



More than what you think.

AWS 101



Advanced
Consulting
Partner



Index

Chapter 1

What can AWS offer?

Chapter 2

Getting Started - AWS Cloud platform

Chapter 3

AWS Analytics

Chapter 4

AWS services in Management and Governance

Chapter 5

AWS Databases

Chapter 6

AWS Storage

Conclusion

Appendix

Resources

Introduction

What is Cloud Computing?

Simply put, cloud computing is a collective use of software and hardware which delivers services over a network - the internet, on demand. The delivery of computing services—including servers, storage, databases, networking, software, analytics, and intelligence happens through exchange of data and information on a cloud to users who have access to it, anywhere, anytime. These services are easily accessible, increase efficiency, and makes tasks easy to users.

The other biggest advantage of deploying a cloud based technology over on premises or server based, computing assistance is its pay-as-you-go pricing. A cloud services platform provides immediate access to flexible and low-cost IT resources, which does not require large investments in terms of hardware and time. The cloud offers a mix and match facility where the enterprise can choose the best computing resources for all its IT needs. It is a simple, efficient way to access servers, storage, databases and a broad set of application services over the Internet and pay only for the amount of resources used .

A cloud services platform, is a service provider that owns, facilitates and maintains the network connected hardware required for these cloud based services. While there are a few players that offer these services, AWS or Amazon Web Services, a platform that which has its beginnings a a firm that offered IT Infrastructural services to enterprises through the web in 2006, has emerged a pioneer.





The Amazon Web Services, with over 50% of the cloud based enterprise market share, provides an accessible, highly reliable, adaptable, cost effective infrastructure platform in the cloud that services businesses in over 170 countries around the world.

With over 140 AWS services available, AWS offers a wide spectrum of cloud-based products from data warehousing to deployment tools, directories to content delivery, data analytics to data security. Further, AWS allows for provisioning of new services quickly, without the upfront hardware or other capital costs. This allows for companies, organisations and enterprises who have had prior exposure to IT or have never had any contact with the world of IT, to easily configure, build and access technology they need to keep in sync with the ever changing business requirements.

In partnership with Amazon, YASH combines its legacy and domain expertise in hosting/migration solutions, to deliver efficient Cloud-hosted environments. It leverages Amazon's Cloud platform solutions that enables applications to be highly available, secure, and scalable.

This e-book provides the reader with a comprehensive overview of the AWS Cloud and introduces the reader to the basic concepts of AWS.

Chapter 1

What can AWS offer?

The beauty of AWS is that, anything that could run on physical hardware or servers can be run on it seamlessly. The solutions provided can be grouped together to offer complete business solutions or meet piece meal requirements.

AWS offers as many services as the clients demand, and it is constantly adding more to its catalogue. They can be roughly grouped into the following categories based on the functionality of services.

Administrative Services:

Managing resources - Identify & Access, Monitor, Deploy & automate

Application Services:

Adding functionalities like search, messaging, workflow etc, to the application

Foundation services:

Basic services like storage, computation, networking etc.,

There are more than one services being offered by AWS in each category. For example, AWS offers five database services, each one optimized for a certain type of use. The idea is complete freedom to the enterprise to choose what they want for themselves, resulting in a best fit business solution.

There is no one size fits all and each enterprise no matter how big or small can get its own custom made application.

YASH is a pioneer that is always upto date in state of the art technologies, investing in AWS training, certifications, and resources to ensure smooth, quality service before, during, and after migration.

Chapter 2

Getting started: Amazon Web Services Cloud Platform

The Amazon Web Services offers many cloud services that can be used in imaginative permutations and combined to achieve desired business results. This chapter elucidates major services AWS has on offer, outlined by category.

All these services can be accessed majorly through Software Development Kits, or the Command Line Interface, or the AWS Management Console, which are briefly

Software Development Kits

which are essentially tailored kits which simplify the usage of AWS services in the application through program interfaces compatible with the application platform.

Command Line Interface

The AWS Command Line Interface is a one stop shop - a single central control interface, through which multiple AWS services can be accessed and controlled through scripts.

Management Console

This is a simple user interface which allows management and access of the AWS. It is also available in a mobile version, which offers quicker access to services.

Infrastructure in AWS:

The AWS Cloud infrastructure is built on a robust model that is classified into Regions and Availability zones. An AWS Region is a physical location which has several availability zones, which consist of one or more data centers. Each Data Center has independent power, connectivity and other hardware facilities. As of now, the AWS has over 20 AWS regions and 60 availability zones, which offer facilities and services that are less susceptible to faults, highly flexible and more available.

Also, each AWS Region is strategically located such that it is isolated from other AWS Regions. Every availability zone in the region is independent too, but connected through low latency links.

Security in AWS:

AWS promises that security - data or otherwise is its highest priority. Modelled after some of the most security sensitive organisations, security tools ensure adequate monitoring and protection of information, resources and data.

The AWS Cloud offers a shared responsibility model, which means that AWS manages the security OF the cloud, while the enterprise is in itself responsible for security WITHIN the cloud. The enterprise is free to choose the level of security it requires, and deploys it accordingly.

However, AWS provides complete advisories and guidance on all security concerns and issues.

Advantages of AWS in security:

Compliance: AWS comes with inbuilt compliance programs in its infrastructure that it constantly upgrades. This means that many crucial statutory compliances are already attended to.

Cost effective: By choosing the level of security, data encryption and accessibility, the enterprise has the comfort of building and innovation while still maintaining a secure environment that AWS offers.

Scalability: Irrespective of the size of business, the AWS infrastructure is designed to keep all data safe.

Compliance in AWS :

With a combination of water tight data governance practices and concurrent audit service features, AWS offers unmatched compliance

Some of the assurance programs that AWS complies with are:

- SOC 1/ISAE 3402, SOC 2, SOC 3
- FISMA, DIACAP, and FedRAMP
- PCI DSS Level 1

Chapter 3

AWS Analytics

Amazon Athena

Amazon Athena is a serverless, interactive query service that uses standard SQL commands to analyse data from Amazon S3. Athena is easy to use, retrieves data within seconds, and does not require complex jobs for analysis of data. When integrated with other services like AWS Glue data Catalog, it can transform into a data repository, where you can play around with data, pulling out informatics and spread sheets to suit requirements.

Amazon EMR

Amazon EMR is a Hadoop framework that is simple, fast and low cost, allowing to manage huge amounts of data across dynamic E2C instances. It also allows for interaction in other data stores, and provides a base for querying and analysis.

Amazon EMR handles a large loads of big data use cases, including log analysis, web indexing, data transformations (ETL), machine learning, financial analysis, scientific simulation, and bioinformatics.

Amazon CloudSearch

Amazon CloudSearch is a managed service that allows to set up, manage a search solution for the application or website. It supports 34 languages and popular search features such as highlighting, autocomplete, and geospatial search.



Amazon Kinesis

Amazon Kinesis is a service that offers collection, processing and analysis of real time data that is streamed into the application from various sources.

With Amazon Kinesis, real-time data such as video, audio, application logs, website click streams, are gathered and analysed as and when it arrives, and also enable to generate a suitable response or reaction.

Amazon Kinesis currently offers four services: Kinesis Data Firehose, Kinesis Data Analytics, Kinesis Data Streams, and Kinesis Video Streams which operate in different capabilities.

Amazon QuickSight

Amazon QuickSight Amazon QuickSight is a business intelligence (BI) service that delivers insights to everyone in the organisation. It allows creation of interactive dashboards that can be embedded into applications, and can be accessed from mobiles or website providing valuable analytical insights.

AWS Glue

AWS Glue is a service that helps retrieve, transform and load data for analytics as required. AWS Glue, when pointed to data, discovers it and stores the associated metadata in its Data catalog, after which the data can be queried, searched or analysed.

Amazon Redshift

Amazon Redshift is a data warehouse that analyzes data from across several data warehouses and data lakes, through machine learning, parallel queries, and still pay as low as 0.25\$ per hour. Redshift allows for queries on petabytes of data, delivering results in record times.

AWS Data Pipeline

AWS Data Pipeline is a service that helps processing and movement of data between different AWS compute and storage services, as well as on-premises data sources, at specified intervals. AWS Data Pipeline also allows you to move and process data that was previously locked up in on-premises data silos.

AWS Lake Formation

AWS Lake Formation is a service that helps set up a secure data lake in days. A data lake is a centralized repository that stores all data in various formats, in its original form and masked form. However, while managing a data lake can be a task in itself, creating a data lake with Lake Formation is easy - data definitions, access and security policies required when clearly defined - Lake formation collects and catalogs data into Amazon S3 data lake, cleans and classifies it using ML, and also secures access to sensitive data.

Chapter 4

AWS services in Management and Governance

Amazon CloudWatch

Amazon CloudWatch is a service where for developers, system operators, and IT managers can monitor and manage the application. It provides data and insights based on set parameters that help optimise resource utilisation, identify potential risk areas, and get an overview of operational robustness.

AWS Auto Scaling

AWS Auto Scaling allows application scaling for multiple resources across multiple services with high speed. It provides a simple, powerful user interface that helps build scaling plans for resources including Amazon EC2 instances, Spot Fleets, and Amazon Aurora Replicas. AWS Auto Scaling makes scaling simple with recommendations that allow optimizing performance, costs, or balance between them.

AWS Control Tower

AWS Control Tower just like it sounds, allows the set up of a landing zone. The configuration of the landing zone is architected on best practices that have been carefully established through time from thousands of enterprise workers, with time tested rules for security, operations and compliance.

Consider an enterprise that wants to migrate to AWS, with distributed teams and large number of applications. It would be easy for the enterprise to create multiple accounts while still monitoring them centrally. Control Tower automates the set-up of their landing zone where distributed teams are able to provision new AWS accounts quickly, while central teams have the comfort that the teams are secure and adherent to policies, with no compromise on speed or agility.

AWS CloudFormation

AWS CloudFormation helps developers and systems administrators to create and manage a collection of related AWS resources, cataloging them in an orderly manner.

It also has sample templates or allows creation of own templates to describe AWS resources, and associated dependencies required to run the application. It also enables visualization of templates as diagrams and enables editing them using a drag-and-drop interface with the AWS CloudFormation Designer.

AWS Trusted Advisor

AWS Trusted Advisor is an online resource to provide guidance and advisory information about the AWS environment. It also provides real-time guidance to help provision resources following AWS best practices.

AWS Personal Health Dashboard

AWS Personal Health Dashboard provides alerts and reminders when AWS senses that there are events that might effect the application. While the Service Health Dashboard displays the general status of AWS services, Personal Health Dashboard gives a personalized view of the performance and availability of AWS services with respect to the underlying AWS resources.

Chapter 5

AWS Databases

Amazon Aurora

Amazon Aurora is relational database engine, compatible with MySQL and PostgreSQL. It offers the speed and availability of high-end commercial databases with the simplicity and cost-effectiveness of open source databases. Amazon Aurora is five times faster than standard databases while providing the reliability of commercial databases at 1/10th of the cost.

It features a fault-tolerant, self-healing storage system that auto-scales up to 64TB per database instance, while delivering high performance and availability and continuous backup to Amazon S3.

Amazon RDS

Amazon Relational Database Service (Amazon RDS) is a tool that makes it easy to set up, operate, and scale a relational database in the cloud.

Amazon RDS provides six database engines to choose from, including Amazon Aurora, PostgreSQL, MySQL, MariaDB, Oracle Database, and SQL Server.

Amazon RDS on VMware

Amazon Relational Database Service (RDS) on VMware offers deployment of managed databases in on-premises VMware environments using the Amazon RDS technology. Amazon RDS automates administrative tasks including hardware provisioning, database set-up, patching, and backups.



Amazon DynamoDB

Amazon DynamoDB is a key-value and document database that delivers single-digit millisecond performance at any scale. which can handle more than 10 trillion requests per day and support peaks of more than 20 million requests per second.

Amazon Neptune

Amazon Neptune is a graph database service that helps build and run applications that work with highly connected datasets. It supports popular graph models like the Property Graph and W3C's RDF, and their respective query languages Apache TinkerPop Gremlin and SPARQL, allowing for easy building of queries that efficiently navigate highly connected datasets.

Amazon Timestream

Amazon Timestream is a time series database service for IoT and operational applications that makes it easy to store and analyze trillions of events per day at 1/10th the cost of relational databases.

Amazon DocumentDB

Amazon DocumentDB is a document database service that supports MongoDB workloads, designed from the ground-up to deliver performance, scalability, and availability at scale.

Amazon ElastiCache

Amazon ElastiCache is a web service that allows to deploy, operate, and scale an in-memory cache in the cloud. The service improves the performance of web applications by allowing retrieval of information from fast, managed, in-memory caches, instead of relying on disk-based databases.

Amazon Quantum Ledger Database (QLDB)

Amazon QLDB is a ledger database that provides a block chain like transparent, immutable, and cryptographically verifiable transaction log owned by a central trusted authority. Reliability and immutability in relational databases is time-consuming and prone to human error. Alternatively, blockchain frameworks, such as Hyperledger Fabric and Ethereum, can also be used as a ledger involves a great deal of complexity.

Amazon QLDB is a new class of database that eliminates the need to engage in the complex development effort of building its own ledger-like applications. With QLDB, the data's change history is immutable – it cannot be altered or deleted – and using cryptography, it can be easily verified that there have been no unintended modifications to your ap-

Chapter 5

Storage

Amazon S3

Amazon Simple Storage Service (Amazon S3) is an object storage service which can store and protect any amount of data. It can manage data from a range of use cases, such as websites, mobile applications, backup and restore, archive, enterprise applications, IoT devices, and big data analytics.

Amazon Elastic Block Store

Amazon Elastic Block Store (Amazon EBS) is a block storage volume provision for use with Amazon EC2 instances in the AWS Cloud. Each volume is automatically replicated within its Availability Zone to protect the application from component failure while offering availability and durability.

Amazon Elastic File System

Amazon Elastic File System (Amazon EFS) is an elastic file system for Linux-based workloads for use with both AWS Cloud services and on-premises resources. It is built to scale on demand to several petabytes without disrupting applications, growing and shrinking automatically. Amazon EFS is a fully managed service that requires no changes to any existing applications and tools, providing access through a standard file system interface for seamless integration.

Amazon FSx for Lustre

Amazon FSx for Lustre is a file system, optimized for compute-intensive workloads. It is ideal for applications that require high-performance and low latencies of scale-out, parallel file systems which require specialized expertise and administrative overhead, and complex performance parameters.

Amazon FSx for Lustre is integrated with Amazon S3, making it easy to link long-term data sets with high performance file systems to run compute-intensive workloads.

Amazon FSx for Windows

Amazon FSx for Windows File Server provides a native Microsoft Windows file system to move Windows-based applications that require file storage to AWS.

Amazon S3 Glacier

Amazon S3 Glacier is a storage service for data archiving and long-term backup. It provides comprehensive security and compliance capabilities costs as low as \$0.004 per gigabyte per month.

AWS Storage Gateway

AWS Storage Gateway is a hybrid storage service that enables on-premises applications to seamlessly use AWS cloud storage, which can be used for backup and archiving, disaster recovery, cloud data processing, storage tiering, and migration.

Conclusion:

AWS is like an exciting set of Lego blocks - which can be quickly assembled and used to create dream applications. AWS through its exhaustive services, provides access to highly durable storage, low-cost compute, high-performance databases, management tools, and much more. All this is provided at low costs that the enterprise only pays as much as it uses for. It is a trusted platform for millions of users worldwide and enjoys a lions share in cloud computing services.

AWS is a platform that is constantly evolving with new services, features and provisions added every day to suit the latest trends in technology and cater to various business requirements.

As a consulting partner and channel reseller of Amazon Web Services, YASH Technologies helps migrate enterprise applications, data, mobile applications, and everything imaginable, to a secure and reliable Cloud infrastructure.

Resources

<https://aws.amazon.com>



Copyright © 2019. YASH Technologies. All Rights Reserved.