

Abdul Afolabi

Data Engineer

London, UK · enquiries@abdulanalysis.com · abdulanalysis.com · github.com/itsafo · linkedin.com/in/abdul-afolabi-804476211

SUMMARY

Data engineer with 4 years building cloud data platforms on Snowflake, Microsoft Fabric, Azure and AWS, after 3 years in data analytics. I specialise in migrating legacy and SaaS systems into governed, Kimball-modelled warehouses whose numbers reconcile to source, delivered entirely through Terraform and CI/CD. Recent work includes auditing and recovering an AI-assisted Fabric migration for a US medical liability insurer, where I cut capacity cost by 80%.

EXPERIENCE

Data Engineer, Contract Consultant · Exequit

Client: US medical professional liability insurer · Jun 2026 – Sep 2026 · Remote

- Led the technical audit and recovery of an AI-assisted (Claude) migration from a legacy IBM i (AS/400, DB2) system of record to a Microsoft Fabric warehouse, deployed entirely through CI/CD.
- Modelled the claims domain for actuarial and finance users (claims, payments, reserves, recoveries, reinsurance and written premium) as a Kimball star schema, reconciling business measures against the legacy system of record after every load, with monetary totals matched to the penny.
- Cut Fabric capacity cost by 80% by moving the always-on silver and gold layers from full reloads to incremental loads on scheduled, rather than always-on, compute.
- Found that automated parity tests conflated measure-level and table-level reconciliation, which overstated migration accuracy. Redesigned validation and widened the goal from parity testing to a governed enterprise warehouse.
- Designed the governance framework: metadata control and load-audit tables, sensitivity and endorsement labels, and dynamic role-based access control through Microsoft Entra ID.
- Restored data recoverability by enabling time-travel retention on 95% of bronze tables that had been left unprotected.
- Defined the delivery standard from ticket to GitLab CI to Fabric deployment, and modularised pipelines so a failed stage reruns on its own, cutting recovery time by up to 24 hours.
- Maintained Terraform across three roots (azurerm and Microsoft Fabric providers), provisioning workspaces, medallion lakehouses, Container App Jobs, Key Vault and managed identities for dev and UAT.
- Enabled self-service analytics with a dedicated workspace on the SQL analytics endpoint, gold tables and semantic models, isolated from the locked-down CI/CD workspaces. Trained 7 analysts in Fabric and Power BI.

Data Engineer · Emeria

London, UK · Jul 2022 – May 2026

- Replaced Salesforce as the reporting source of truth for properties with a Snowflake warehouse that consolidates Qube MRI, Reapit and DataStation. Built the Power BI reporting layer, which saved 8 hours a week of manual processing and gave stakeholders one view of property volumes.
- Designed reusable, parameterised ELT pipelines using Fivetran metadata, dbt vars, macros and artifacts, and Python UDFs, on a medallion architecture (AWS for internal consumers, Microsoft stack for external).
- Cut data processing time by 80% through table redesign, clustering, micro-partition pruning, materialized views, parallel jobs and orchestration redesign.
- Led the migration and integration of legacy systems, including on-premises SQL Server: batch from Salesforce CRM and Reapit, real-time from Qube MRI, with risk-based access and minimal disruption, in line with ISO 27001 and ISO 9001.
- Introduced Kimball dimensional modelling across the organisation, cutting report load time from 2 hours to 30 minutes and improving lineage and consistency.
- Built the first historical trend model of property volume by type, status, sector and manager, going beyond creation date, which increased trust in and adoption of reporting.
- Reduced health & safety data retrieval from 5 hours to 10 minutes with centralised reporting datasets.
- Implemented pipeline monitoring with automated email alerts to data stewards.
- Set engineering standards for version control, testing, deployment and IaC (GitHub, Terraform, CI/CD, Jira, Confluence). Mentored junior analysts and trained stakeholders on reporting tools.

Data Analyst Lead · **Modulr Finance**

UK · Jul 2021 – Jul 2022

- Built management-information packs and dashboards for leadership. Worked with data engineering on ML pipelines and owned data governance through a CRM warehouse migration.

Data Analyst Specialist · **Tutora**

UK · Feb 2019 – Jul 2021

- Automated manual data collection and reporting. Built a Python + Power BI customer segmentation used for targeted campaigns.

SKILLS

- **Domains:** Insurance (claims, reserves, reinsurance, premium, actuarial reporting), property management, fintech
- **Languages:** Python, SQL, PySpark, DAX, Jinja
- **Platforms:** Microsoft Fabric, SQL Server, IBM i (DB2), Azure (Container Apps, Key Vault, Entra ID), Snowflake, AWS (Glue, Athena, Redshift, Lambda)
- **Data:** dbt, Fivetran, Delta Lake, Kimball modelling, CDC, data governance, Power BI / TMDL
- **Delivery:** Terraform, Docker, GitLab CI, GitHub, Jira, Confluence
- **AI engineering:** Claude agents, sub-agents, plugins and MCP servers

CERTIFICATIONS

- **Anthropic Claude Academy (2026):** Claude Code in Action · Claude Code 101 · Introduction to Model Context Protocol · Building Effective Human-Agent Teams · AI Fluency for Builders · Claude Platform 101 · Claude 101
- Snowflake Data Warehouse · AWS Cloud Practitioner · Google Cloud ML · CFI SQL Fundamentals · GDPR (iHasco)

EDUCATION

- MSc Oil and Gas Engineering, Robert Gordon University (2019–2020)
- Diploma, Geographic Information Systems, University of Aberdeen (2020–2021)
- BEng Mechanical Engineering, Ahmadu Bello University (2010–2016)