



Case Study

HOW DID ARIVIHAN AND AWS BUILD A REAL-TIME NEET QUESTION AI ENRICHMENT PIPELINE?

Solution At a Glance

Arivihan implemented a **serverless, decoupled content-processing pipeline** on AWS to ingest academic questions, enrich them with AI and persist results for downstream use. The solution uses **DynamoDB + Streams** to capture new items, **Lambda** to route workload into **SQS subject queues** and **ECS Fargate** workers to orchestrate Amazon Bedrock (explanations/embeddings) and Amazon SageMaker (custom inference) before storing enriched outputs into an **Amazon S3** data lake. Everything is governed and monitored with **CloudWatch, IAM and CI/CD**.

Key Outcomes Delivered

- **Real-time ingestion to AI enrichment** via DynamoDB → Streams → Lambda → SQS → ECS Fargate, enabling low-touch, scalable processing.
- **Higher resilience and throughput** with subject-wise SQS queues and DLQs to isolate failures and prevent pipeline back-pressure.

- **Accurate, explainable enrichment** by combining **Bedrock** for generation/embeddings and SageMaker endpoints for domain models.
- **Clean data lifecycle** in an **S3 data lake** (raw → processed → enriched → analytics) with versioning and Athena access for analysis.
- **Operational clarity & CI/CD velocity** using **CloudWatch** (metrics/logs/alarms) and CodeCommit/CodeBuild/CodePipeline for automated builds and deployments.
- **Defined DR posture** with Terraform-driven rebuild, DynamoDB **PITR**, S3 versioning and documented RTO/RPO targets.

About Client and their Mission

Arivihan is building India's first fully automated online learning platform to provide one-on-one, interactive and personalized high-quality lectures in vernacular languages. Targeting 200 million low-income and middle-class K-12 students in Tier 2 and smaller cities, Arivihan offers education at just 10% of the current cost, making learning extremely affordable and accessible.

Driven by a mission to bridge the educational gap, Arivihan empowers students from lower-income families who are often limited to local coaching classes. By delivering cost-effective, high-quality education, Arivihan enables these students to pursue academic excellence and competitive exams with confidence.

Architectural Diagram

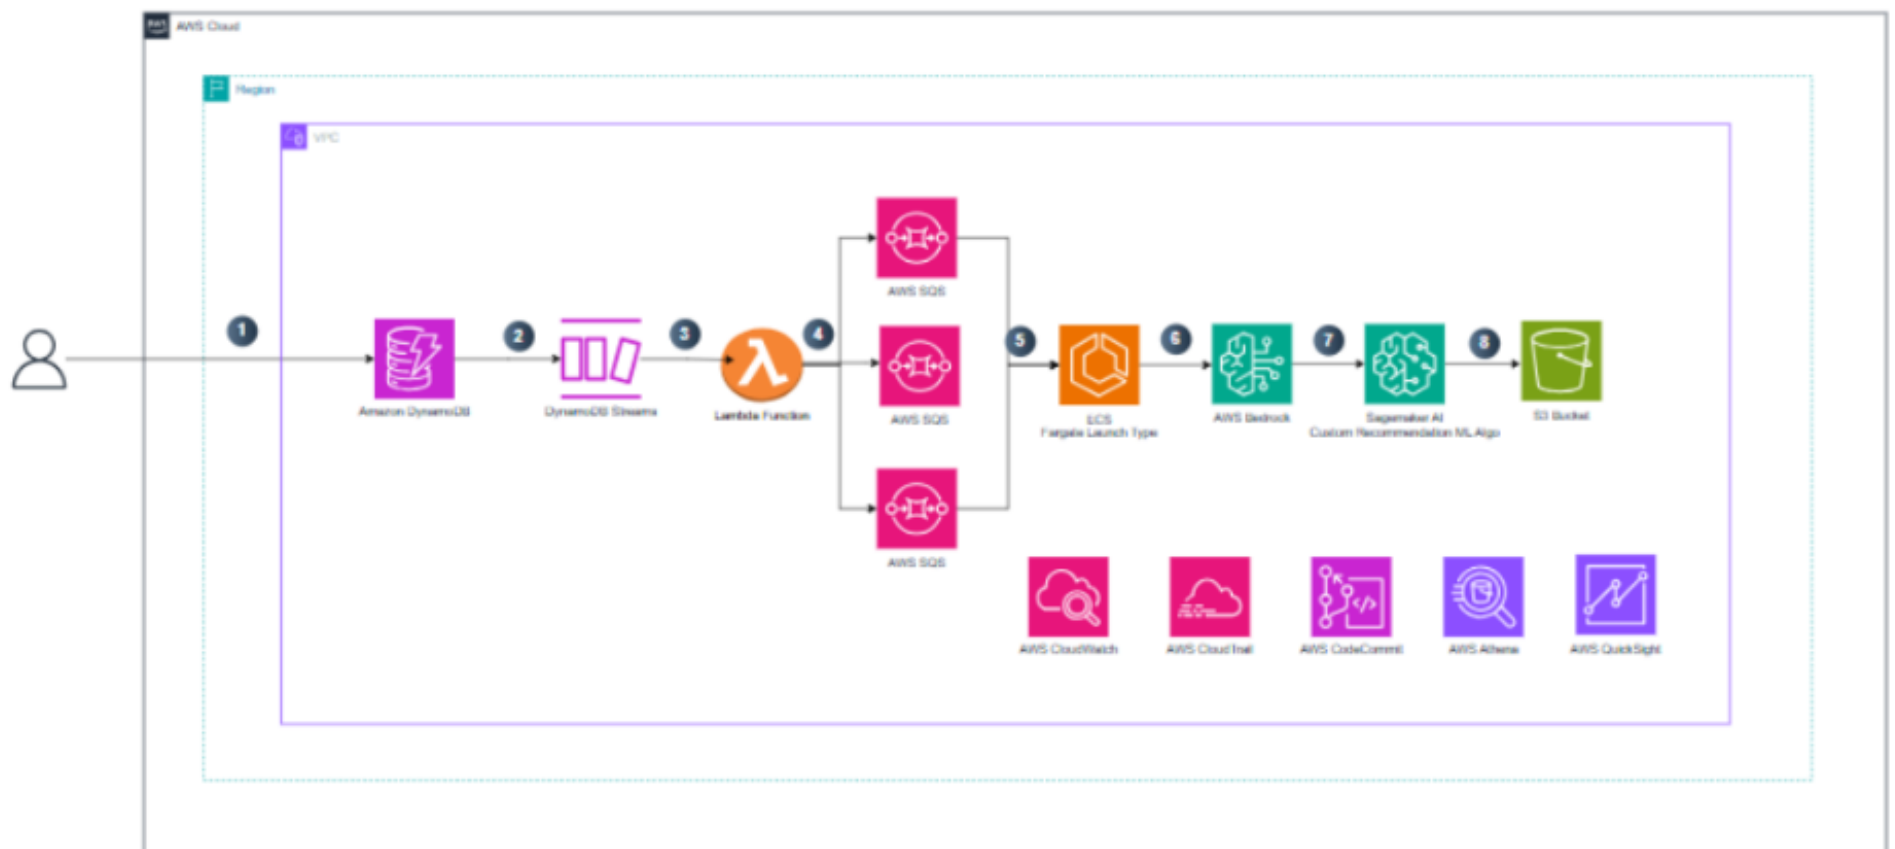
Arivihan needed to **ingest large volumes of questions** reliably, **decouple enrichment steps** to avoid bottlenecks and **blend foundation models with custom ML** while ensuring **least-privilege access, observability and a repeatable CI/CD path**.

Additional requirements included **failure isolation (DLQs)**, **quota-aware Bedrock/SageMaker usage** and a **clear DR plan** with verifiable end-to-end tests.

Key Outcomes Delivered

- Built a **modular, event-driven NEET assessment pipeline** on AWS using **DynamoDB** for question ingestion and **DynamoDB Streams** for real-time triggers.
- Used **AWS Lambda** to validate new questions and route them into **subject-wise SQS queues** (Physics, Chemistry and Biology) with **DLQs** for failure isolation and smooth scaling.
- Ran **ECS Fargate** workers (images in **ECR**) to consume SQS messages and orchestrate AI enrichment via **Amazon Bedrock** (explanations, summaries and embeddings) and **Amazon SageMaker** (custom inference like difficulty scoring).
- Stored enriched metadata back in **DynamoDB** and persisted datasets in a versioned **Amazon S3 data lake** with clear **raw, processed and enriched** layers for auditability and analytics.
- Automated builds and releases with **CI/CD** using **CodeCommit, CodeBuild and CodePipeline** for Lambda packaging, container builds and ECS deployments.
- Monitored health and performance with **CloudWatch** (Streams iterator age, SQS backlog, ECS CPU and memory, Bedrock and SageMaker latency) and secured access with **least-privilege IAM**.
- Enabled fast recovery using **Terraform-driven DR** backed by **DynamoDB PITR, S3 versioning** and reusable model artifacts for quick SageMaker endpoint redeployments within defined RTO and RPO.

End-to-End Architecture Diagram



Result and Impact

- **Scalable, Decoupled Processing** Subject-wise SQS queues with DLQs absorb spikes, isolate failures and keep enrichment moving without blocking upstream ingestion.
- **Better Quality with Composable AI Bedrock** enables high-quality generation (explanations/embeddings) while **SageMaker** serves custom inference (e.g., difficulty), producing richer metadata for each question.
- **Faster, Safer Releases** A first-class **CI/CD pipeline** automates builds, tests and blue/green style updates for Lambda and Fargate tasks, reducing manual errors and time-to-change.
- **Clear Operations & Governance CloudWatch** dashboards and alarms track stream iterator age, queue depth, ECS utilization and model latencies; runbooks define thresholds and actions.

About Us



Years of Exp

15+



Data Center

10+



Awards

100+

AceCloud is a leading provider of end-to-end cloud computing solutions to global organizations across industries at scale. It offers a full spectrum of cloud services, including Public Cloud, Application Hosting, AWS Services, Managed Security Services, and Hosted Virtual Desktop Solutions.

AceCloud is a brand of Real Time Data Services (RTDS) group of companies – a leading provider of information technology capabilities specializing in Cloud Computing and Cloud Telephony. RTDS Group has an employee base of 600+ across India, the US, and the UK, supporting over 20,000+ customers and IT infrastructure in 10+ data centers spanning the globe. It offers industry-leading technological solutions that help customers streamline their operations and enhance efficiency.

As an advanced consulting and certified AWS Well-Architected Partner, AceCloud recently achieved the prestigious AWS Storage Competency and is also recognized as an Amazon RDS Service Delivery Partner. These achievements underscore its deep expertise in AWS technologies and commitment to delivering high-quality cloud solutions.



- Public Sector
- Storage Consulting Competency
- Small and Medium Business Services Competency
- Amazon RDS Delivery
- AWS Database Migration Service Delivery
- Amazon CloudFront Delivery
- AWS Well-Architected Partner Program
- Authorized Commercial Reseller



Contact Us

Are you ready to cut down your cloud spending? Contact AceCloud Experts today to get a customized quote.



awspartner@acecloud.ai



<https://acecloud.ai/aws/>

Our Location

**2069 Oneill Dr, Bethel
Park, Pennsylvania,
15102-6602**

**2637 E Atlantic Blvd
#23519 Pompano Beach,
FL 33062, Florida**

**Ace Tower, 809-A, Udyog
Vihar, Phase 5, Gurugram
122016, Haryana, India**