

DANISH ANSARI

Associate Software Engineer | Data Engineer | Backend & Cloud | Full-Stack + AI

danishansari96571@gmail.com | +91-6292022250 | Bengaluru, India | [linkedin.com/in/danish-ansari-05241b201](https://www.linkedin.com/in/danish-ansari-05241b201)

PROFESSIONAL SUMMARY

Associate Software Engineer, **2+ years** of production experience delivering **LLM powered applications**, large-scale **data pipelines**, and **cloud-native backend systems**. Architected an enterprise **GenAI assistant** processing **10K+ daily queries** over a multi TB **Snowflake** warehouse at **99.9% uptime**, and engineered **ML pipelines** processing **500K+ records/day** on **Apache Spark** for Fortune 500 clients. Proficient in **Python, SQL, REST APIs, FastAPI, AWS/GCP, Docker, and Kubernetes**. Adept at independently delivering complex projects with a focus on scalability, reliability, and code quality in Agile environments.

PROFESSIONAL EXPERIENCE

Associate Software Engineer | Relanto Global, Bengaluru

Sep 2024 – Present

Visibility AI Service (Client: Cisco) | Python | FastAPI | LangChain | Snowflake | AWS | Kubernetes | RAG | LLM | PostgreSQL

- Designed and deployed a production-grade **GenAI assistant** using **LangChain** with a hybrid **RAG** and **NL-to-SQL** architecture over a multi TB **Snowflake** data warehouse, reducing analyst query turnaround by **60%**.
- Engineered an **LLM inference** pipeline with entity extraction, **prompt optimization**, and injection-safe **SQL** validation, achieving **sub-50ms** cached response times for over 40% of **10K+ daily requests**.
- Integrated **AWS** services including **Lambda, S3, API Gateway, SQS**, and **CloudWatch** to support data ingestion, service orchestration, and real-time monitoring across the platform.
- Deployed a fault-tolerant distributed backend using async **FastAPI, Kubernetes HPA** autoscaling, and **LLM provider failover**, sustaining **99.9% uptime** under 3x peak traffic load.
- Implemented **blue-green CI/CD pipelines** using **Terraform** and **GitHub Actions** with over 80% automated test coverage, reducing deployment time by **70%** and production incidents by **45%**.
- Collaborated with Data Analysts and Solution Architects in **Agile** sprint cycles to define requirements, review code, and document APIs ensuring high standards of code quality, maintainability, and stakeholder alignment.

Technical Trainee | Relanto Global, Bengaluru

Feb 2024 – Aug 2024

Life Cycle CS&R (Client: Cisco) | Python | Apache Spark | SQL | Hive | GCP | MLflow | Docker | Bash | REST APIs | CI/CD

- Developed large-scale **ETL/ELT batch pipelines** using **Apache Spark** and **GCP**, processing **500K+ records/day** from 3+ heterogeneous data sources improving supply chain forecast accuracy by **22%** and reducing decision latency by **35%**.
- Architected **Hive**-based schema management workflows and automated job scheduling using **Bash/Unix shell scripting**, enabling reliable orchestration of batch workloads across the data platform.
- Trained 3+ **ML models** using **SMOTE** balancing, **hyperparameter tuning**, and **cross validation**; improved anomaly detection recall by **28%** and sentiment analysis **F1 score** by **18%**, with all experiments tracked in **MLflow**.
- Constructed and consumed **REST APIs** for cross system data integration, connecting **PostgreSQL** and **MongoDB** backends to downstream ML services with robust error handling and structured logging.
- Implemented an **NLP inference pipeline** using **spaCy** for structured entity extraction (skills, education, experience) achieving **92% field level accuracy**; containerized training and inference jobs with **Docker**, cutting release cycles from 2 hours to under **15 minutes** via **CI/CD** automation.

TECHNICAL SKILLS

Languages: Python (advanced), SQL, Java, Bash / Unix Shell, JavaScript (React, Angular), C / C++

AI / ML & GenAI: LangChain, LangGraph, RAG, LLM Agents, Prompt Engineering, NL-to-SQL, Embeddings, Vector Databases, Fine-tuning, TensorFlow, PyTorch, scikit-learn, MLflow

Data Engineering: Apache Spark, Snowflake, Hive, Hadoop, ETL/ELT, Data Warehousing, Batch & Stream Processing, Iceberg, Delta Lake, REST API Integration

Backend & Cloud: FastAPI, Flask, REST APIs, AWS (Lambda, ECS, S3, Glue, API Gateway, SQS, DynamoDB, CloudWatch), GCP, Azure, Django

DevOps & Infra: Kubernetes (HPA), Docker, Terraform, GitHub Actions, CI/CD, Git, PostgreSQL, MongoDB Performance Profiling, Agile / Scrum

EDUCATION & CERTIFICATIONS

B.Tech. Computer Science & Engineering | Vellore Institute of Technology

2020 – 2024 | CGPA: 8.35

IBM Certified Artificial Intelligence Analyst