

Moving to Microservice Architecture: A case Study

TABLE OF CONTENTS

1 About Our Client	1
2 The Problem Statement	1
3 ACC Solution	2
4 Amazon EC2	2
5 Amazon ECS	3
6 Amazon RDS	4
7 Amazon S3	7
8 Load Balancer	8
9 VPC	9

About our Client: -

Our client is the leading stock exchange of India. It is the world's largest derivatives exchange by number of contracts traded based on the statistics maintained by Futures Industry Association (FIA), a derivatives trade body. Our client was the first exchange in the country to provide a modern, fully automated screen-based electronic trading system that offered easy trading facilities to investors spread across the length and breadth of the country.



Problem Statement:-

The client followed a monolithic approach where the overhead of installing and configuring applications. It took a large amount of time and effort and to overcome this they moved from monolithic architecture to microservices architecture. The basis of microservices architecture is a container. A container is a standard unit of software that packages up code and all its dependencies so the application runs quickly and reliably from one computing environment to another.

A container image is a lightweight, standalone, executable package of software that includes everything needed to run an application: code, runtime, system tools, system libraries and settings. Container images become containers at runtime.

The client settled on using cloud vendors, container platforms and Platform as a Service (PaaS) to have their own built-in container engine that uses OCI-compatible container images.

ACC successfully helped the client deploy a microservices architecture application that implemented the mentioned base stack technologies.

Solution:-

- EC2
- ECS
- Elastic Cache
- Elastic Search
- S3
- Load Balancer
- VPC

**Amazon EC2:-**

The ecosystem of Amazon Web Services includes Elastic Compute Cloud (EC2). EC2 offers scalable, on-demand computing capacity in the AWS cloud. With Amazon EC2 instances, there is no longer a need to pay for leasing hardware up front or keep it maintained. You can develop and launch applications more quickly thanks to it. With AWS' EC2, you can start as many virtual servers as you require. You can scale up or down based on the volume of website visitors. The term "elastic" used to describe Elastic Compute Cloud refers to the capability of the system to adapt to shifting workloads and provision or de-provision resources in response to demand.

Using Amazon EC2, you can access the following features:

- Virtual computing environments are instances.
- Amazon Machine Images (AMIs) are already-configured templates for your instances that contain the server components you require (including the operating system and additional software)

Amazon Elastic Container Service: -

Users can rent virtual computers on which to execute their own computer programmes through Amazon Elastic Compute Cloud (EC2), a component of Amazon.com's cloud computing infrastructure, Amazon Web Services (AWS). By offering a web service through which a user can boot an Amazon Machine Image (AMI) to configure a virtual machine, or what Amazon refers to as a "instance," containing any required software, EC2 promotes the scalable deployment of applications. The phrase "elastic" refers to the ability of a user to construct, launch, and terminate server instances as needed while paying per second for active servers. By giving users control over the physical location of instances, EC2 enables users to reduce latency and increase redundancy. Amazon migrated to EC2 and AWS as their retail website platform in November 2010.

Features: -

- Using your choice continuous integration and delivery (CI/CD) and automation technologies, launch thousands of containers across the cloud.
- By removing the need to configure and manage the control plane, nodes, and instances, AWS Fargate serverless compute for containers can help you make the most of your time.
- Autonomous provisioning, auto-scaling, and pay-as-you-go pricing can reduce compute expenses by up to 50%.
- Easily integrate with AWS management and governance tools, which are standardised for compliance with almost every regulatory body worldwide.

Amazon RDS: -

The managed SQL database service known as Amazon Relational Database Service (RDS) is offered by Amazon Web Services (AWS). To store and manage data, Amazon RDS offers a variety of database engines. Additionally, it supports activities related to relational database maintenance, including data migration, backup, recovery, and patching.

The setup and upkeep of relational databases in the cloud are made easier by Amazon RDS. A relational instance of a cloud database is set up, run, managed, and scaled by a cloud administrator using Amazon RDS. Amazon RDS is a service used to manage relational databases; it is not a database in and of itself

Features of RDS: -

- Less administrative requirements
- **Simple to use**
- You can quickly gain access to the features of a production-ready relational database via the Amazon Web Services Management Console, the Amazon RDS Command Line Interface, or straightforward API calls.
- **Scalability**
- Push-button compute scaling
- You can scale the compute and memory resources powering your deployment up or down, up to a maximum of 32 vCPUs and 244 GiB of RAM. Compute scaling operations typically complete in a few minutes.
- **Easy storage scaling**
- As your storage requirements grow, you can also provision additional storage. The Amazon Aurora engine will automatically grow the size of your database volume as your database storage needs grow, up to a maximum of 64 TB or a maximum you define. The MySQL, MariaDB, Oracle, and PostgreSQL engines allow you to scale up to 64 TB of storage and SQL Server supports up to 16 TB. Storage scaling is on-the-fly with zero downtime.

Accessibility and Robustness

Scripted backups

Point-in-time recovery for your database instance is made possible via Amazon RDS' automated backup capability. Your database and transaction logs will be backed up via Amazon RDS and stored for a user-specified retention period. This enables you to restore your database instance to any second up to the previous five minutes of your retention period. You have the option of configuring your automatic backup retention time to last up to 35 days.

Snapshots of databases

Database snapshots are user-initiated copies of your instance that are retained in Amazon S3 until they are specifically deleted. Anytime you choose, you can start a fresh instance from a database snapshot. You are only charged for incremental storage utilisation even if database snapshots function functionally as full backups.

Amazon ElasticCache

A fully managed in-memory caching solution that supports numerous real-time use cases is called Amazon ElastiCache. ElastiCache can be used as a primary data store for use cases including session stores, gaming leaderboards, streaming, and analytics or for caching, which improves application and database performance. Redis and Memcached are compatible with ElastiCache.

Amazon Elasticsearch

Built on Apache Lucene, Elasticsearch is a distributed search and analytics engine. Elasticsearch, which was introduced in 2010 and has since grown to be the most widely used search engine, is frequently used for use cases involving log analytics, full-text search, security intelligence, business analytics, and operational intelligence.

Benefits: -

- A quick time to value
- Elasticsearch makes it simple to get started and quickly construct applications for a range of use-cases by providing simple REST-based APIs, a simple HTTP interface and uses schema-free JSON documents.
- Extreme performance
- Elasticsearch can swiftly locate the best matches for your queries thanks to the distributed nature that allows it to handle huge volumes of data in parallel.
- Tools and plugins that are free to use
- Kibana, a well-liked reporting and visualisation tool, is integrated with Elasticsearch. You can simply transform source data and load it into your Elasticsearch cluster thanks to its connection with Beats and Logstash. In order to provide your app with comprehensive functionality, you may also leverage a variety of free Elasticsearch plugins, such as language analyzers and suggesters.

Amazon S3: -

A web service interface-based object storage service offered by Amazon Web Services is called Amazon S3, also known as Amazon Simple Storage Service (AWS). Amazon S3 makes use of the same scalable storage technology that Amazon.com utilises to power their e-commerce network. In addition to hybrid cloud storage, backups, disaster recovery, data archiving, and data lakes for analytics, Amazon S3 can store any type of item, making it useful for a wide range of purposes.

The object storage architecture that Amazon S3 uses to manage data aims to be scalable, highly available, low latency, and highly durable. Objects are grouped into buckets as the basic storage units in Amazon S3. Each object is identified by a unique key that the user has chosen. You can use the Amazon console.

Benefits of S3: -**Dependable security**

Only the identity that originally generated the Amazon S3 buckets can use them (IAM policy grants are the exception). You have total control over how, where, and who can access the data on a regular basis. For each file, each bucket, or using IAM, you can specify access rights (Identity access management). By employing this set of rules and permissions, you can make sure that no one else has access to your data.

Availability

Every Amazon S3 user has access to the same highly scalable, dependable, efficient, and cost-effective data storage infrastructure that Amazon uses to power its own global network of websites. S3 Standard and Standard - IA are designed to be accessible 99.9% of the time and 99.99 percent of the time, respectively. Both are supported by Amazon's strict adherence to the SLA for its Amazon S3 service.

Load Balancer:-

Elastic load balancing is used to automatically divide incoming traffic among various targets, such as EC2 instances, containers, and IP addresses in one or more Availability Zones. It monitors the condition of the registered targets, sending traffic only to those who are healthy. Elastic load balancing scales your load balancer as your incoming traffic changes over time. It has the ability to scale most workloads automatically.

A load balancer serves as clients' only point of contact that you specify. The rules that you define for a listener determine how the load balancer sends requests to its registered destinations. Each rule consists of a priority, one or more actions, and one or more conditions. When a rule's conditions are met, the rule's actions are taken.

Amazon VPC: -

You can launch Amazon resources into a defined virtual network using the Amazon Virtual Private Cloud (Amazon VPC). With the advantages of utilising Amazon's scalable infrastructure, this virtual network closely mimics a conventional network that you would operate in your own data centre.

Each VPC builds an own virtual network environment for your AWS account in the cloud. In order to offer cloud services, other AWS resources and services run inside of VPC networks.

Anyone used to managing a physical Data Center will recognise AWS VPC (DC). A VPC functions similarly to a conventional TCP/IP network that may be grown and expanded as necessary. However, a VPC does not explicitly contain any of the DC parts that you are accustomed to working with, such as routers, switches, VLANs, etc. They were redesigned and abstracted into cloud software.

A virtual network architecture that AWS instances can be launched into can be quickly created using VPC. Each VPC specifies the requirements for your AWS resources, such as:

- Subnetworks
- Routing
- Security
- Networking

ABOUT ACC

ACC is an AWS Advance Partner with AWS Mobility Competency. Awarded The Best BFSI industry Consulting Partner for the year 2019, ACC has had several successful cloud migration and application development projects to its credit.

Our business offerings include Digitalisation, Cloud Services, Product Engineering, Big Data & Analytics and Cloud Security. ACC has developed several products to its credit. These include Ottohm – Enterprise Video and OTT Platform, Atlas API – API Management and Development Platform, Atlas CLM – Cloud Life Cycle Management, Atlas HCM – HR Digital Onboarding and Employee Management, Atlas ITSM – Vendor Onboarding and Service Management and Smart Contracts – Contract Automation and Management.



www.appliedcloudcomputing.com

Shubho Pramanik

(+91) 90297 20294

shubho@acc.ltd

Rogin Rappai

(+91) 8828478321

rogin.rappai@acc.ltd

also reach us at

