

Orchestrating Machine Learning Models on Kubernetes

CAPABILITIES SHOWN



Advisory
Services



Data
Science



Kubernetes
Service



Machine
Learning



Accelerated
DevSecOps

ABOUT THE COMPANY

Navitas Business Consulting is an award-winning, woman and minority-owned, SBA 8(a) certified IT Consulting & Solutions Company. We value passion, integrity, commitment, and collaboration, provide the perfect combination of technology, people, innovation, and expertise, and are committed to helping you improve speed, achieve desired outcomes, and reduce costs. Our core capabilities include multi-cloud enablement, digital transformation, and data and intelligence.



CUSTOMER CHALLENGE

A large banking company is running multiple risk, fraud, and credit line increase Machine Learning (ML) Models on virtual machines (VMs) in its on-premises data center. There was no ability to build and deploy code on-demand and no environment parity between model training and production. Capabilities of the current infrastructure imposed limitations on accuracy, leading to bias and pattern misidentification during recurring cycles. The customer needed an enterprise-grade end-to-end automated data science solution with the flexibility to run on-premises, in multiple public clouds, and support every stage of the ML lifecycle.

NAVITAS SOLUTION

Navitas led the effort by defining the architecture for building a Kubernetes-based platform on the cloud using fully open-source technologies. To deliver agility and speed to the ML lifecycle, we employed MLOps practices with automated pipelines from integrating with the model generation, orchestration, and deployment to health, diagnostics, governance, and business metrics. Our solution provided a consistent approach to deploy, maintain, and monitor Machine Learning (ML) models by multiple application teams. Our solution is deployed on AWS EC2 instances, providing the ability to create multiple Kubernetes clusters with various features across different AWS accounts. Our team developed a stable and resilient ML model serving in Kubernetes clusters with optional customizable features. We continuously monitored the performance and stability of the model using metrics to detect model drifts. Navitas also helped train data scientists on Docker, optimized the sharing of data and cloud costs, and enabled feature-based dataset access across teams.

THE BENEFITS

Developed

the ability to run stateful workloads on Kubernetes

Ensured

fully monitored Kubernetes clusters

Removed

the use of on-demand EC2 instances across various teams

Created

a consistent ML operating model