MOHAMMED NISARUDDIN

Senior Software Engineer — Data Engineering, Cloud Platforms & AI-Enabled Systems

Hyderabad, Telangana, India | +91-8247555607 | mohammednisar623@gmail.com | linkedin.com/in/Mohammed-Nisaruddin | github.com/Nisar-1234

Notice Period: 60 days | 9+ years experience | Open to Hyderabad / Bengaluru / Hybrid / Remote

PROFESSIONAL SUMMARY

9+ years of software and data engineering experience, with deep hands-on expertise in PySpark and distributed data processing on AWS Glue, including ETL/ELT pipeline design, data quality validation, and schema/column standardization. Databricks experience migrating a data lake to a Lakehouse architecture. Own delivery of a production RAG data pipeline (Azure OpenAI, LangChain, Azure AI Search) processing 5,000+ documents/month, and contribute to a large-scale SQL Server-to-PostgreSQL banking data migration platform built on AWS Glue and PySpark.

CORE SKILLS

Data Engineering & Pipelines: PySpark, AWS Glue, Databricks, Apache Airflow (MWAA), batch ETL/ELT design, data quality validation gates, schema/column standardization

Lakehouse & Data Platforms: Databricks Lakehouse migration, Unity Catalog fundamentals, distributed Spark processing

Cloud & DevOps: AWS (Glue, MWAA, API Gateway, S3), Azure (OpenAI, AI Search, Data Factory, Synapse Analytics), Docker, Kubernetes, Azure DevOps CI/CD

AI/ML & Data Governance: RAG pipelines, LLM-based extraction, MCP servers/clients, LLMOps, RBAC, SAML SSO (Okta), data validation/lineage controls

Integration & Languages: REST API design, RabbitMQ; Python, FastAPI, Flask, Django, PostgreSQL, SQL Server

PROFESSIONAL EXPERIENCE

Senior Software Engineer — Solugenix

Nov 2023 – Present

Mercer (Marsh McLennan) — Benefits Document Intelligence Platform

- Owned end-to-end architecture and delivery of an AI-ready data pipeline turning unstructured insurance documents into structured, retrieval-ready data, using Azure OpenAI, LangChain, and FastAPI.
- Migrated document data storage from a data lake to a Databricks Lakehouse architecture, building ETL transformations on Databricks to support downstream retrieval and analytics.
- Built a production RAG pipeline pairing Azure AI Search semantic/vector retrieval with LLM-based extraction, cutting policy extraction time by 93% at a throughput of 5,000+ policies/month.
- Instituted LLMOps practices — prompt versioning, model evaluation, cost tracking, semantic caching, and observability — improving throughput by 25%.
- Built REST APIs connecting the platform to 3+ enterprise systems, cutting integration delivery time by 35%; containerized with Docker and deployed to Kubernetes.

Deluxe (Payments & Data) — R360/R360+ Banking Data Migration Platform

- Contribute to a large-scale SQL Server-to-PostgreSQL banking data migration platform built on AWS Glue and PySpark, including diagnosing and fixing data-integrity issues in production migration batches.
- Designed and built a column-standardization service (7 REST endpoints, bulk spreadsheet upload, YAML export, validation/completeness gates) consumed directly by the Glue transformation layer.
- Implemented RBAC and SAML SSO (Okta) integration, and worked with RabbitMQ and API Gateway to connect the platform with upstream and downstream systems.

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; introduced Azure Data Services, speeding up data processing by ~60%.
- Added GenAI-powered predictive analytics, improving forecasting accuracy by 30%.

Technical Consultant — Netscore Technologies Limited

Dec 2020 – Apr 2023

- Cleaned and prepared data for Monte Carlo simulations in portfolio management (2M+ records, 99% accuracy), improving investment decision accuracy by ~20%.

Earlier Experience

Software Engineer, Privus Technologies (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.

EDUCATION

B.E., Electronics and Communication Engineering — Osmania University, Hyderabad, India | Apr 2012 – Jun 2016 | CGPA: 8.20