CASE STUDY

Secure Agentic AI Architecture for Marketing Insights and Personalization on AWS Bedrock

Mactores built a secure, scalable Agentic AI platform on AWS Bedrock for Total Expert to automate marketing insights, improve personalization, and ensure SOC2-aligned, responsible AI operations.

Case Study
Summary

- Delivered a governed Agentic AI architecture using AWS Bedrock, Lambda, Step Functions, and ECS Fargate.
- Automated generation of marketing insights, summaries, and campaign analysis for Total Expert.
- Implemented SOC2-aligned security with private VPC endpoints, granular IAM, and encrypted data flows.
- Embedded Responsible AI guardrails, validation workflows, and explainability mechanisms.
- Enabled accurate, safe, and contextual AI outputs across customer engagement journeys.

About
The Customer

TotalExpert®

Total Expert provides a leading marketing and customer engagement platform used by financial institutions to deliver personalized communications. Their cloud-native ecosystem powers data-driven customer insights, orchestrated campaigns, and compliance-aligned engagement across the customer lifecycle.

Customer
Situation

Total Expert needed to modernize how marketing insights, summaries, and recommendations were generated across customer journeys. Existing processes required heavy manual analysis, lacked real-time context, and created operational bottlenecks for relationship managers. At the same time, the company needed to maintain stringent SOC2 Type II controls while handling sensitive CRM and behavioral data. Any GenAI solution requires high accuracy, low latency, traceability, and alignment with internal governance.

Total Expert also needed seamless integration with its AWS analytics platform, proprietary APIs, and identity providers, while ensuring full explainability and auditability of AI-driven decisions. The challenge was to balance autonomy and control through a secure, scalable Agentic AI framework.

Our
Approach

Business
Outcomes

Technical
Outcomes

Mactores designed a modular Agentic AI architecture using AWS Bedrock models, Claude 3 Sonnet, Titan Text Lite, and Titan Embeddings, accessed through private VPC endpoints. Lambda, Step Functions, and ECS Fargate powered agent workflows for data retrieval, reasoning, validation, and content generation. A SOC2-aligned IAM model enforced granular access, short-lived credentials, mTLS, and per-agent isolation.

Responsible AI guardrails, prompt governance, validation layers, and audit logging ensured ethical, accurate, and explainable outputs. Observability with CloudWatch and OpenSearch enabled continuous monitoring and improvement.

Total Expert reduced manual review workload by **65%** and accelerated customer insight delivery with real-time summarization and recommendation generation. The platform delivered consistently accurate, explainable outputs with zero safety or bias incidents. Marketing and analytics teams gained reliable automation while maintaining full governance, auditability, and operational confidence.

The solution provided **sub-2-second** average inference latency, secure VPC-contained Bedrock access, and auto-scaling compute using Lambda, Fargate, and Step Functions. IAM boundaries enforced zero-trust agent interactions, while KMS encryption, audit logs, validation layers, and moderation guardrails ensured security and compliance. Unified observability enabled rapid debugging, drift detection, and continuous optimization.

Getting
Started

Mactores began with discovery workshops to understand Total Expert's business goals, compliance constraints, and integration dependencies. A benchmarking phase evaluated Bedrock models for accuracy, cost, and latency. Architectural design established secure connectivity, IAM boundaries, and orchestration patterns. A pilot deployed MVP agent workflows using synthetic data before scaling to production. Mactores delivered IaC templates, governance documentation, and training to enable seamless adoption and ongoing iteration.

MACTORES

Our
Solutions
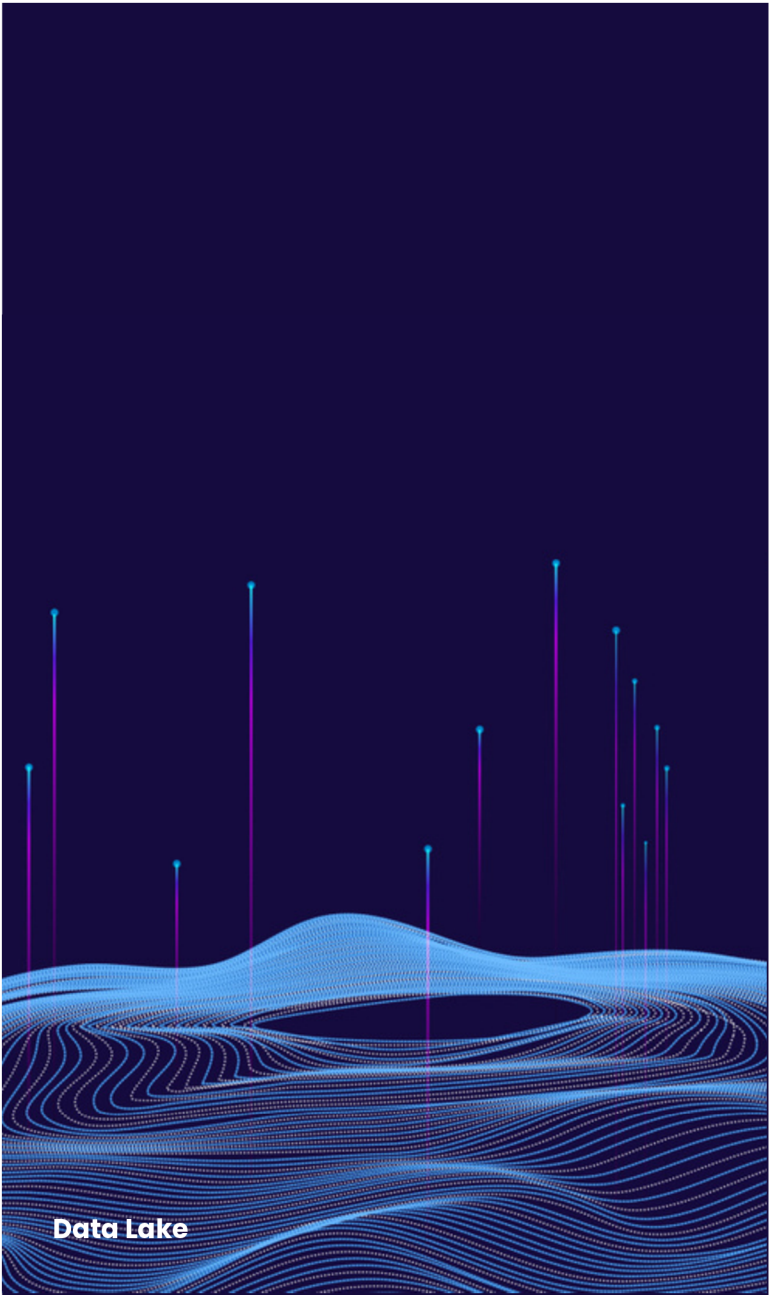
Data Lake

Data Lake

Databases

Data Ops

Generative AI

Our **Process**

Digital transformation via assessments, migration or modernization

We work alongside your tech team to assess and strategize what you need and how to implement the right data solutions on time, on budget and with c-suite buy in.

Assess

- Discovery Automation
- Future State Assessment
- GAP Analysis
- End State
- Road Map
- TCO

Migrate

- Strategy
- Execution
- Migrate
- Migration Acceleration

Modernize

- StrategyFuture State
- Design
- Build
- Automate

MACTORES

AWS Validated
Competencies

- aws PARTNER — Advanced Tier Services
- aws PARTNER — ML Services Competency
- aws PARTNER — Data & Analytics Services Competency
- aws PARTNER — DevOps Services Competency
- aws PARTNER — Migration Services Competency
- aws PARTNER — Healthcare Services Competency
- aws PARTNER — AI Services Competency
- aws PARTNER — Generative AI Services Competency
- aws 200 CERTIFIED
- aws PARTNER — Immersion Day

AWS Validated
Service Deliveries

- aws PARTNER — Amazon Redshift Delivery
- aws PARTNER — Amazon RDS Delivery
- aws PARTNER — AWS Database Migration Service Delivery
- aws PARTNER — Amazon EMR Delivery
- aws PARTNER — AWS CloudFormation Delivery
- aws PARTNER — AWS Glue Delivery
- aws PARTNER — Amazon QuickSight Delivery
- aws PARTNER — Amazon Kinesis Delivery
- aws PARTNER — AWS Config Delivery
- aws PARTNER — Amazon MSK Delivery
- aws PARTNER — Amazon DynamoDB Delivery
- aws PARTNER — AWS Service Catalog Delivery

MACTORES

www.mactores.com/lets-talk

Mactores Cognition Inc

2018 156th Ave NE

Bellevue, WA 98007