

HIKMAT ULLAH

Senior Data Engineer • Snowflake • Azure Databricks • Azure Data Factory • Data Mentor

Lahore, Pakistan | +92 324 4297623 | uhikmat036@gmail.com | linkedin.com/in/hikmat-ullah-a72729180

SUMMARY

Senior Data Engineer with 7+ years of experience in SQL, Python and cloud data platforms. At Algo I build and run the backend data layer of a production supply chain intelligence platform, focusing on pipeline reliability, SLA monitoring, cost and performance optimization, and the migration of workloads from Databricks to Snowflake.

Skilled in end-to-end ETL/ELT design with Azure Data Factory, Snowflake and Azure Databricks, backed by strong SQL, Python and PySpark. Alongside my role, I have mentored 100+ aspiring data engineers one-to-one since 2021, and 80+ of them now work in data roles.

EXPERIENCE

Algo	Hybrid
Senior Data Engineer	2024 – Present
Data Engineer	2021 – 2024
<i>AlgoCore product team: backend data systems and pipeline reliability for large-scale supply chain analytics.</i>	
Data pipeline development & optimization	
- Designed and optimized end-to-end ETL/ELT pipelines for high-volume production data. - Improved the performance of core feed-processing pipelines, reducing delays and increasing reliability. - Designed and built an end-to-end Prod-to-Dev data sync pipeline. - Built automated workflows to unload data from Databricks and load it into Snowflake. - Contributed to backend enhancements of the Algo product data layer.	
Monitoring, alerts & reliability	
- Implemented an SLA monitoring system for critical pipelines using logging frameworks. - Built automated alerts for job failures, long-running pipelines and empty-file processing, with team notifications for proactive issue resolution. - Resolved production incidents, runtime dependency and authentication issues on Azure Databricks.	
Migration & platform modernization	
- Led the migration of pipelines from Databricks to Snowflake, converting notebooks and scripts and updating ADF orchestration. - Ensured data consistency and performance throughout the migration.	
Snowflake cost, performance & access	
- Monitored warehouse credits and compute utilization to identify cost optimization opportunities. - Optimized SQL transformations to cut repeated processing and improve reliability. - Managed Snowflake user creation and role-based access control.	
Data integration & quality	
Built and maintained pipelines for ingestion, transformation and validation, with data cleaning and quality checks to keep data accurate.	
Softinto	2019 – 2021
SQL / Python Developer	
- Wrote T-SQL stored procedures, views and reporting queries on SQL Server. - Built Python automation scripts for recurring data tasks. - Tuned queries and indexes to improve report and application performance.	

TRAINING & MENTORSHIP

Independent Data Engineering Mentor (1:1, online)	2021 – Present
Trained 100+ students one-to-one, from their first SQL query to their first data job; 80+ are now placed in data roles.	

- Teach 12 courses: SQL, Python, Snowflake, Azure Databricks, Azure Data Engineering, Microsoft Fabric, Power BI, Git & GitHub, Azure Fundamentals, Azure Data Fundamentals, Prompt Engineering and Generative AI for Data Engineers.
- Live sessions with assignments, real projects, quizzes and interview preparation delivered through Google Classroom.

KEY PROJECTS

- **SLA & Alerting Framework:** logging-based SLA tracking with automated alerts for failures, long runs and empty files.
- **Cloud Migration Initiative:** moved production pipelines from Databricks to Snowflake with consistent data and performance.
- **Standardized Data Pipelines:** reusable ingestion, transformation and validation patterns across feeds.
- **Prod-to-Dev Data Sync:** end-to-end pipeline keeping development environments in sync with production data.

TECHNICAL SKILLS

Data engineering & cloud: Azure Data Factory, Snowflake, Azure Databricks, ADLS Gen2, Azure Key Vault, Azure SQL, Microsoft Fabric, Redshift, SQL Server

SQL & databases: SQL, T-SQL, MS SQL Server, MySQL, PostgreSQL, Oracle SQL, stored procedures, triggers, UDFs, indexing, query optimization, data modeling, data warehousing

Programming: Python, PySpark, Pandas, NumPy

Reporting & BI: Power BI, SSRS, drill-down, drill-through and parameterized reports, dashboards

Practices: ETL/ELT design, SLA monitoring, alerting, data validation, performance tuning, cost optimization, automation

Tools: Git, GitHub, CI/CD pipelines, Jira, Draw.io, Lucidchart

CERTIFICATIONS

- Snowflake: Advanced Data Engineer
- Microsoft Certified: Fabric Data Engineer Associate
- Generative AI for Data Engineers Specialization (Coursera)
- Meta Database Engineer Specialization
- Microsoft Azure SQL · Data Storage in Microsoft SQL
- Version Control with Git (Coursera)

EDUCATION

BS Software Engineering

International Islamic University, Islamabad

References available upon request.