

Case Study



















Multi-Cloud Hosted Data Lake Platform

Client Overview

A California-based hosted data lake provider offers custom pipelines for Apache Pulsar and Trino, helping enterprises adopt AI. They needed a Managed Service Provider (MSP) to build a cloud-hosted platform that automates deployment and management on AWS, GCP, or Azure, based on their end clients' infrastructure. They also required self-service provisioning, automated scaling, and marketplace integration (AWS & GCP) with ongoing support.

Challenges

- Building a scalable, fault-tolerant infrastructure for distributed messaging and SQL queries.
- Seamless integration with AWS, GCP, and Azure while ensuring compliance and listing on AWS and GCP marketplace.
- Automating provisioning, billing, and monitoring for a usage-based model.

Infrastructure Tech Stack	
CI/CD	  
Database	
Cloud Service Provider	 
Version Control System	 
Infrastructure Provisioning & Automation	
Orchestration & Management	
Monitoring Tools	  
3rd Party tools / Others	    

Our Contribution

- **Multi-Cloud Deployment Support** – Designed a platform that allows deployment on AWS, GCP, or Azure, based on client requirements.
- **Infrastructure as Code (IaC) & Automation** – Automated provisioning and setup for each cloud provider.
- **Autonomous Self-Healing Infrastructure** – Implemented automated monitoring and recovery mechanisms for resilience.
- **Cloud-Specific Compliance & Security** – Ensured adherence to best practices for AWS, GCP, and Azure environments.
- **Integrated Observability** – Provided real-time performance insights tailored for each cloud provider.
- **Optimized Developer Experience** – Enabled a self-service UI for seamless cloud-specific deployments.

Impact Delivered

- 99.9% uptime, ensuring high availability for end clients on AWS, GCP, or Azure.
- 70% reduction in provisioning time through automation and self-service capabilities.
- 2x scalability, allowing seamless workload expansion on any supported cloud.
- Improved client satisfaction, with a streamlined, cloud-specific managed service experience.

