

MOHAMED ATEF FAHMI

Senior Data Engineer

+971 507 422 363 • muhamedfahmy7474@gmail.com • Abu Dhabi, UAE | [GitHub](#) • [LinkedIn](#) • [Portfolio](#) • [Medium](#)

PROFESSIONAL SUMMARY

Senior Data Engineer with 6+ years of experience building production-grade Lakehouse architectures, dbt modelling layers, and batch/streaming pipelines at scale. Proven track record owning end-to-end data products in regulated, statistics-driven environments — designing Airflow orchestration, enforcing data quality through CI/CD and automated testing, and leading platform migrations. Hands-on with Trino/Presto distributed query engines, Apache Iceberg and Delta Lake table formats, and AWS and Azure cloud stacks. Published author on Analytics Engineering and dbt. Fluent in English; open to relocation to Barcelona.

PROFESSIONAL EXPERIENCE

Statistics Centre — Abu Dhabi (SCAD) — Senior Data Engineer Sep 2025 – Present | Abu Dhabi, UAE

- Design and manage 95+ daily production data pipelines (Python, Scrapy, Selenium) ingesting national CPI and retail price index data into Apache Iceberg and Delta Lake tables on ADLS Gen2, feeding government statistical reporting systems.
- Built MARSAD/WSMT, a centralised monitoring platform covering all production pipeline runs — triggering, execution tracking, failure alerting, and data quality controls; reduced pipeline failure MTTR by 60%.
- Led source system integrations into a centralised data hub orchestrated by Apache Airflow, including schema mapping, ingestion testing, and deployment governance in a regulated statistical environment.
- Leveraged Trino for large-scale distributed SQL transformations on the Databricks Lakehouse layer, replacing ad-hoc notebook queries with governed, version-controlled transformation models.
- Mentored a team of data engineers on dbt best practices, code review standards, and CI/CD workflows; introduced peer review gates that cut production incidents by 35%.

e& (Etisalat) — Senior Data Engineer

Nov 2024 – Sep 2025 | Dubai, UAE

- Re-engineered the Spark/Scala ETL framework powering customer segmentation, reducing end-to-end processing time from 16 hours to 2 hours — an 87.5% reduction — across a dataset spanning tens of millions of customer records.
- Designed and maintained batch data workflows on HDFS using Spark, Hive, Impala, and Presto for large-scale campaign management and finance reporting; introduced Delta Lake as the storage layer for reliable, versioned data ingestion.
- Built a custom ETL framework with dynamic Scala compilation enabling analysts to rapidly define new pipeline logic without engineering support; collaborated with product and finance stakeholders across international teams.

LigaData — Data Engineer

Oct 2022 – Nov 2024 | Remote

- Built and optimised ETL pipelines using Python, Hive, and Airflow ingesting large datasets from diverse sources into HDFS; reduced report generation time by 40% through workflow automation and query optimisation.
- Contributed to a data fabric product deployed on Kubernetes, using Delta Lake for schema-enforced, scalable data integration across cloud environments; delivered Power BI dashboards for business decision-making.

Upwork — Freelance Data Engineer & Analyst

Jun 2022 – Dec 2023 | Remote

- Delivered ETL processes and API integrations for clients across e-commerce and logistics; developed ML models achieving 85% prediction accuracy; maintained consistent 5-star feedback across all engagements.

KEY PROJECTS

dbt Analytics Engineering Project | github.com/Muhammadataf

- End-to-end dbt project covering staging, intermediate, and mart layers with data tests, documentation, Jinja macros, and CI/CD integration — the foundation for a Medium article on Analytics Engineering published in Towards Dev (Feb 2023).

AWS Retail Data ETL Pipeline | github.com/Muhammadataf/AWS-Retail-Data-ETL-Pipeline

- Complete AWS data engineering pipeline for retail sales analysis: MySQL on RDS → AWS Glue ETL (star schema transformation) → S3 (Parquet) → Amazon Athena for SQL querying. Infrastructure provisioned as code with Terraform. Demonstrates IaC discipline and cloud-native ELT design.

Kakashi — Open-Source PII Masking CLI | github.com/Muhammadataf/kakashi

- Python CLI tool for privacy and PII masking in AI agent workflows, supporting 50+ file formats and 20+ agent integrations. Demonstrates data governance mindset directly applicable to regulated fintech environments.

Covid-19 End-to-End Pipeline — Azure Lakehouse

- Full Azure pipeline: ADF, ADLS Gen2, Azure SQL, Databricks for transformation; CI/CD via Azure DevOps; Power BI dashboards for reporting. Demonstrates Azure Lakehouse architecture and DevOps discipline.

EPL Analytics Platform — Arsenal FC | github.com/Muhammadataf/ArsenalFC-Data-Pipeline-Project

- End-to-end football analytics pipeline: PySpark, Docker, PostgreSQL, Apache Airflow, Power BI. Medallion Architecture with 11 dashboards covering player, goalkeeper, and match performance data (2017–2023).

TECHNICAL SKILLS

Orchestration	Apache Airflow, Dagster (asset-based), dbt (Core & Cloud)
Lakehouse Formats	Apache Iceberg, Delta Lake, Apache Hudi, Parquet
Query Engines	Trino, Presto, Apache Spark, Hive, Impala, HDFS
Streaming	Apache Kafka, Apache Flink, Spark Streaming
Cloud & Platforms	Azure (ADF, ADLS Gen2, Databricks, Synapse, DevOps), AWS (S3, Glue, RDS, Athena), Kubernetes
Data Modelling	Dimensional modelling, star schema, SCD, dbt staging / intermediate / marts
Data Quality	dbt tests, Great Expectations, data contracts, CI/CD validation
Languages	Python (Pandas, NumPy, Scrapy, Selenium), SQL, Scala, Bash
Data Warehousing	Databricks Lakehouse, Azure Synapse, Azure SQL, Snowflake
DevOps / CI-CD	Azure DevOps, Git, GitHub Actions, Docker, Kubernetes, Terraform
BI & Reporting	Power BI, Matplotlib, Seaborn

EDUCATION & CERTIFICATIONS

B.Sc. Computer Engineering — Mansoura University

Microsoft Azure Data Engineer Associate (DP-203) • AWS Cloud Practitioner • Data Engineering ZoomCamp • Udacity AWS Data Engineering Nanodegree • Udacity Data Streaming Nanodegree

Published author: "Introduction to dbt and Analytics Engineer" (Towards Dev, Feb 2023) • "How Hive Transformed the World of Data Engineering" (Towards Dev, Apr 2024) • "Building a Football Data Pipeline" (Art of Data Engineering, Oct 2024)