

CLOUD COMPUTING

Unit III — Cloud Platforms and Applications

Topics Covered (09 Hours)

AWS Infrastructure: Regions, Availability Zones, & Edge
AWS Components: EC2 (Compute Instances, AMIs, Groups)
Amazon SimpleDB, S3 Bucket & Elastic Block Store (EBS)
Amazon Relational Database Service (RDS) & DynamoDB
Microsoft Azure Core Concepts: Groups and Subscriptions
SQL Azure Databases & Managed App Runtime Services
Open Source Platforms: OpenStack Compute Architecture
CloudStack & Eucalyptus Private Cloud Operations
Emerging Cloud Applications: IoT Integrations in Smart Cities
Centralized Cloud Dashboards (Traffic, Waste, Power Grids)
Artificial Intelligence and Machine Learning in the Cloud

*Compiled in a simple, easy-to-understand, book style format
for students of Computer Science and Engineering*

TABLE OF CONTENTS

3.1 Industrial Cloud Platforms: Amazon Web Services	3
AWS Global Infrastructure	3
Core AWS Compute & Storage Components	3
3.2 Microsoft Azure & Open Source Platforms	4
Microsoft Azure Core Concepts	4
Open Source Private Cloud Platforms	4
3.3 Cloud Applications: IoT, Smart Cities & AI	5
Smart Cities & IoT Sensor Integration	5
AI and Machine Learning in the Cloud	5
Unit Summary	6

3.1 Industrial Cloud Platforms: Amazon Web Services

Amazon Web Services (AWS) is the world's most comprehensive and broadly adopted cloud platform. It provides computing, storage, database, networking, and analytics services across a global infrastructure.

AWS Global Infrastructure

- **Regions:** Physical geographic locations around the world where AWS clusters data centers. Each Region is completely isolated from other Regions to ensure high fault tolerance and compliance.
- **Availability Zones (AZs):** One or more discrete data centers within an AWS Region. Each AZ has redundant power, cooling, and network connectivity, connected to other AZs in the same Region via low-latency fiber-optic links. This allows users to design highly available architectures by replicating applications across multiple AZs.
- **Edge Locations:** Cache locations deployed globally to deliver content (such as videos, data, and APIs) to end users with low latency using CloudFront (CDN).

Core AWS Compute & Storage Components

- **Elastic Compute Cloud (EC2):** Provides secure, resizable compute capacity (Virtual Machines) in the cloud. Users select CPU, memory, and storage configurations, booting instances using pre-configured templates called Amazon Machine Images (AMIs). Security Groups act as virtual firewalls to control traffic.
- **Amazon Simple Storage Service (S3):** An object storage service offering industry-leading scalability, data availability, security, and performance. Data is stored in containers called buckets, providing 99.999999999% (11 9s) of durability.
- **Elastic Block Store (EBS):** Provides persistent block-level storage volumes for use with EC2 instances (resembling raw physical hard drives).
- **Database Services (SimpleDB, RDS, DynamoDB):** AWS offers Amazon SimpleDB (a legacy, lightweight NoSQL database), RDS (Relational Database Service supporting MySQL, PostgreSQL, Oracle, etc., automating patching and backups), and DynamoDB (a fully managed, high-performance NoSQL database offering single-digit millisecond latency).

3.2 Microsoft Azure & Open Source Platforms

Microsoft Azure is the second largest public cloud platform, heavily integrated with enterprise Windows software, while open-source platforms enable building private clouds.

Microsoft Azure Core Concepts

- **Resource Groups:** Logical containers that hold related Azure resources (VMs, databases, virtual networks). Resources are managed, deployed, and billed collectively.
- **SQL Azure:** A fully managed relational cloud database service built on Microsoft SQL Server, providing automated scaling, high availability, and threat detection.
- **Application Services:** Azure App Service is a PaaS platform for hosting web applications, REST APIs, and mobile backends, supporting managed runtimes like .NET, Java, Node.js, and Python.

Open Source Private Cloud Platforms

- **OpenStack:** A powerful open-source cloud operating system that controls large pools of compute, storage, and networking resources throughout a data center. Key components include: Nova (Compute orchestration), Neutron (Software-Defined Networking), Swift (Object storage), Cinder (Block storage), Keystone (Identity management), and Horizon (Web Dashboard).
- **CloudStack:** An open-source turn-key PaaS/IaaS platform used to deploy and manage large networks of virtual machines, supporting multiple hypervisors (KVM, Xen, VMware).
- **Eucalyptus:** An open-source private cloud software that is fully API-compatible with AWS. This allows hybrid deployments where applications can be moved between Eucalyptus private clouds and AWS public clouds without code changes.

3.3 Cloud Applications: IoT, Smart Cities & AI

Cloud platforms enable modern data-intensive applications by aggregating telemetry from millions of edge devices and scaling heavy AI/ML calculations.

Smart Cities & IoT Sensor Integration

A Smart City utilizes internet-connected sensors (IoT) to gather real-time data across urban infrastructures. The cloud acts as the centralized data aggregator:

- **Traffic Management:** IoT cameras and road sensors stream active traffic volumes to the cloud. Machine learning models analyze the data and dynamically adjust traffic light timings to reduce congestion.
- **Waste Management:** Trash bins equipped with ultrasonic level sensors report their capacity to the cloud, allowing waste services to dynamically route trash collectors only to full bins, saving fuel.
- **Power Grids:** Smart meters report electricity usage every second. The cloud analyzes this big data to predict load spikes and balance supply across the city.

All these data streams are integrated into a single, centralized Cloud Dashboard, providing administrators with real-time operational views.

AI and Machine Learning in the Cloud

AI/ML model training requires massive computational power (GPUs and TPUs) that is expensive to own. Cloud platforms provide fully managed ML environments (like AWS SageMaker, Azure ML, and Google Vertex AI). Developers can scale compute clusters instantly, train deep neural networks on big datasets stored in S3 or Blob storage, and deploy models as scalable web APIs with a single click.

Unit Summary

This unit explored industrial and open-source cloud platforms and applications.

- AWS global infrastructure uses Regions, Availability Zones, and Edge content caches.
- AWS core compute is EC2, supported by SimpleDB, S3, EBS, and RDS databases.
- Microsoft Azure provides managed App Services, Resource Groups, and SQL Azure.
- Open-source private clouds are built using OpenStack, CloudStack, and Eucalyptus.
- Smart Cities integrate IoT traffic, waste, and power grid sensor data into central cloud dashboards.
- Managed cloud environments scale GPUs/TPUs to train and host complex AI/ML models.

© Copyright — abhyasmitra.in — All Rights Reserved