MOHAMMED NISARUDDIN

Senior AI Engineer | Agentic Systems, RAG & GenAI in Production | AWS Data Engineering

Hyderabad, Telangana, India • +91-8247555607 • mohammednisar623@gmail.com • LinkedIn • GitHub

PROFESSIONAL SUMMARY

Senior AI Engineer with 9+ years in software and data engineering, the last three focused on shipping production-grade GenAI systems for regulated enterprise clients. Hands-on builder of RAG pipelines, multi-agent systems (LangChain, LangGraph, AWS Bedrock AgentCore, Model Context Protocol), and the evaluation/observability layer (Langfuse, LLM-as-judge) that keeps them trustworthy in production. Delivered a document-intelligence platform processing 5,000+ insurance policies/month, cutting extraction time by 93% and manual review by 80%+. Complements this with a solid AWS data engineering background — PySpark/AWS Glue pipelines, large-scale SQL Server-to-PostgreSQL migrations, and a Delta Lake-to-Databricks Lakehouse migration — giving end-to-end ownership from raw data to deployed AI agent.

CORE SKILLS

Agentic AI: LangChain, LangGraph, AWS Bedrock & Bedrock AgentCore SDK, Multi-Agent Orchestration, Model Context Protocol (MCP) servers & clients, ReAct patterns, AI agents & tool use

Generative AI / LLMs: Azure OpenAI, OpenAI APIs, Retrieval-Augmented Generation (RAG), Prompt Engineering & Versioning, Structured Outputs, Meta LLaMA, Rasa

RAG & Retrieval: Azure AI Search, Multi-Query & Hybrid Retrieval, Semantic Search & Embeddings, Vector Databases, Intelligent Document Chunking, Text2SQL concepts

AI Evaluation & LLMOps: Langfuse, LLM-as-Judge, Three-Layer Scorer Design (programmatic → LLM judge → human calibration), Model Evaluation, Observability, Cost Tracking, Semantic Caching, PyTorch/scikit-learn (working knowledge)

AWS & Data Engineering: AWS Glue, PySpark, Apache Airflow / AWS MWAA, AWS Bedrock, SQL Server → PostgreSQL Migration, Databricks / Delta Lake Lakehouse, ETL/ELT Design, Pandas

Cloud, APIs & DevOps: Azure, AWS, GCP, Docker, Kubernetes, CI/CD, Python, SQL, JavaScript, REST APIs, Microservices, FastAPI, Flask, Django, React

PROFESSIONAL EXPERIENCE

Senior Software Engineer / Senior AI Engineer — Solugenix **Nov 2023 – Present**

GenAI & Agentic AI — Mercer (Marsh McLennan) Benefits Document Intelligence & Deluxe Finance Platform

- Own end-to-end delivery of a Benefits Document Intelligence platform for Mercer — from architecture through production rollout — using Azure OpenAI, LangChain, and FastAPI, turning an ambiguous healthcare-insurance extraction requirement into a scalable technical design that now processes 5,000+ policies/month.
- Built a production RAG pipeline pairing Azure AI Search semantic retrieval with LLM-based extraction, cutting policy extraction time by 93% and manual review effort by 80%+; traced and fixed benefits-schedule field hallucinations by moving from single-query to LangChain multi-query retrieval combined with hybrid search.
- Implemented Model Context Protocol (MCP) servers and clients to expose internal data sources and tools to LLM agents in a standardized, reusable way, improving agent tool-integration speed and reliability.
- Built and deployed production AI agents on AWS Bedrock AgentCore for the Deluxe Finance Platform, applying the same agentic-tool-use patterns across a second enterprise cloud stack.
- Designed and stood up an LLM/AI evaluation framework for the Deluxe Finance Platform from scratch on Langfuse (self-hosted on AWS), using a three-layer scorer model — programmatic evaluators, LLM-as-judge, and human calibration — to assess agent, RAG, and extraction outputs.
- Instituted LLMOps practices — prompt versioning, model evaluation, cost tracking, prompt-response and semantic caching, observability — cutting overhead and boosting throughput by 25%.
- Built RESTful APIs connecting the platform to 3+ enterprise systems, cutting integration delivery time by 35%; containerized services with Docker and deployed to Kubernetes for enterprise-grade production.

AWS Data Engineering — R360+ Banking Data Migration & Mercer Lakehouse Platform

- Core contributor on the R360/R360+ banking data migration (SQL Server → PostgreSQL via AWS Glue/PySpark), debugging production data-integrity issues in migrated batches and fixing Spark JDBC SSL and non-deterministic query bugs.
- Built a column standardization mapping engine (7 REST endpoints: single/bulk mapping, list, update, delete, YAML export) that drives a config-as-data AWS Glue job for matching and harmonizing bank source-field names at scale.
- Led the Mercer client platform's data layer migration from Delta Lake to a Lakehouse architecture on Databricks and provides ongoing platform support.

Senior Software Developer — Pinnacle Data Reliability **Apr 2023 – Aug 2023**

- Built an analytics and visualization platform ingesting real-time data from 10+ sources, used daily by ~200 people for business decision-making.
- Added GenAI-powered predictive analytics to the dashboards, improving forecasting accuracy by 30%.
- Introduced Azure Data Services for ingestion and transformation, speeding up data processing by ~60%.

Technical Consultant — Netscore Technologies Limited **Dec 2020 – Apr 2023**

- Cleaned and prepared data for Monte Carlo simulations in portfolio management — 2M+ records at 99% accuracy — delivering simulation models that improved investment decision accuracy by ~20%.

Software Engineer — Privus Technologies Limited **Apr 2019 – Feb 2020**

- Built secure backend systems for business-critical operations, improving performance by ~15%.

Associate — Amazon **Dec 2018 – Apr 2019**

- Validated and quality-checked large training datasets for ML model development.

Senior Associate — Wipro Limited (Google account) **Mar 2017 – Oct 2018**

- Used SQL and Looker to detect and prevent fraud across millions of transactions on Google's transaction-monitoring systems.

EDUCATION

B.E., Electronics and Communication Engineering

Osmania University, Hyderabad, India • Apr 2012 – Jun 2016 • CGPA: 8.20