

Fouzan Ahmed

Data Engineer

📍 VIC 3012 ✉️ fouzanf.ahmed@gmail.com ☎️ 0451 673 122 🔗 [fouzanf-ahmed](#) 🌐 [fouzanahmed](#)

Professional Summary

Data Engineer with experience across financial services and community analytics, building the whole stack from pipelines and warehouses through to Power BI dashboards and ML models. At Principal Financial Group, I found the cause of an ingestion failure that had blocked a team for months, a one-line pagination fix that restored the pipeline. Since then I have built an automated AWS pipeline over APRA superannuation data and set up a not-for-profit's first analytics infrastructure from scratch.

Experience

UMNT Inc., Data Analyst

Darwin, NT (Remote)
June 2025 - present

- Built the organization's first automated data pipeline, bringing registration, attendance, feedback and finance data from four separate systems into one cloud database using Python and SQL.
- Created a single source of truth for 800+ members by matching student records across different classes and systems.
- Built Power BI dashboards for leadership to track engagement, retention and funding trends, cutting manual reporting time by 40%.
- Analyzed enrolment, attendance and donation data to identify growth opportunities, participation gaps and funding shortfalls.

Principal Financial Group, Data Engineer

Hyderabad, India
Jan 2023 - Feb 2024

- Found the cause of a months-old ingestion failure blocking Morningstar ratings data in non-prod. The S3 listing call was not paginated, so it stopped at 1,000 objects, and only the environment with the largest file backlog ever hit that limit. A one-line fix restored the pipeline
- Owned daily ingestion of Morningstar fund ratings data feeding the product tables used by client-facing investment rating teams across equities, fixed income and real estate
- Built AWS infrastructure as code with CDK in Python across S3 buckets, Lambda functions and Glue jobs, with assertion tests checking runtime versions, IAM permissions and encryption settings on every synthesized template
- Maintained a three-stage Lambda ingestion pattern covering pre-processing validation, transformation and post-processing quality checks, with pipeline state tracked in DynamoDB so reruns were always safe
- Rewrote Informatica mappings and SQL workflows and changed how the jobs read from S3, which cut ETL runtimes

Education

Monash University, Master of Data Science

Clayton, VIC
Mar 2024 - Mar 2026

- Relevant Coursework: Data Wrangling, Cloud Computing and Security, Foundations of Data Science, Statistical Data Modelling, Databases, Big Data Processing

Vasavi College of Engineering, Bachelor of Engineering in Information Technology

Hyderabad, India
June 2019 - May 2023

- Relevant Coursework: Deep Learning, Machine Learning, Natural Language Processing, Artificial Intelligence

Projects

APRA Superannuation Analytics Platform

2026

- Built a weekly Apache Airflow pipeline on AWS EC2 that ingests APRA regulatory filings, pushes raw files to an S3 data lake, loads cleaned data into AWS RDS PostgreSQL with idempotent upsert logic, runs dbt transformations, and sends Slack alerts on failure, with no manual steps
- Modeled 4 dbt data marts covering fund performance rankings with RANK() window functions over 1, 3 and 5 year returns, fee versus return analysis, and quarterly member flow tracking
- Connected Power BI to AWS RDS over DirectQuery for three live dashboard pages, published to Power BI Service
- Developed a natural language to SQL API in FastAPI using GPT-4, which takes a plain English question, injects the dbt mart schemas into a structured prompt, and returns generated SQL plus live results as JSON

Market Data Warehouse on Snowflake

2025

- Set up the whole Snowflake environment from scratch, covering database and schema layout, separate warehouses for loading and querying, custom engineer and analyst roles, 60-second auto-suspend, and a resource monitor to cap spend
- Loaded five years of prices across 1,000 tickers through an internal stage with COPY INTO, then built the staging and feature-engineering models on top

StreamPredict, Real-Time Delivery Prediction

2025

- Built a Kafka and Spark Structured Streaming pipeline for real-time food delivery time prediction using PySpark MLlib regression

Technical Skills

Data & Analytics: Python, SQL, pandas, Power BI, DAX, Tableau, Snowflake

Data Engineering: ETL pipelines, Apache Spark, Kafka, dbt, Informatica, PostgreSQL, DynamoDB

Cloud & Infrastructure: AWS (Lambda, S3, EC2, RDS, CDK, Glue, CloudWatch), Docker, GitHub Actions, CI/CD

AI & ML: NLP, LangChain, RAG, agentic workflows

Development: FastAPI, REST APIs, React, Git, Agile

Certifications

- AWS Certified Developer - Associate
- Snowflake SnowPro Core Associate
- Microsoft Power BI Data Analyst Associate (PL-300)
- Microsoft Certified Azure Fundamentals (AZ-900)

Leadership and Activities

Australian Bureau of Statistics (Community Field Officer)

- Followed ABS collection procedures and privacy requirements, and handled daily administrative reporting through ABS mobile applications

7-Eleven (Assistant Manager)

- Supervised and trained a team of 5+, improving operational efficiency and customer satisfaction

Achievements

- Top 4% nationally, 96th percentile in JEE Mains (1.2M+ test takers)
- Multiple Toastmasters awards for public speaking and structured debate