

JOSHUA MUSIRA

Senior Software Engineer · Industrial Automation & IoT · AI/ML Solutions
Architect · Full-Stack Developer

✉ joshuamusira01@gmail.com

☎ +254 111 779 815

■ Nairobi, Kenya

■ joshmusira.netlify.app

PROFESSIONAL SUMMARY

Results-driven Senior Software Engineer and Data Scientist with 3+ years of progressive experience architecting enterprise-grade SaaS platforms, industrial IoT systems, cloud-native data pipelines, and AI/ML solutions in oil & gas and industrial engineering. Proven ability to deliver end-to-end solutions — from statistical modeling and predictive analytics to production full-stack systems — across AWS, Azure, and containerized microservice environments. Deep expertise in Generative AI (LLMs, RAG, Agentic workflows), real-time IoT data infrastructure, SCADA integration (Modbus, MQTT, OPC-UA), multi-tenant SaaS with OAuth 2.0, Stripe billing, Terraform IaC, and compliance-aware design (SOC 2, NIST CSF). Passionate about applying advanced machine learning, NLP, and engineering best practices to drive measurable business outcomes.

CORE COMPETENCIES

■ Industrial IoT & SCADA Integration (Modbus, MQTT, OPC-UA)	■ Enterprise SaaS & Multi-Tenant Architecture
■ Real-Time Data Pipelines (Kafka, InfluxDB, TimescaleDB)	■ Generative AI, LLMs & RAG Systems (OpenAI, Claude, LangChain)
■ Full-Stack Engineering (React / Next.js / FastAPI / Node.js)	■ AWS Infrastructure (ECS/Fargate, Lambda, RDS/Aurora, SageMaker)
■ Predictive Maintenance & LSTM Time-Series Modeling	■ Infrastructure as Code (Terraform, CloudFormation)
■ NLP, Feature Engineering & Advanced ML (Scikit-Learn, TensorFlow)	■ CI/CD & DevOps (Docker, Kubernetes, GitHub Actions)
■ Security (OAuth 2.0, SAML/SSO, TLS, RBAC, SOC 2, NIST CSF)	■ Data Visualization & BI (Power BI, Plotly, Seaborn)

PROFESSIONAL EXPERIENCE

Senior Software Engineer

Landstar Engineers Ltd · Aug 2025 – Present · Nairobi, Kenya | Oil & Gas / Industrial Automation & SaaS

- Architected a cloud-native, multi-tenant SaaS operations platform on AWS (ECS/Fargate, RDS/Aurora, S3, CloudFront) serving 10+ enterprise client organizations with role-based tenant isolation, processing 10,000+ concurrent real-time sensor streams at sub-100ms latency.
- Implemented Infrastructure as Code using Terraform and AWS CloudFormation to provision repeatable cloud environments across dev, staging, and production — enabling zero-drift deployments and full environment parity.
- Deployed production LSTM-based predictive maintenance models integrated into live AWS pipelines via SageMaker-compatible workflows, reducing unplanned equipment downtime by ~30% and cutting manual site inspection frequency by 40%.
- Built a Retrieval-Augmented Generation (RAG) AI assistant using Claude API and OpenAI GPT-4 with Qdrant vector store and LangChain orchestration, enabling natural language querying across 500,000+ equipment manuals with sub-second contextual retrieval.
- Engineered event-driven alert and automation pipelines using Kafka, AWS Lambda, and N8N — automating 20+ hours/week of operational workflows including alert escalation, report generation, and subscription billing notifications.
- Led full-stack development of an industrial operations dashboard (React.js + Node.js + InfluxDB) providing real-time KPI visibility across remote oil & gas field sites; implemented OAuth 2.0, TLS, and RBAC security across all API endpoints and data stores.
- Established CI/CD pipelines via GitHub Actions for containerized deployment to AWS ECS/Fargate, compressing release cycles from days to hours; leveraged Claude Code and AI-augmented code review workflows to accelerate feature delivery.

ICT Engineer (Software)

Landstar Engineers Ltd · May 2025 – Aug 2025 · Nairobi, Kenya | Oil & Gas / Industrial Automation

- Designed and deployed RESTful APIs on AWS Lambda bridging front-end operator interfaces with IoT hardware and SCADA backend services; implemented Modbus TCP/RTU and MQTT communication protocols for real-time sensor integration.

- Developed automation scripts reducing manual data entry and cross-system synchronization tasks, improving operational team productivity by an estimated 35%.
- Documented API specifications and integration protocols for industrial control systems in adherence with secure coding best practices and SOC 2-aware design standards.

Data Science Intern

Sankhyana Consultancy Services Ltd · Mar 2024 – Sep 2024 · Remote | Enterprise Analytics

- Executed end-to-end enterprise analytics projects spanning credit risk modeling, customer segmentation, financial forecasting, and NLP-driven market intelligence for Fortune-level clients.
- Built and validated production ML pipelines in Python (Scikit-Learn, TensorFlow) using Information Value (IV), Weight of Evidence (WOE), and Fisher score feature selection; improved model accuracy by 25% through iterative feature engineering and hyperparameter optimization, delivering 10% cost savings.
- Delivered geospatial analysis and LDA topic modeling that directly informed strategic market expansion decisions, translating complex data insights into actionable executive recommendations.

Machine Learning Engineering Intern

Mentorness · Feb 2024 – Mar 2024 · Remote | ML Engineering

- Developed and optimized ML models for multiple real-world applications; handled data preprocessing, feature engineering, model training with Scikit-Learn, and cloud deployment on Microsoft Azure.

Data Science Intern

InforTech · May 2024 – Jun 2024 · Remote | Machine Learning

- Completed four end-to-end ML projects: car price prediction, shopper sentiment analysis, cinema ticket demand forecasting, and fake news detection using NLP — responsible for full data lifecycle from ingestion and EDA through model training, evaluation, and deployment.

KEY ENGINEERING & DATA SCIENCE PROJECTS

Industrial IoT SaaS Monitoring Platform

Landstar Engineers Ltd · 2025

- Production-grade multi-tenant SaaS on AWS (ECS/Fargate, RDS/Aurora, S3, CloudFront) ingesting real-time sensor telemetry from oil & gas field equipment via Modbus TCP and MQTT over TLS, with tenant-isolated data stores, RBAC, and Stripe subscription billing.
- Provisioned all infrastructure using Terraform; event-driven alert system via Kafka + AWS Lambda auto-triggers shutdown sequences on threshold breaches with engineer notification workflows.

Stack: React.js · Node.js/Express · FastAPI · InfluxDB · Kafka · MQTT · Modbus · Docker · AWS ECS/Fargate · RDS/Aurora · Terraform · GitHub Actions

AI-Powered Operations Intelligence System (RAG)

Internal R&D · 2025

- Engineered a production RAG system enabling field engineers to query 500,000+ equipment manuals and maintenance logs via natural language using OpenAI GPT-4 and Claude APIs; Qdrant vector store with LangChain orchestration delivers sub-second contextual answers.
- Secured with OAuth 2.0, encrypted data pipelines, and role-gated access; built using Claude Code AI-assisted development workflows.

Stack: OpenAI API · Claude API · Claude Code · LangChain · Qdrant · FastAPI · React.js · AWS Lambda · Azure

LSTM Predictive Equipment Failure Model

Landstar Engineers Ltd · 2024–2025

- Designed and deployed an LSTM time-series neural network trained on historical sensor telemetry to predict equipment failures 72+ hours in advance with 87% precision; integrated confidence intervals and recommended maintenance actions directly into the operations dashboard.
- CI/CD pipeline via GitHub Actions automated model retraining and deployment; SageMaker-compatible pipeline for production inference at scale.

Stack: Python · TensorFlow/Keras · InfluxDB · FastAPI · React.js · AWS Lambda/S3/RDS · SageMaker · GitHub Actions

Credit Risk & Loan Approval Intelligence Engine

Sankhyana Consultancy · 2024

- Delivered a production ML pipeline with IV, WOE, and Fisher score feature selection; built a credit scoring SaaS microservice with a PostgreSQL backend, Stripe-compatible billing integration hooks, and a React.js interface — model predicts credit scores and auto-classifies approval, review, or rejection.

Stack: Python · Scikit-Learn · FastAPI · React.js · PostgreSQL · Docker · GitHub Actions

Algorithmic Trading & Market Sentiment Analyzer

Personal / Research · 2024

- End-to-end system correlating real-time stock market movements with social media sentiment using NLP on live Twitter/X streams; deployed as a Next.js SaaS application with Node.js backend, subscription-gated access, and Monte Carlo simulation for option pricing.

Stack: Python · LSTM · yfinance · NLP · Monte Carlo · React.js (Next.js) · Node.js · AWS

Company Analytics Dashboard (ETL + Power BI)

Sankhyana / Personal · 2024

- Built an enterprise BI dashboard in Power BI visualizing key business metrics after designing a normalized schema and orchestrating a full ETL pipeline from MySQL source data; published for live stakeholder access.

Stack: Power BI · MySQL · ETL · Schema Design · Python

TECHNICAL SKILLS

Data Science & ML	Python, Scikit-Learn, TensorFlow/Keras, LSTM, XGBoost, NLP (spaCy, NLTK), Topic Modeling (LDA), Feature Engineering (IV/WOE/Fisher), Geospatial Analysis, Time-Series Forecasting, Monte Carlo Simulation, Pandas, NumPy, Seaborn, Matplotlib
Generative AI	OpenAI GPT-4, Claude API, Gemini, LangChain, RAG, Agentic Frameworks, Claude Code, Qdrant, Pinecone, Vector Stores, Prompt Engineering
Full-Stack	React.js, Next.js, TypeScript, JavaScript, Vue.js, Redux Toolkit, TanStack Query, Tailwind CSS, Node.js, Express.js, FastAPI, GraphQL, WebSocket APIs, Prisma ORM, Redis
Cloud & DevOps	AWS (ECS/Fargate, Lambda, S3, RDS/Aurora, CloudFront, SageMaker, IoT Hub), Microsoft Azure (Azure ML, Azure IoT Hub), Terraform, CloudFormation, Docker, Kubernetes, GitHub Actions, Jenkins, CircleCI, GitOps
Industrial & IoT	MQTT, Modbus TCP/RTU, OPC-UA, WebSockets, SCADA Integration, InfluxDB, TimescaleDB, Kafka, RabbitMQ, N8N Workflow Automation, Microsoft Graph API, ETL Pipelines
Databases	PostgreSQL, MySQL, MongoDB, Redis, InfluxDB, TimescaleDB, Neo4j, Vector Stores (Qdrant, Pinecone)
Security	OAuth 2.0, SAML/SSO, TLS (Encryption in Transit & At Rest), RBAC, Multi-Tenant Isolation, SOC 2 Awareness, NIST CSF, Secure Coding Practices
Visualization & BI	Power BI, Matplotlib, Seaborn, Plotly, Word Cloud, ETL & Schema Design for BI reporting
Payments & APIs	Stripe API, Daraja/M-Pesa, yfinance, NewsAPI, Algorithmic Trading

KEY ACHIEVEMENTS

- Reduced unplanned equipment downtime by **~30%** via production ML-based predictive maintenance models deployed in live oil & gas field operations.
- Achieved **40% reduction** in manual site inspection frequency by implementing real-time IoT sensor monitoring across distributed industrial field sites.
- Improved data model accuracy by **25%** and delivered **10% cost savings** through advanced feature engineering and hyperparameter optimization across enterprise analytics engagements.
- Automated **20+ hours/week** of operational workflows using N8N and AWS Lambda, significantly reducing engineering team overhead.
- Architected a RAG-based AI assistant (powered by Claude API) enabling natural language querying across **500,000+ technical documents**, built with Claude Code AI-assisted development workflows.
- Built and shipped a multi-tenant SaaS operations platform on AWS (ECS/Fargate, RDS/Aurora, CloudFront) with Terraform-managed infrastructure, GitHub Actions CI/CD, and OAuth 2.0 + SOC 2-aware security.
- Developed LSTM neural network achieving **87% precision** in predicting equipment failures 72+ hours in advance — model integrated into live production dashboard with automated maintenance recommendations.

EDUCATION & PROFESSIONAL TRAINING

B.Sc. Software Engineering

Kirinyaga University, Kerugoya, Kenya · Sep 2021 – Apr 2025

- Full-Stack Software Development — SERN Stack, AWS, Docker · *TheJitu, 2023*
- Data Science & Machine Learning — Comprehensive Program · *Sankhyana Consultancy, 2023–24*
- Machine Learning Engineering & Azure Deployment · *Mentorness, 2024*

REFERENCES

Technical & Operations Manager

Landstar Engineers Ltd

+254 787 433 188

agunda@landstargroup.com

Pavitra Biswas

Data Science Lead, Sankhyana Consultancy

+91 9560450026

pavitra.stat@gmail.com

Tapaswani Nayak

Senior Analyst, Sankhyana Consultancy

+91 8217426430

tapaswani.nayak@sankhyana.com

Kelvin Comba

Instructor, TheJitu Nyeri

0715364294

kelvin.gatimu@thejitu.com

Mentorness Supervisor

Mentorness, India

+91 6351567156

mentorness.help@gmail.com