

ABHIJIT KUNJIRAMAN

Boston, MA | +1 (857) 250-5275 | abhijitrk97@gmail.com | [LinkedIn](#) | [Github](#) | [Portfolio](#)

Data Engineer specializing in cloud-native ingestion platforms and modern Snowflake architecture—designing replay-safe, incremental pipelines with rigorous validation, observability, and release automation across AWS and Azure.

TECHNICAL SKILLS

Languages: Python, SQL, TypeScript

Data Platforms: Snowflake, PostgreSQL (AWS RDS), MS SQL Server, Oracle, BigQuery, Redshift

Data/Orchestration: Azure Data Factory, dbt, Airflow, Spark, Kafka

Cloud/Tools: AWS, Azure, GCP, Terraform, Jenkins, Azure DevOps, Git, Linux, Liquibase, BigQuery, EC2, Lambda

BI/Monitoring: Power BI, Tableau, Grafana | Other: Postman, Confluence, Liquibase, Pydantic, sqlalchemy

Concepts/Practices: Data Modeling, ELT/ETL, CDC Patterns, CI/CD, IaC, Monitoring, APIs (REST/GraphQL), Streamlit

EXPERIENCE

ENGIE Impact - Senior Data Engineer | Boston, MA

Aug 2023 - Present

- Owned an event-driven ingestion platform on **AWS** using **S3, EventBridge, SQS, Lambda, Redis, Glue, Athena** and **Step Functions**, storing data in **PostgreSQL, Influx** processing 500M records/day to power reporting for 150+ client's dashboards
- Acted as the single-threaded owner for Product delivery (estimates, feasibility, resourcing, and execution), consolidating work previously handled by an outsourced 7-person data engineering team into one accountable engineering lead.
- Established schema governance with versioned migrations and CI deployments using Liquibase, Python and Jenkins reducing breaking-change incidents and standardized development with databases for devs.
- Built a CDC-style ingestion pattern for batch files (inserts/updates/deletes) with hash-based change detection, replacing full partition rewrites and reducing reprocessing from 10.5 hours weekly to 3 hours weekly.
- Productized FX and geography enrichment as a batch **REST** and **GraphQL API**, and it became the company-wide standard teams reused it instead of using one-off lookup queries on Snowflake in every pipeline severely improving performance.
- Implemented **Kafka** and **Spark** ingestion for event data so we could process it at scale and still keep reporting stable like schema validation, dedupe, replay-safe writes improving 30 - 50% performance.
- Authored and delivered an architecture to source directly from **Snowflake**, enabling **Matillion** retirement and \$40–50K annual savings, while reducing data availability lag refreshing data once every day to whenever needed.
- Orchestrated on-prem **MS SQL Server** extraction into **Azure Blob** with **Azure Data Factory**, creating a durable landing zone that decoupled source pulls from downstream processing and enabled repeatable replays and backfills.
- In Snowflake, implemented external stages on Azure Blob and a RAW view “contract” layer over staged files, then used **dbt** to transform into Stage, Dim, Fact models with consistent parsing/typing and backfill-friendly reprocessing.
- Partnered with DevOps on Disaster Recovery, parameterizing **Terraform** to deploy across Envs and regions.
- Replaced legacy “deploy everything” Snowflake releases with a manifest-based, change-only CI/CD pipeline (**Azure DevOps YAML + Terraform**), reducing script size from ~95K to ~1.5K–3K lowering deployment risk via repeatable, deterministic runs.
- Unblocked two long-stalled initiatives (RDS migration, Unified UOM) by driving root-cause analysis across data and models, providing technical fixes and data model changes moving delivery from ‘stuck for 12 months’ to ‘release in 2 months’.

Technocrat Services - Data Engineer | India

Sep 2019 - Sep 2021

- Migrated sales ETL from Amazon Redshift to PostgreSQL, improving performance ~50% and reducing runtime ~60% by redesigning extraction workflows
- Built Airflow, Spark, Python, SQL pipelines across S3, Athena, Glue, Redshift delivered analytics reducing wastage by ~6%

EDUCATION

Northeastern University - M.S., Data Architecture and Management | Boston, MA

Sept 2021 - May 2023

NMIMS University - B.Tech, Electronics and Telecommunication Engineering | Mumbai, India

Jun 2016 - Jul 2020

PROJECTS

Amazon Product Recommendation (NLP/MLOps): Sentiment-based recommendations with ETL and deployment workflows (GCP)

IMDB Data Integration: Built DW-style dataset and BI dashboards using SQL/ETL tools (Power BI/Tableau)