

DANIEL MALUNGU

Senior Data & AI Infrastructure Engineer | PostgreSQL · RAG Pipelines · LLM Inference

+254 721 555 986 · danielkmalungu@gmail.com · linkedin.com/in/danielkmalungu · github.com/DKMalungu

PROFESSIONAL SUMMARY

Senior Data & AI Infrastructure Engineer with 9+ years building production data platforms and AI infrastructure. Currently Lead Data Engineer at ArdhiSasa, Kenya's national land management platform, owning the full data stack: PostgreSQL 17 at scale, ETL pipelines for land transaction digitization, PeerDB change data capture, CI/CD pipelines, and production infrastructure on Rocky Linux and Docker Compose. The platform handles digitization of millions of official government land records.

Founded Sheria Platform — an AI-powered judicial intelligence system for Kenya's court system, built entirely from scratch. The architecture combines a LangGraph agentic RAG pipeline, Qdrant vector search, Neo4j citation graph traversal, self-hosted Ollama inference, semantic caching, and Next.js streaming frontend — all orchestrated by FastAPI with JWT RBAC, Prometheus metrics, and OpenTelemetry tracing. Shortlisted in Kenya's NIRU National AI Hackathon 2025 (Governance & Policy track) from 2,467 proposals out of 5,656 national registrants.

Core expertise: PostgreSQL administration, CDC replication, LangGraph agentic pipelines, vector databases, dbt, Airflow, Snowflake, self-hosted LLM inference, GitHub Actions, and Linux server hardening. Open to senior remote roles in data engineering, platform engineering, or AI infrastructure with US or EU-based teams.

PROFESSIONAL EXPERIENCE

ArdhiSasa

Lead Data Engineer

December 2021 - Present | Nairobi, Kenya

- Designed and administer PostgreSQL 17 production clusters with hybrid master-replica topology using repmgr HA and PgBouncer connection pooling
- Scaled cluster throughput from 1,000 TPS to 10,000 TPS through architectural redesign and query optimization
- Hardened production databases with SCRAM-SHA-256 authentication, audit triggers, and automated backup scripts across 31+ databases
- Built and maintain ETL/ELT pipelines in Python, Apache Airflow, and SQL for land transaction and records data
- Implemented PeerDB change data capture replication from production OLTP to data warehouse with custom schema mapping
- Led digitization of 5 million+ land records through a custom EDMS including bulk document ingestion ETL tool
- Designed and implemented CI/CD pipelines with GitHub Actions for automated data pipeline testing and production deployments
- Served as Scrum Master for the data engineering team — sprint planning, backlog grooming, and agile coaching

Stack: PostgreSQL 17 · PeerDB · Airflow · Python · dbt · Airbyte · MinIO · Docker Compose · GitHub Actions · Rocky Linux · nginx · Nagios · ELK Stack

Sheria Platform

Founder & Lead Engineer

October 2024 - Present | Independent Project

- Shortlisted from 2,467 proposals across 5,656 national registrants — NIRU National AI Hackathon 2025, Governance & Policy track
- Architected a 3-node Ollama GPU inference cluster (RTX 3050s) running Qwen2.5:7b — fully self-hosted with no external API dependency
- Built a LangGraph agentic RAG pipeline: Planner → Retriever → Responder with IRAC legal prompting and streaming NDJSON responses
- Implemented semantic caching (Qdrant + Redis, cosine > 0.95, 30-day TTL) to reduce redundant inference
- Qdrant vector search across Kenya Law Reports with Neo4j citation graph (CITES, OVERRULES, DISTINGUISHES relationships)
- FastAPI microservices with JWT RBAC, Prometheus metrics, and OpenTelemetry tracing
- M-Pesa STK Push payments via Safaricom Daraja API with FastAPI callback on Railway
- Next.js streaming frontend with admin panel, ingestion job monitor, and health dashboard

Stack: Python · FastAPI · LangGraph · Ollama · Qdrant · Neo4j · PostgreSQL · MinIO · Next.js · Docker Compose · Redis

CloudFactory

Data Engineer

March 2022 - November 2022 | Kenya

- Led full migration of all data pipelines from Xplenty to Prefect orchestration on AWS EC2 — achieving centralized logging and significant cost reduction
- Wrote dbt transformations following the Kimball dimensional modelling approach for an enterprise Snowflake data warehouse
- Built a CI pipeline from scratch in GitHub Actions for automated dbt model testing and fast feedback for contributors
- Introduced SQLFluff for automated SQL quality enforcement, improving team delivery speed by 80%

Stack: dbt · Snowflake · Prefect · Python · SQL · GitHub Actions · AWS (EC2) · SQLFluff

Olive Technologies

Full Stack Software Engineer

January 2016 - December 2018 | Nairobi, Kenya

- Built a chatbot system for back-office customer support using TensorFlow, deployed on a microservices architecture
- Developed and launched an online academic marketplace, growing it to 200+ active users from zero
- Designed deployment architecture using microservices principles for scalability and maintainability

Stack: Python · TensorFlow · SQL · Microservices

Open Arch Systems

Software Engineer

November 2015 - September 2016 | Nairobi, Kenya

- Built a smart water meter system on an IoT/NB-IoT platform with bi-directional communication between field sensors and back-office
- Reduced water losses for two water utility companies by 30% through real-time sensor-based incident detection

Stack: Python · Django REST Framework · PyQt · NB-IoT

EDUCATION

MSc Data Analytics — KCA University (May 2019 - November 2022)

BSc Information Technology — KCA University (May 2010 - November 2013)

TensorFlow Developer Certification — Google Mobile Academy (August 2021 - September 2022)

Diploma in Business Management — Kenya Institute of Management (January 2017 - November 2018)

CORE TECHNICAL SKILLS

Data Engineering: PostgreSQL 17 (administration, performance tuning, HA with repmgr), PeerDB CDC replication, Apache Airflow, dbt (Kimball modeling), Snowflake, ETL/ELT pipelines, Python, SQL

AI/ML Infrastructure: LangGraph agentic pipelines, Ollama (self-hosted LLM inference), Qdrant (vector search), Neo4j (graph databases), FastAPI, RAG systems, semantic caching, ColBERT/RAGatouille reranking

DevOps & Infrastructure: Docker Compose, Rocky Linux, nginx, GitHub Actions CI/CD, MinIO (S3-compatible storage), Prometheus, OpenTelemetry, Nagios, ELK Stack, server hardening, backup automation

Development: Python, TypeScript, Next.js, FastAPI, TensorFlow, Django REST Framework, JWT RBAC, microservices architecture

AWARDS & RECOGNITION

NIRU National AI Hackathon 2025 — Shortlisted from 2,467 proposals (5,656 total registrants) in the Governance & Policy track. Advanced to prototyping and mentorship phase. Backed by Safaricom, Equity Bank, KCB, Strathmore University, and USIU-Africa.