Case Study: Real-Time Personalization at Scale with AWS ML & Analytics

Case Study Short Description

Graaho Technologies transformed its personalization engine, ALGOREC, using AWS's scalable infrastructure and machine learning services. This enabled the delivery of real-time, context-aware recommendations to users—boosting engagement, increasing sales, and reducing operational costs through automation and intelligent infrastructure.

Problem Statement / Definition

To meet the rising demand for hyper-personalized digital interactions, Graaho needed to:

- Build a robust recommendation engine based on machine learning and behavioral data.
- Enable seamless data processing for real-time predictions and scheduled model training.
- Ensure cost-effective, scalable infrastructure capable of handling fluctuating loads.
- Maintain strict security and data compliance throughout the ML lifecycle.

Proposed Solution & Architecture

Graaho Technologies built a modular, full-stack ML architecture with real-time inference and analytics using the following AWS services:

User Authentication & API Management:

- Amazon Cognito for secure user sign-in and access control.
- Amazon API Gateway to expose and manage backend APIs.

Compute & ML Orchestration:

• Amazon EC2 for hosting backend logic, inference services, and orchestration tools like Apache Airflow.

Storage & Data Processing:

- Amazon S3 for storing training datasets, artifacts, and behavioral data.
- Amazon RDS (PostgreSQL) for transactional data and feature generation.

ML Training & Inference Pipelines:

- Apache Airflow orchestrates training pipelines.
- EC2-based inference servers deliver real-time recommendations based on user activity.

Messaging & Communication:

• Amazon SQS for decoupled messaging between services.

Monitoring & Observability:

• Amazon CloudWatch for performance monitoring, logging, and alerting.

Outcomes of Project & Success Metrics

Area	Before AWS	After AWS Migration
App Performance	Unstable during traffic	40% lower latency with consistent
	surges	response times
User Engagement	Low due to static content	30% increase in user interaction
Revenue Growth	Limited by generic listings	20% increase via personalized
		recommendations
Scalability	Bottlenecks under peak load	Seamlessly handled 50% increase in
		user traffic
Operational Costs	High due to manual	25% reduction via automation and
	provisioning	dynamic scaling
Security &	Disparate, manual processes	Unified, automated controls using AWS
Compliance		services

Total Cost of Ownership (TCO) Analysis Performed

- **Reduced TCO:** Eliminated hardware CAPEX and minimized OPEX with elastic resources.
- Improved Agility: Automated pipelines and patching saved engineering time.
- **Governance Simplified:** IAM, encryption, and compliance tools built into AWS architecture.
- **Support:** AWS Support Services accelerated architecture validation and performance tuning.

Lessons Learned

- Modular, Cloud-Native architecture is key to agile scaling and faster iteration.
- Early integration of orchestration tools like Apache Airflow significantly reduces ML ops complexity.
- Real-time personalization success relies heavily on robust data pipelines and observability.
- AWS support and architectural best practices were crucial to rapid deployment and optimization.

Industry Vertical

Software & Internet / SaaS / Artificial Intelligence

Industry (Other)

Personalization Technology / Digital Experience Platforms

Use Case

Real-time content and product personalization using ML-driven recommendation engines.

ISV Tools and Technology Used

- Apache Airflow for ML pipeline orchestration
- Celery for distributed task execution
- **PostgreSQL** for feature engineering and training data
- Custom ML models deployed on EC2-based inference servers

Related Services

- Amazon EC2
- Amazon S3
- Amazon RDS
- Amazon Cognito
- Amazon SQS
- Amazon API Gateway
- Amazon CloudWatch
- Amazon SageMaker (for future integrations)

Related Competencies

- Machine Learning
- Data & Analytics
- DevOps
- SaaS on AWS

Customer Testimonial

"AWS enabled us to bring ALGOREC to life—faster, smarter, and more securely than we could have imagined. The support from AWS architects, combined with the flexibility of AWS services, empowered our team to focus on innovation while confidently scaling to meet real-time personalization demands."

— ALGOREC

Summary:

Partner: Graaho Technologies

Customer: ALGOREC

AWS Competency Areas: Machine Learning & Analytics

Partner Role: Design and deployment of real-time ML-powered recommendation engine, secure data pipelines, scalable infrastructure, monitoring and orchestration automation
Key AWS Services: Amazon SageMaker, Amazon EC2, Amazon S3, Amazon RDS, Amazon API Gateway, Amazon Cognito, Amazon SQS, Amazon CloudWatch, Apache Airflow on EC2
Business Outcomes: 20% increase in sales, 30% boost in user engagement, 40% lower inference latency, 25% reduction in operational costs, enhanced scalability and security compliance