Case Study: Scaling Outreach & Cutting AWS Spend

Role: Lead Software Developer / Software Architect • AWS Serverless • Data Platforms • LLM Automation

Client / Org	Confidential (B2B Marketing Automation)
Timeframe	2019–2024 (multi-phase)
Scale	Throughput: 2M+ prospects/month • Real-time data QA • 24×7 operations
Role & Team	Lead Developer/Architect • 4 engineers • cross-functional with marketing/ops

Problem

Lead generation throughput was capped at ~100k prospects/month with significant manual QA bottlenecks, rising AWS costs, and limited visibility into pipeline health. Leadership needed a step-change in scalable automation, lower unit costs, and reliable metrics for decision-making.

Approach

• Designed event-driven, serverless architecture (API Gateway, Lambda, SQS/SNS, S3, DynamoDB).
• Implemented IaC & CI/CD (CloudFormation + GitHub Actions) with environment isolation and automated rollbacks.
• Built ETL/ELT pipelines supporting near real-time enrichment and BI (QuickSight dashboards for exec reporting).
• Integrated LLMs (OpenAI) for text-to-JSON extraction and data validation guardrails as well as a custom, trained ML model to replace manual QA.
• Established observability (CloudWatch metrics/alarms, Sentry tracing) and cost allocation tags with weekly reviews.

Results (Quantified)

• Scaled from ~100k → 2M+ prospects/month; appointments increased by ~112%.
• Cut AWS spend by ~40% via right-sizing, leveraging on-demand/serverless, and reducing wasteful data flows.
• Improved data quality by replacing manual QA with automated LLM validation (higher throughput, lower error rates).
• Shipped exec-facing BI with daily rollups and drill-through into pipeline health and outliers.

Architecture Snapshot (Summary)

Ingress (APIs/webhooks) → Event Bus /SQS (FIFO where needed) → Fan-out workers (Lambda) → Data stores (DynamoDB/S3) → Validation/Enrichment (LLM microservice) → Analytics (Athena/QuickSight). Emphasis on idempotency, retries, and DLQs.

Key Responsibilities

• End-to-end solution design; hands-on implementation and code reviews.
• Security/IAM baselines, VPC design, and cost controls with budgets and anomaly alerts.
• **Operations & reliability**: Sole 24/7 on-call; authored runbooks (SQS/Lambda triage, partial-batch failures, DLQ replays, safe rollbacks, back-pressure) and built CloudWatch alerts/Sentry tracing to cut MTTR.
• **Vendor leadership & light team management**: Selected and procured AWS/OpenAI/QuickSight for best cost/perf; negotiated controls delivering ~40% savings; hired/managed three contract frontend engineers for routine UI components and set review/architecture guardrails.

Tech Stack

• AWS: API Gateway, Lambda, Step Functions, SQS/SNS, S3, DynamoDB, RDS/Aurora, CloudWatch, IAM
• Data/BI: Postgres, ETL/ELT, Athena/Glue, QuickSight
• AI/LLM: OpenAI API, prompt patterns, validation guardrails, text-to-JSON transforms
• Dev: Node.js/TypeScript, Python (FastAPI and for ML model creation / training), Docker, GitHub Actions, CloudFormation

Reference and System Diagram

Reference is available upon request. System diagram is below (click for an interactive version):