

SATISH FULWANI

Senior Data Engineer | Snowflake, Kafka, Azure & AWS | 3x AWS Certified
Pune, India · MTech Data Science & Engineering, BITS Pilani

SUMMARY

Staff Data Engineer with 8+ years building reliable, cloud-native data platforms. Seeking Senior Data Engineer roles at large investment firms. Strengths include Snowflake warehousing, realtime pipelines (Kafka, Lambda), batch ETL/ELT on Azure Data Factory and Databricks, and production GenAI/RAG on Snowflake Cortex AI and LangChain.

Led a team of 3 to modernize a legacy platform on Azure and Snowflake, migrating 1+ TB with zero downtime for a major regulated insurance program. Delivered agent skills and automations and event-driven streaming into Snowflake for analytics. Mentor engineers; partner with product and business on data quality and SLAs.

Proven delivery in regulated, high-volume environments (insurance, media, nutrition); comfortable partnering with business and compliance stakeholders on data governance, quality SLAs, and platform reliability.

Rapidly adapts to new stacks, tooling, and domains - learn, unlearn, and relearn is a working principle.

TECHNICAL SKILLS

DATA & STREAMING:

Snowflake, Kafka, PySpark, Databricks, SQL, Azure Data Factory, Pandas, data modeling

CLOUD:

AWS (Lambda, Glue, S3, API Gateway, DynamoDB), Azure (ADF, Databricks, Synapse, Container Apps), Terraform, Kubernetes, Jenkins CI/CD, Serverless

GENAI & TOOLING:

Python, LangChain, RAG, Snowflake Cortex AI, vector search, agent skills & automations, data governance & lineage

AWARDS & RECOGNITION

- Brightest Mind - Nagarro's company-wide engineering excellence award.
- 2nd place, Accenture AI Hackathon (Aug 2019) - AI solution pitched and demoeo to senior leadership.
- APEX Award x2 - Accenture's highest performance award, for Delivery and for Margin.
- Built HR mail automation system in Python, adopted company-wide by the HR team.

CERTIFICATIONS

- GCP Certified Cloud Digital Leader | GCP | Dec'22
- AWS Certified Machine Learning - Specialty | AWS | Mar'20
- AWS Certified Cloud Practitioner | AWS | Mar'20
- AWS Certified Developer - Associate | AWS | Sep'19
- Machine Learning Advanced Nanodegree | Udacity | Jan'19
- Machine Learning Foundation Nanodegree | Udacity | Apr'18

EDUCATION

BITS PILANI

MTECH DATA SCIENCE & ENGINEERING | APR'21-APR'23 | PILANI, RJ
CGPI - 7.18/10

MUMBAI UNIVERSITY

B.E. COMPUTER SCIENCE | JUN'14-JUN'18 | VESIT, MUMBAI, MH
CGPI - 7.19/10

PROFESSIONAL EXPERIENCE

Nagarro India Pvt Ltd

Feb'23 - Present | Pune, IN

Role: STAFF DATA ENGINEER

Industry: Large Consultancy Firm

Project: AWS Data Platform, GenAI Bot, Realtime Streaming & Agent Automations

Technologies: AWS, Terraform, LangChain, Kafka, Lambda, Snowflake Cortex AI, S3, API Gateway, Python

Led architecture and delivery of a Terraform-managed AWS data platform, a production GenAI bot, and realtime Kafka streaming. Mentored engineers and drove adoption of agent-tooling patterns across the team.

Key Achievements:

- Designed and provisioned end-to-end AWS infrastructure with Terraform (Lambda, S3, API Gateway, IAM policies), ensuring reproducible, secure deployments.
- Architected and deployed a production RAG chatbot on Snowflake Cortex AI; engineered LangChain-driven document parsing, chunking, and embedding into Snowflake native vector tables for semantic retrieval.
- Built a Kafka event-ingestion layer with Lambda consumer workers curating near-real-time datasets in Snowflake for analytics and downstream BI dashboards.
- Designed and published reusable agent skills and automations covering code review, documentation, and operational runbooks - measurably reducing team delivery cycle time.
- Mentored engineers on GenAI patterns and RAG design; established team-wide best practices for agent tooling and observability.

Role: ASSOCIATE STAFF DATA ENGINEER

Industry: Asia's Largest Insurance Company

Project: Azure Cloud Data Pipeline Modernization

Technologies: Azure Data Factory, Databricks, PySpark, Azure SQL, Snowflake

Led a team of 3 to modernize Asia's largest insurance company's on-premises data estate onto Azure, delivering reliable pipelines, regulatory-grade governance, and executive analytics.

Key Achievements:

- Led a team of 3 to migrate 1+ TB of legacy insurance data to Azure SQL and Snowflake with zero downtime, coordinating cutovers across business and ops stakeholders.
- Designed and delivered Azure Data Factory orchestration pipelines and Databricks / PySpark ETL workflows processing large-scale, regulated insurance data at production SLAs.
- Implemented data quality frameworks, pipeline observability, and lineage tracking across 20+ ADF pipelines to meet regulatory compliance and audit requirements.
- Partnered with solution architects and business stakeholders to translate complex insurance-domain requirements into reliable, versioned data contracts and SLA-bound delivery.
- Enforced data governance policies and access controls across Azure and Snowflake workloads; reduced pipeline incident response time through alerting and runbook automation.

Deloitte ILLP

Aug'21 - Feb'23 | Mumbai, IN

Role: DATA ENGINEER CONSULTANT

Industry: Dutch Nutrition Company

Project: Solving Complex Big Data Issues Based on Business Needs

Technologies: Spark (PySpark), Python, AWS Glue, Hadoop, AWS S3, Microsoft Dynamics (CRM)

Aligned data engineering with evolving business needs: data quality, optimized Glue/Athena/S3 lakehouse patterns, and SageMaker notebooks for ML workflows serving a global nutrition client.

Key Achievements:

- Architected and delivered ETL on AWS Data Lake (Glue, Crawlers, Athena) and ML-bound pipelines feeding multiple production models.
- Delivered near-real-time executive dashboards; enabled SageMaker ML workflows on lakehouse data.

Accenture Pvt Ltd

Oct'18 - Aug'21 | Bangalore, IN

Role: BIG DATA ENGINEER

Industry: US based Media & Entertainment Client

Project: Implemented migration of static files to AWS S3, optimizing storage and delivery efficiency

Technologies: AWS Lambda, SQS, DynamoDB, S3, EC2

Designed economical AWS migration architecture for on-premises media ingestion using serverless, fault-tolerant workflows while preserving data integrity for a US media and entertainment client.

Key Achievements:

- Architected event-driven pipelines with Lambda, SQS, and DynamoDB for resilient, scalable ingestion.
- Migrated 1.5 million media assets to S3 with ~\$100 migration overhead via serverless design.

SELECTED PROJECTS

Personal Projects

AI Travel Planning Assistant

LangChain · Azure OpenAI · ChromaDB · Streamlit · Docker

- TravelAgent orchestrator: RAG (ChromaDB) + MCP tools; Streamlit, Docker Compose, pytest.

Event Management Platform

AI-assisted development

- Event platform with AI chatbots, connections, attendance, and networking on Azure Container Apps / Kubernetes; end-to-end AI-assisted development.

Fruit Classifier (MTech)

Deep Learning · AWS/Keras | Apr'23

- MTech thesis: CNN fruit classification via transfer learning (ResNet/VGG) on AWS/Keras; compared accuracy and cost across architectures.