Reducing operational complexity and cost with serverless computing

Table of contents

Introduction

To transform application development into a true competitive advantage, organizations need to enable their developers to experiment frequently and iterate often. Achieving this requires implementing development practices that allow teams to focus on building differentiated applications, rather than handling undifferentiated infrastructure management tasks. One method of doing so is the implementation of a serverless architecture. This eBook will explore the advantages of running applications on a serverless architecture, detail how others have been successful, and explain how you can easily get started by using native cloud services on Amazon Web Services (AWS).

The challenges of managing monolithic architectures

As organizations look for ways to build and maintain a competitive edge, many should look no further than modernizing their application architectures. That's because many businesses are still running monolithic architectures and doing so creates responsibilities that add friction to failing fast. This includes the need to:

- **Architect application environments for reliability and scalability**
- **Patch and maintain server fleets**
- **Elastically scale server fleets in a manner that doesn't create unutilized capacity or hurt the end-user experience**

All this equates to a significant amount of heavy lifting that diverts developers away from building and delivering exceptional applications. Moreover, this heavy lifting massively increases as developers are faced with architecting environments to support more and more applications and large-scale data processing systems, while managing their backends. Organizations can offload many of these undifferentiated tasks by implementing a serverless architecture on the cloud.

What is serverless computing, and why is it valuable?

In a serverless architecture, you leverage a variety of fully managed services that take care of various application concerns. This includes data persistence, API management, messaging and application integration, as well as the core compute platform that runs your code. By using fully managed services instead of hosting application components on virtual machines (VMs) or containers that you run yourself, backend infrastructure management tasks are completely abstracted away from your team. This means your engineers no longer have to worry about provisioning servers or clusters, patching, operating system maintenance, capacity provisioning, and more. By stripping infrastructure management tasks away from the developers, businesses free up their teams to focus on applications and thus foster innovation. Moreover, this also makes it easier to get applications to market faster.

Another key advantage of serverless architectures is that they are typically much less expensive to run than VM- or container-based implementations. Most servers spend the majority of their cycles waiting for work to do, but you pay for that capacity whether it's being used or not. Because the managed services used in serverless architecture are priced based on actual usage, you no longer waste money due to underutilization.

Constructing a serverless architecture with AWS

As organizations look to implement their own serverless architecture, many choose AWS to do so. In large part, that's because AWS provides a comprehensive suite of fully managed services that can help you architect a complete serverless architecture on the cloud. These services are updated continuously based on customer needs and feedback to give users the latest technology innovations. Additionally, AWS Partner Network (APN) Partners provide tools and resources that simplify your adoption of serverless.

Compute

Run code and containers on AWS or at the edge without having to provision or manage servers.

Storage

Use object and elastic file storage services to simply, securely, and elastically store data at any scale.

Data Stores

Support your applications with NoSQL or relational database services that start up, shut down, and elastically scale capacity based on your app's needs.

API Proxy

Create, publish, maintain, monitor, and secure APIs at any scale.

Interprocess Messaging

Decouple and scale microservices, distributed systems, and serverless applications.

Orchestration

Easily coordinate the components of distributed applications and microservices using visual workflows.

Analytics

Collect, process, and analyze streaming data, or analyze data residing in Amazon S3 without managing infrastructure.

Developer Tooling

Leverage tools and services from AWS and its partner ecosystem for support in developing serverless applications.

Benefits of running a serverless architecture

Running applications on a serverless architecture on AWS has helped many businesses modernize their development practices and foster innovation. Additional key benefits include:

No servers to manage:

In a serverless architecture, there is no need to provision, deploy, update, monitor, or otherwise manage servers – this is all handled by AWS. As a result, your developers can focus on building and delivering applications.

Continuous scaling:

Your applications scale automatically, mirroring their actual usage, allowing you to meet demand without hurting end-user experiences or impacting the integrity of internal systems.

Cost-efficiency:

Because applications scale with their actual usage, you also eliminate the need to pre-provision or over-provision capacity, thus reducing costs. Additionally, you only pay per request, based on the duration it takes for your code to run.

Customer success story: FINRA

About the customer

The Financial Industry Regulatory Authority (FINRA) is a not-for-profit organization that regulates broker-dealers to protect investors and ensure market integrity. One of FINRA's primary workloads is its Order Audit Trail System (OATS), which is used along with other market data to monitor trading practices of member firms.

The challenge

OATS processes over 50,000 files from broker-dealers daily. Once the files are received, the data in them is validated for completeness and formatting according to a set of more than 200 rules. This adds up to half a trillion validations each day. Processing demand varies significantly over time and can double or triple in response to market conditions that drive higher trading volumes.

Given the sheer volume OATS handles daily, FINRA sought a new architecture that would accelerate data processing. They also were looking for greater scalability, data partitioning, robust monitoring, high performance, lower cost and maintenance requirements, and strong security. Based on a number of proof-of-concepts, FINRA found an architecture powered by AWS Lambda best fit their needs.

The AWS solution

FINRA's new AWS Lambda architecture was developed in only three months, including testing to make sure it could handle peak loads. Key services used in this architecture include Amazon S3 for data ingestion, Amazon EC2 to run the controller, Amazon SQS for input/output notifications, and Amazon Virtual Private Cloud (Amazon VPC) to partition the system into separate test and production accounts. Since implementation, FINRA has seen a number of benefits including reduced costs and operational simplicity. Tim Griesbach, Senior Director of Technology at FINRA reports, "using AWS Lambda, we've increased cost efficiency by a factor of two. We only pay for what we use, and we don't have to manage on-premises server infrastructure."

Furthermore, the organization has accelerated processing times, helping them meet their goal of real-time processing: "With our previous system, validation results were returned to broker-dealers within three minutes on average, and up to seven minutes at peak times. With serverless processing based on AWS Lambda, we are returning results in less than one minute regardless of volume." Now, FINRA is looking to move more of the application to AWS Lambda and to migrate additional applications to AWS.

Getting started with our APN consulting partners

About Cloudreach

As a native Cloud integrator, our Cloud Enablement and Cloud Operations team at Cloudreach, work with our clients to help them deploy the right tooling, re-design business processes, show the 'art of the possible,' and create a bespoke Cloud Adoption Program to structure the transformation. Learn how CloudReach can help you build modern applications with increased agility and lower total cost of ownership

Learn More >>

Contact CloudReach>>

Learn more

Ready to learn more? Visit our website to explore additional resources to help you implement a serverless architecture with AWS and APN Partners.

Visit our website to learn more >>