

Asiri Jayawardena

Software Engineer | Data Platforms & Agentic AI

Mount Lavinia, Dehiwala, Sri Lanka • +94 75 826 4839 • asirijayawardena920@gmail.com
GitHub • LinkedIn

Summary

Software Engineer with professional experience building and operating enterprise-scale, event-driven data platforms on AWS, high-throughput Kafka streaming pipelines, observability, and resilient microservices across staging and production environments. Strong full-stack foundation (MERN, Java) and an active researcher in agentic AI.

Experience

Associate Software Engineer

September 2025 - July 2026

Sysco LABS Sri Lanka

- Developed and operated Kafka-based consumer and producer services on Confluent Cloud (Avro, Schema Registry), delivering resilient, high-throughput streaming across non-production and production environments.
- Maintained the platform's ETL data flows and built an application to actively monitor these services, reducing operational overhead and manual intervention.
- Built internal developer tooling and agentic systems, including an AWS Self-Service Deployment Platform (IDP) that standardized deployments and plugin-based agentic workflows (Google ADK) that automated engineering tasks and reduced onboarding effort.
- Owned end-to-end delivery of the services I built, authoring Jenkins and CodeBuild CI/CD pipelines and provisioning AWS infrastructure (ECS Fargate, Lambda, Step Functions, DynamoDB, CloudFormation, SAM).
- Built and tuned platform observability with Datadog across ECS, RDS, APM, and Kafka consumers for STG and PROD; developed performance dashboards and conducted multi-phase Kafka load testing (JMeter + Groovy, AWS SSM).

Quality Engineering Intern

September 2024 - September 2025

Sysco LABS Sri Lanka

- Created and maintained automated test suites for internal applications, expanding regression coverage and reducing manual testing effort.
- Developed internal testing tools to validate Kafka-based microservices, exercising both producer and consumer paths against expected message contracts and schemas.
- Built a data-verification tool that connected to multiple IBM Db2 databases concurrently, executed cross-database queries, and performed automated validation to confirm data integrity across systems.

Projects

Enterprise-Grade Kafka Data Consumer

Sysco LABS

High-throughput, configuration-driven streaming ingestion

- Implemented a Kafka consumer from the ground up, sustaining ingestion of **100M+ records per hour** into the pricing platform.
- Designed it as a configuration-driven service, where onboarding a new topic requires only a JSON configuration, with no code changes, making it reusable across any topic.
- Owned resilience (retry handling), authored the Jenkins deployment pipelines, and built end-to-end Datadog monitoring and alerting.

Enterprise-Grade Kafka Producer

Sysco LABS

Configuration-driven publishing service (designed and developed; pending production deployment)

- Designed and developed a configuration-driven Kafka producer that publishes data to multiple Kafka topics on a configurable schedule, with all topics, mappings, and publishing behaviour defined through a single JSON configuration.
- Implemented the Transactional Outbox pattern with a polling publisher, enabling reliable event delivery, configurable topic routing, and synchronization between the database and Kafka.

ETL Monitoring & Recovery Tool

Sysco LABS

Serverless ETL tracking and automated recovery service

- Built an advanced monitoring tool that tracks ETL file processing across services, aggregating state from multiple DynamoDB and MySQL databases.
- Added a recovery mechanism that automatically retries failed AWS Step Functions.
- Deployed the serverless backend on AWS Lambda via AWS SAM and the UI on AWS ECS, authoring the end-to-end Jenkins CI/CD pipelines for both.

OMEGA Agentic System

Sysco LABS

Plugin-based agentic workflow platform (Google ADK)

- Built from the ground up on Sysco's internal developer platform (SAGE) to author and run agentic workflows, using Google's Agent Development Kit (ADK) and following industry-standard architecture.

- Designed around a plugin architecture so new agentic workflows can be onboarded with minimal code, significantly reducing the effort to extend the system.
- Built a flagship autonomous workflow that identifies and cross-references relevant Confluence requirements with the codebase without requiring explicit links, detects implementation gaps, generates an analysis report, and, upon approval, creates a Jira ticket and opens a pull request with the proposed changes.

INTELLEA - Multi-Agent Adaptive Learning Platform

2025 - 2026

Final-year research project, SLIIT

Team Lead

- Led the team and architected an AI-powered adaptive learning platform as four independently designed modules, with fully modular agentic pipelines built on a plugin-based architecture and production-oriented patterns: stateful orchestration, retrieval grounding, and fault isolation.
- Personally built **ATLAS**, the platform's intelligence unit, comprising a shared knowledge base, an assignment-creation engine, and a virtual tutoring component; combined Bayesian Knowledge Tracing with LangGraph-orchestrated agent workflows and Graph RAG over a Pinecone vector store.
- Backed the automated grading and document-ingestion pipelines with Apache Kafka, giving the platform an event-driven architecture.
- Designed and integrated the entire CI/CD, building the Jenkins pipelines and owning the full deployment of the platform to AWS.
- Integrated OpenAI APIs, ElevenLabs TTS, and Firebase, with fallback handling for graceful degradation.

AWS Resource Allocator

Personal Project

Self-service AWS provisioning platform

- Built a self-service platform for managing and provisioning AWS resources within a single account: pre-approved CloudFormation templates for different resources are versioned in S3, and developers request resources through a web UI that provisions the stacks automatically once an admin approves.
- Version-controlled the templates through Git: changes merged to the develop branch sync automatically to S3, with the system creating new folders as they are added to the repository, plus support for custom resource tags.
- Powered the backend with AWS Lambda and deployed the frontend on ECS.

Academic Projects

EV Charging Station Application (*Team project; React Native, .NET*) - cross-platform mobile application for EV charging stations, built with a React Native client and a .NET backend.

Multi-Vendor Food Delivery System (*Team project; AWS, microservices, Kafka*) - helped architect the full system through to deployment - a microservices cluster on AWS ECS with Jenkins CI/CD self-hosted on EC2, and a monorepo whose services each deploy independently via their own pipeline. Supports multi-vendor onboarding, cross-vendor order management, and delivery management, with Kafka streaming powering driver allocation, real-time order tracking, and notifications.

Education

Bachelor of Science in Information Technology - Specializing in Software Engineering

Sri Lanka Institute of Information Technology (SLIIT), Malabe

June 2022 - September 2026

- **First Class** degree, awaiting convocation.
- CGPA: 3.81 / 4.0
- Dean's List: Y1S1, Y1S2, Y2S1, Y4S1, Y4S2
- Relevant coursework: Java OOP, Web Programming, OSSA, Data Structures & Algorithms, Android Development and Agentic AI

Extra-Curricular

Vice President - Software Engineering Student Community (SESC), SLIIT

July 2025 - Present

- Led multiple community-driven initiatives, including industry visits, technical workshops, and hackathons.

Technical Skills

Languages: Java, Python, TypeScript/JavaScript, SQL

Backend: Spring Boot, FastAPI, NestJS, Node.js, Express, .NET/C#, Java Servlets/JSP

Frontend: React, MERN, HTML/CSS, Bootstrap, Three.js, React Native/Expo, Vite

Databases & Streaming: MongoDB, DynamoDB, Aurora/PostgreSQL, IBM Db2, MySQL, Apache Kafka, Confluent Cloud, Avro, Schema Registry, Athena, S3

Cloud & Infra: AWS (ECS Fargate, Lambda, Step Functions, EC2, CloudFormation, SAM, CodeBuild, SSM), Azure (Logic Apps, Power Automate), Firebase

Observability & AI: Datadog, CloudWatch, LGTM stack; LangGraph, Google ADK, RAG, Pinecone, MCP, OpenAI APIs

CI/CD & Practices: Jenkins, CodeBuild, Git, Agile, Data Structures & Algorithms **Testing:** JMeter, Groovy

References

Dr Kalpani Manathunga - Head of Department, Software Engineering, SLIIT • 077 350 3324 • kalpani.m@slit.lk

Lasitha Weerasinghe - Software Architect, Sysco LABS • 077 465 4123

Gayumi Kottage - QA Engineering Lead, Sysco LABS • 076 323 2279