Capgemini

MAINFRAME MODERNIZATION & MIGRATION TO AWS CASE STUDY - LARGE U.S. INVESTMENT COMPANY

About the Customer

A Large U.S. Investment Company

The customer is one of the world's largest investment companies, offering a large selection of low-cost mutual funds, ETFs, advice, and related services. They provide expertise, stability, and reliable investment approach to individual investor, financial professional, or corporate or institutional investor. Services provided include brokerage services, variable and fixed annuities, financial planning, asset management, and trust services.

Challenges Faced by the Customer

The customer wanted to modernize their systems for one of their main lines of business by migrating data and refactoring business functionality to the AWS Cloud. The customer had a heavy dependency on mainframe systems across all their IT products with more than 40 million online/batch lines of code, and more than 5,000 IBM Data Management (IBM DB2) tables. They wanted to reduce mainframe MIPS (“Millions of Instructions Per Second”) and reduce dependencies on mainframe experts several of whom were retiring/nearing retirement.

The customer had developed a Java middleware over the span of 20 years which was hard to maintain and upgrade. The Java middleware was built to expose mainframe services to be consumed by the front-end applications. Previously, the customer had developed several Private Cloud Foundry (PCF) services, but since they had made the decision to standardize all their workloads on AWS, they decided to retire PCF services and move everything to Amazon ECS based microservices.

To break up their monolithic applications (mainframe/Java), the customer had identified strategic bounded contexts based on expert knowledge. The customer then was seeking help to migrate the bounded contexts to AWS in an accelerated manner using the new ways of working to bring increased agility and speed. Capgemini was selected to be this partner.

Capgemini's Specific Solutions to Help the Customer Solve Their Challenges

Capgemini partnered with the customer to identify bounded contexts/user journeys that were prioritized for migration to AWS. For the identified bounded contexts, Capgemini performed reverse engineering of Java and Mainframe COBOL code to identify functional and mainframe data dependencies for a successful migration.

With the goal to breakdown the monolithic applications, the customer faced two main challenges. The first was to maintain a single source of truth when moving the data consumption (read) onto AWS while keeping the IBM DB2 as the gold copy (write). The second was to set up a real-time pipeline to move the SQL databases to NoSQL (Amazon DynamoDB) databases. Capgemini proposed solutions for both.

Capgemini implemented its homegrown CAP360 solution, which is Capgemini's mainframe code analysis tool to help identify technical debt and assist in impact analysis during the reverse engineering phase. To migrate data to AWS and set up real-time change data capture pipelines between IBM DB2 and Amazon DynamoDB, Capgemini used a solution based on AWS Kinesis, AWS Lambda, and AWS Glue. Finally, APIs were designed using REST best practices and were implemented using docker containers on ECS. The APIs were exposed using Apigee which allowed self-discovery by the consuming applications and front-end.

Outcomes of the Solution Benefited by the Customer

The development of microservices on AWS granted the customer's business stakeholders the flexibility and agility to stitch microservices together to provide the business capabilities desired. This led to a reduction in IBM DB2 footprint and Mainframe MIPS (i.e. leading to cost reductions as well as risk reductions due to heavy dependencies on scarce mainframe experts) as functionalities and data were moved to cloud-native microservices on AWS.

An API-led design allowed for the self-discovery of the microservices that can now be consumed by multiple application/front-end teams. The customer was also able to retire their PCF-based services and standardize everything on Amazon ECS for their APIs.

In order to accelerate migration across the enterprise, Capgemini helped the customer setup data replication-as-a-service and helped speed up application development and eliminate the need for dedicated data resources on application development projects. The expertise of the data pipeline was complex, and the customer valued creating a central group that can be leveraged by multiple teams internally.

About Capgemini

Capgemini is a global leader in partnering with companies to transform and manage their business by harnessing the power of technology. The Group is guided everyday by its purpose of unleashing human energy through technology for an inclusive and sustainable future. It is a responsible and diverse organization of 290,000 team members in nearly 50 countries. With its strong 50 year heritage and deep industry expertise, Capgemini is trusted by its clients to address the entire breadth of their business needs, from strategy and design to operations, fueled by the fast evolving and innovative world of cloud, data, AI, connectivity, software, digital engineering and platforms. The Group reported in 2020 global revenues of €16 billion.

Get the Future You Want | www.capgemini.com