choice for many applications. The db.m5 instance classes provide more computing capacity than the previous db.m4 instance classes. They are powered by the AWS Nitro System, a combination of dedicated hardware and lightweight hypervisor.
  - db.m4 – General-purpose DB instance classes that provide more computing capacity than the previous db.m3 instance classes.

  For the RDS for Oracle DB engines, Amazon RDS has started the end-of-life process for db.m4 DB instance classes using the following schedule, which includes upgrade recommendations. For RDS for Oracle DB instances that use db.m4 instance classes, we recommend that you upgrade to a db.m5 instance class as soon as possible.

- Security
  Cloud security at AWS is the highest priority. As an AWS customer, you benefit from a data center and network architecture that is built to meet the requirements of the most security-sensitive organizations.
  Security is a shared responsibility between AWS and you. The shared responsibility model describes this as security of the cloud and security in the cloud:
  - Security of the cloud – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the AWS compliance programs. To

learn about the compliance programs that apply to Amazon Route 53, see AWS Services in Scope by Compliance Program.

- Security in the cloud – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

- Parameter Group

  The following list contains the parameters that all actions use for signing Signature Version 4 requests with a query string. Any action-specific parameters are listed in the topic for that action. For more information about Signature Version 4, see Signature Version 4 Signing Process in the Amazon Web Services General Reference.

  Action
  The action to be performed.
  Type: string
  Required: Yes
  Version
  The API version that the request is written for, expressed in the format YYYY-MM-DD.
  Type: string
  Required: Yes
  X-Amz-Algorithm
  The hash algorithm that you used to create the request signature.
  Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.
  Type: string
  Valid Values: AWS4-HMAC-SHA256
  Required: Conditional
  X-Amz-Credential
  The credential scope value, which is a string that includes your access key, the date, the region you are targeting, the service you are requesting, and a termination string ("aws4_request"). The value is expressed in the following format: access_key/YYYYMMDD/region/service/aws4_request.
  For more information, see Task 2: Create a String to Sign for Signature Version 4 in the Amazon Web Services General Reference.
  Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.
  Type: string
  Required: Conditional
  X-Amz-Date
  The date that is used to create the signature. The format must be ISO 8601 basic format (YYYYMMDD'T'HHMMSS'Z'). For example, the following date time is a valid X-Amz-Date value: 20120325T120000Z.
  Condition: X-Amz-Date is optional for all requests; it can be used to override the date used for signing requests. If the Date header is specified in the ISO 8601 basic format, X-Amz-Date is not required. When X-Amz-Date is used, it always overrides the value of the Date header. For more information, see Handling Dates in Signature Version 4 in the Amazon Web Services General Reference.
  Type: string
  Required: Conditional
  X-Amz-Security-Token
  The temporary security token that was obtained through a call to AWS Security Token Service (AWS STS). For a list of services that support temporary security credentials from AWS Security Token Service, go to AWS Services That Work with IAM in the IAM User Guide.

Condition: If you're using temporary security credentials from the AWS Security Token Service, you must include the security token.
Type: string
Required: Conditional
X-Amz-Signature
Specifies the hex-encoded signature that was calculated from the string to sign and the derived signing key.
Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.
Type: string
Required: Conditional
X-Amz-SignedHeaders
Specifies all the HTTP headers that were included as part of the canonical request. For more information about specifying signed headers, see Task 1: Create a Canonical Request For Signature Version 4 in the Amazon Web Services General Reference.
Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.
Type: string
Required: Conditional

- Monitoring Resourcing
  Monitoring is an important part of maintaining the reliability, availability, and performance of your AWS solutions. You should collect monitoring data from all of the parts of your AWS solution so that you can more easily debug a multi-point failure if one occurs. However, before you start monitoring, you should create a monitoring plan that includes answers to the following questions:

  - What are your monitoring goals?
  - What resources will you monitor?
  - How often will you monitor these resources?
  - What monitoring tools will you use?
  - Who will perform the monitoring tasks?
  - Who should be notified when something goes wrong?

- Dynamo DB
  - Amazon DynamoDB is a fast and flexible NoSQL database service for all applications that require consistent single-digit millisecond latency at any scale.
  - It is a fully managed database that supports both document and key-value data models.
  - Its flexible data model and performance makes it a great fit for mobile, web, gaming, ad-tech, IOT, and many other applications.
  - It is stored in SSD storage.
  - It is spread across three geographically data centres.

  Because of its availability in three geographically data centres, It consists of two different types of consistency models:

  Eventual Consistent Reads
  It maintains consistency across all the copies of data which is usually reached within a second. If you read a data from DynamoDB table, then the response would not reflect the most recently completed write operation, and if you repeat to read the data after a short period, then the response would be the lattest update. This is the best model for Read performance.

  Strongly Consistent Reads

A strongly consistent read returns a result that reflects all writes that received a successful response prior to the read.

- Elasticache

  Whether serving the latest news, a top-10 leaderboard, a product catalog, or selling tickets to an event, speed is the name of the game. The success of your website and business is greatly affected by the speed at which you deliver content.

  In "For Impatient Web Users, an Eye Blink Is Just Too Long to Wait," the New York Times noted that users can register a 250-millisecond (1/4 second) difference between competing sites. Users tend to opt out of the slower site in favor of the faster site. Tests done at Amazon, cited in How Webpage Load Time Is Related to Visitor Loss, revealed that for every 100-ms (1/10 second) increase in load time, sales decrease 1 percent.

  If someone wants data, you can deliver that data much faster if it's cached. That's true whether it's for a webpage or a report that drives business decisions. Can your business afford to not cache your webpages so as to deliver them with the shortest latency possible?

  It might seem intuitively obvious that you want to cache your most heavily requested items. But why not cache your less frequently requested items? Even the most optimized database query or remote API call is noticeably slower than retrieving a flat key from an in-memory cache. Noticeably slower tends to send customers elsewhere.

- Cloud Watch

  Amazon CloudWatch is basically a metrics repository. An AWS service—such as Amazon EC2—puts metrics into the repository, and you retrieve statistics based on those metrics. If you put your own custom metrics into the repository, you can retrieve statistics on these metrics as well.

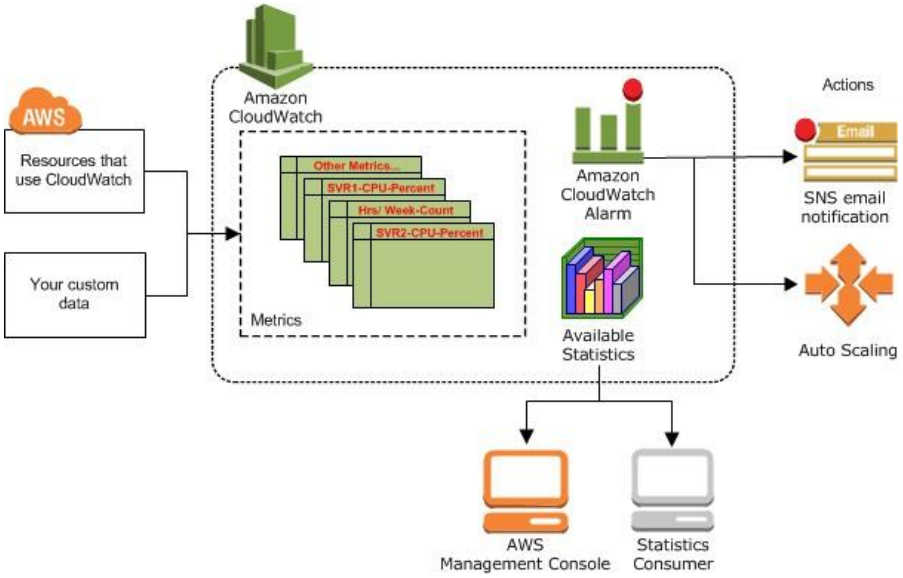

  You can use metrics to calculate statistics and then present the data graphically in the CloudWatch console. For more information about the other AWS resources that generate and send metrics to CloudWatch, see AWS services that publish CloudWatch metrics.

  You can configure alarm actions to stop, start, or terminate an Amazon EC2 instance when certain criteria are met. In addition, you can create alarms that initiate Amazon EC2 Auto Scaling and Amazon Simple Notification Service (Amazon SNS) actions on your behalf. For more information about creating CloudWatch alarms, see Alarms.

  AWS Cloud computing resources are housed in highly available data center facilities. To provide additional scalability and reliability, each data center facility is located in a specific geographical area, known as a Region. Each Region is designed to be completely isolated from the other Regions, to achieve the greatest possible failure isolation and stability. Metrics are stored separately in Regions, but you can use CloudWatch cross-Region functionality to aggregate statistics from different Regions. For more information, see Cross-account cross-Region CloudWatch console and Regions and Endpoints in the Amazon Web Services General Reference.

- Matrices

  The following matrix lists the supported deployment strategies for Amazon Elastic Container Service (Amazon ECS), AWS Lambda, and Amazon EC2/On-Premise.

  Amazon ECS is a fully managed orchestration service.

  AWS Lambda lets you run code without provisioning or managing servers.

  Amazon EC2 enables you to run secure, resizable compute capacity in the cloud.

- Alarm & Notification

  You can configure CloudWatch Logs to send a notification whenever an alarm is triggered for CloudTrail. Doing so enables you to respond quickly to critical operational events captured in CloudTrail events and detected by CloudWatch Logs. CloudWatch uses Amazon Simple Notification Service (SNS) to send email.

  Amazon CloudWatch uses Amazon SNS to send email. First, create and subscribe to an SNS topic. When you create a CloudWatch alarm, you can add this SNS topic to send an email notification when the alarm changes state.

  Alternatively, if you plan to create your CloudWatch alarm using the AWS Management Console, you can skip this procedure because you can create the topic when you create the alarm.

- Log & Billing Monitoring

  Monitoring is an important part of maintaining the reliability, availability, and performance of your AWS account. There are several tools available to monitor your Billing and Cost Management usage.

  AWS Cost and Usage Reports

  AWS Cost and Usage Reports tracks your AWS usage and provides estimated charges associated with your account. Each report contains line items for each unique combination of AWS products, usage type, and operation that you use in your AWS account. You can customize the AWS Cost and Usage Reports to aggregate the information either by the hour or by the day.

  AWS CloudTrail

  Billing and Cost Management is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in Billing and Cost Management. CloudTrail captures all write and modify API calls for Billing and Cost Management as events, including calls from the Billing and Cost Management console and from code calls to the Billing and Cost Management APIs.

- Other AWS Monitoring

  Monitoring is a critical part of any application. This article discusses X AWS monitoring tools that you can integrate with your AWS account.

  Cloud computing has taken over the IT world by storm. We have moved from the traditional system of in-house servers to virtual infrastructure in our cloud platforms. We use our cloud providers for all kinds of applications, and there are hardly any problem statements that we cannot solve on the cloud. The scalability, ease of use, and high availability make the cloud an obvious and an important choice for most modern applications.

  AWS is the leading cloud provider in the world. It has over 200 services and controls a massive chunk of the cloud market share. Thousands of startups and MNCs trust AWS as their cloud provider. With such a high demand for the AWS Cloud platform arises a need for the monitoring of our services in our account.

  AWS has a shared responsibility model. Some services in AWS are automatically managed, while others have to be managed by the user. Even if AWS was to do its part to the fullest, there could be scenarios where the hardware of your service malfunctions or your application crashes. All these scenarios and possible crashes should have a monitoring system.

  Monitoring your AWS resources is one of the best ways to ensure that your resource (or your application) is efficiently performing to its capacity. Monitoring is a major topic in most of the AWS certification exams.

  AWS understands the importance of monitoring and why it can be critical for applications to have an efficient monitoring system. It is not only meant for application or system failures. It can also help you make important changes to your application architecture based on your monitoring insights.

- Eucalyptus

  Eucalyptus in cloud computing is an open-source software platform for carrying out IaaS or Infrastructure-as-a-Service in a hybrid cloud computing or private cloud computing environment.

  Eucalyptus in cloud computing pools together existing virtualised framework to make cloud resources for storage as a service, network as a service and infrastructure as a service. Elastic Utility Computing Architecture for Linking Your Programs To Useful Systems is short known as Eucalyptus in cloud computing.

  Eucalyptus in cloud computing frameworks declared a conventional concurrence with AWS or Amazon Web Services in March 2012, permitting overseers to move cases between an Amazon Elastic Compute Cloud and the Eucalyptus private cloud to make a hybrid cloud. The organisation additionally permits Eucalyptus to work with Amazon's product groups to create interesting Amazon Web Services viable highlights.

  It tends to be effortlessly sent in existing IT frameworks to appreciate the advantages of both eucalyptus private cloud and eucalyptus public cloud models.

  - History
  - Eucalyptus architecture
  - Eucalyptus components

- Other tools
- The advantages of the Eucalyptus cloud
- Eucalyptus vs other IaaS private clouds
- What is the use of eucalyptus in cloud computing?

1. History

Improvement on Eucalyptus started as an examination project at US-based Rice University in the year 2003. In the year 2009, an organisation named Eucalyptus Systems was framed to market Eucalyptus software. Afterwards, in the year 2012, the firm went into a concurrence with Amazon Web Services for keeping up similarity and Application Programming Interface support. In the year 2014, it was procured by Hewlett-Packard or HP, which unexpectedly has its own cloud contributions under the HPE eucalyptus. The Helion portfolio has an assortment of cloud-related items, which incorporates HP's own kind of OpenStack called HP Helion OpenStack. Presently, Eucalyptus is a piece of the HPE portfolio and is known as HPE Helion Eucalyptus.

2. Eucalyptus architecture

Eucalyptus CLIs can oversee both Amazon Web Services and their own private occasions. Clients can undoubtedly relocate cases from Eucalyptus to Amazon Elastic Cloud. Network, storage, and compute are overseen by the virtualisation layer. Occurrences are isolated by hardware virtualisation. The following wording is utilised by Eucalyptus architecture in cloud computing.

1. Images: Any software application, configuration, module software or framework software packaged and conveyed in the Eucalyptus cloud is known as a Eucalyptus Machine Image.
2. Instances: When we run the picture and utilise it, it turns into an instance.
3. Networking: The Eucalyptus network is partitioned into three modes: Static mode, System mode, and Managed mode.
4. Access control: It is utilised to give limitation to clients.
5. Eucalyptus elastic block storage: It gives block-level storage volumes to connect to an instance.
6. Auto-scaling and load adjusting: It is utilised to make or obliterate cases or administrations dependent on necessities.

3. Eucalyptus components

Components of eucalyptus in cloud computing:

1. Cluster Controller: It oversees at least one Node controller and liable for sending and overseeing occurrences on them.
2. Storage Controller: It permits the making of depictions of volumes.
3. Cloud Controller: It is a front end for the whole environment.
4. Walrus Storage Controller: It is a straightforward record storage framework.
5. Node Controller: It is an essential part of Nodes. It keeps up the life cycle of the occasions running on every node.

4. Other tools

Numerous other tools can be utilised to associate with AWS and Eucalyptus in cloud computing, and they are recorded below.

1. Vagrant AWS Plugin: This instrument gives config records to oversee AWS instances and oversee VMs on the local framework.
2. s3curl: This is a device for collaboration between AWS S3 and Eucalyptus Walrus.
3. s3fs: This is a FUSE record framework, which can be utilised to mount a bucket from Walrus or S3 as a local document framework.
4. Cloudberry S3 Explorer: This Windows instrument is for overseeing documents among S3 and Walrus.

5. The advantages of the Eucalyptus cloud

The benefits of Eucalyptus in cloud computing are:

- Eucalyptus can be utilised to benefit both the eucalyptus private cloud and the eucalyptus public cloud.
- Clients can run Amazon or Eucalyptus machine pictures as examples on both clouds.
- It isn't extremely mainstream on the lookout yet is a solid contender to CloudStack and OpenStack.
- It has 100% Application Programming Interface similarity with all the Amazon Web Services.
- Eucalyptus can be utilised with DevOps apparatuses like Chef and Puppet.

Features of eucalyptus in cloud computing are:

- Supports both Windows and Linux virtual machines.
- API is viable with the Amazon EC2 platform.
- Viable with Simple Storage Service (S3) and Amazon Web Services (AWS).

6. Eucalyptus vs other IaaS private clouds

There are numerous Infrastructure-as-a-Service contributions accessible in the market like OpenNebula, Eucalyptus, CloudStack and OpenStack, all being utilised as private and public Infrastructure-as-a-Service contributions.

Of the multitude of Infrastructure-as-a-Service contributions, OpenStack stays the most well-known, dynamic and greatest open-source cloud computing project. At this point, eagerness for OpenNebula, CloudStack and Eucalyptus stay strong.

7. What is the use of eucalyptus in cloud computing?

It is utilised to assemble hybrid, public and private cloud. It can likewise deliver your own datacentre into a private cloud and permit you to stretch out the usefulness to numerous different organisations.

Conclusion

Eucalyptus in cloud computing is open-source programming that carries out an AWS viable cloud, which is financially savvy, secure and flexible. It tends to be effectively sent in existing IT frameworks to appreciate both private and public cloud models' advantages.

- Microsoft Azure

Microsoft Azure, formerly known as Windows Azure, is Microsoft's public cloud computing platform. It provides a broad range of cloud services, including compute, analytics, storage and networking. Users can pick and choose from these services to develop and scale new applications or run existing applications in the public cloud.

The Azure platform aims to help businesses manage challenges and meet their organizational goals. It offers tools that support all industries -- including e-commerce, finance and a variety of Fortune 500 companies -- and is compatible with open source technologies. This gives users the flexibility to use their preferred tools and technologies. In addition, Azure offers four different forms of cloud computing: infrastructure as a service (IaaS), platform as a service (PaaS), software as a service (SaaS) and serverless functions.

Microsoft charges for Azure on a pay-as-you-go (PAYG) basis, meaning subscribers receive a bill each month that only charges them for the specific resources and services they have used.

How does Microsoft Azure work?

Once customers subscribe to Azure, they have access to all the services included in the Azure portal. Subscribers can use these services to create cloud-based resources, such as VMs and databases. Azure resources and services can then be assembled into running environments used to host workloads and store data.

In addition to the services that Microsoft offers through the Azure portal, a number of third-party vendors also make software directly available through Azure. The cost billed for third-party applications varies widely but may involve paying a subscription fee for the application, plus a usage fee for the infrastructure used to host the application.

Microsoft provides the following five different customer support options for Azure:

Basic

Developer

Standard
Professional Direct
Enterprise (Premier)
These customer support plans vary in terms of scope and price. Basic support is available to all Azure accounts, but Microsoft charges a fee for the other support offerings. Developer support costs $29 per month, while Standard support costs $100 per month and Professional Direct support is $1,000 per month. Microsoft does not disclose the pricing for Enterprise support.
What is Microsoft Azure used for?
Because Microsoft Azure consists of widely varied resource and service offerings, its use cases are extremely diverse. Running virtual machines or containers in the cloud is one of the most popular uses for Microsoft Azure. These compute resources can host infrastructure components, such as domain name system (DNS) servers; Windows Server services, such as Internet Information Services (IIS); networking services such as firewalls; or third-party applications. Microsoft also supports the use of third-party operating systems, such as Linux.

Azure is also commonly used as a platform for hosting databases in the cloud. Microsoft offers serverless relational databases such as Azure SQL and non-relational databases such as NoSQL.

In addition, the platform is frequently used for backup and disaster recovery. Many organizations use Azure for archival storage in order to meet their long-term data retention or disaster recovery (DR) requirements.

- Amazon EC2
  Amazon EC2 is short for: Amazon Elastic Compute Cloud.
  Amazon EC2 provides cloud hosted virtual machines, called "instances", to run applications.
  Amazon EC2 is a cloud computing platform that can be auto-scaled to to meet demand.
  Different hardware and software configurations can be selected. Different geographical locations can be selected be closer to users, as well as providing redundancy in case of failures.
  Persistent storage can be provided by Amazon EBS (Elastic Block Storage). Amazon S3 (Simple Storage Service) data can also be accessed with Amazon EC2 instances, and is free if they are in the same region.