AKSHAT SINGH CHANDEL

Data Engineering | Azure Databricks, Unity Catalog, ADF, PySpark

Mumbai, India (Remote / Hybrid) | akshatchandel92@gmail.com | +91 8828054292 | linkedin.com/in/akshatchandel

WORK EXPERIENCE

Novigo Solutions (Remote) **Oct 2024 - Present**

Programmer Analyst (Data Engineer), promoted Jul 2026 - Present

Senior Software Engineer Oct 2024 - Jun 2026

- Migrated a Saudi retail group from Azure Synapse to Azure Databricks on a 10-developer team, decommissioning Synapse onto one Azure platform and helping cut its largest company's monthly cost from about $40K to $27K with legacy and medallion both live.
- Helped build a reusable medallion framework (bronze, silver, gold) on Unity Catalog, reused to onboard about 1,500 tables and 2-3 TB across three group companies from SQL Server, BigQuery, Dynamics 365 and Oracle.
- Owned the sales vertical's transformation layer, writing the SQL stored procedures behind its 12 GB sales table on Databricks, and worked across 20-25 fact and dimension stored procedures in the legacy and medallion builds.
- Delivered one group company's data platform end to end as sole owner, from ingestion through bronze, silver and gold, building a star schema of 11 fact and 8 dimension tables that feeds 6 Power BI reports the business previously built by hand every day.
- Built and scheduled that company's three Databricks jobs: an hourly Shipsy load, a daily Oracle load, and a file-arrival trigger that runs only the matching stored procedure when a file lands in ADLS.
- Resolved T-SQL to Databricks SQL migration defects by rewriting UPDATE ... FROM logic as MERGE INTO with source deduplication and fixing case-sensitivity mismatches left by SQL Server collation.
- Validated data parity between Databricks and Synapse before cutover so business users saw no change in their reports, and traced legacy vs medallion gaps to deleted, modified or stale source records.
- Helped replace about 900 one-pipeline-per-table ADF pipelines with a metadata-driven design of one pipeline per source (10 for medallion, about 30 total), and built the Databricks jobs it triggers for stored procedure execution and materialized view refresh.
- Built a local ADF lineage accelerator with Claude Code that turns ARM templates into a searchable HTML dashboard of pipeline dependencies, master/child links and triggers, used by the team to map about 900 legacy pipelines.
- Leading production support for the sales vertical through go-live, diagnosing and fixing pipeline, job and data issues across Azure Data Factory and Databricks.
- Built solo and co-presented to the client a multi-agent retail decision-support demo on Databricks Agent Bricks, with a supervisor routing across 2 Genie spaces, 3 Unity Catalog functions and a live XGBoost Model Serving endpoint.
- Built a voice AI service agent POC for a car-servicing client on the Databricks Mosaic AI Agent Framework, with tools that identify returning callers, onboard new ones and book service slots in Unity Catalog Delta tables; demoed to the client.
- Designed a two-agent setup to keep LLM cost low: a fast, low-cost front agent, steered by a custom routing system prompt, handled conversation and escalated data tasks to the Databricks supervisor agent.
- Led model and architecture research for the voice agent, comparing standalone STT/TTS models against Twilio ConversationRelay with Databricks Genie, and selected the final stack.
- Rebuilt a Databricks synthetic data accelerator for Snowflake in Python, Streamlit and Snowflake Cortex, generating a populated demo data model of 7-8 or 20-25 tables from a prospect's company profile for presales client demos.
- Co-built a Snowflake Cortex and Streamlit agentic AI demo, a ship-routing app over simulated fleet data that classified routes by weather, used in Novigo's demo of Snowflake agentic capabilities.

Infosys Ltd - Digital Specialist Engineer (Navi Mumbai) **Jun 2022 - Oct 2024**

- Built data management pipelines on a 6-person team for a leading pharmaceutical client's platinum-tier program, part of about 40 Azure Data Factory pipelines across ADLS and Databricks feeding downstream analytics.
- Engineered end-to-end ETL solutions on Azure, using PySpark and SQL for data aggregation and transformation behind client reporting.
- Built reusable PySpark templates that let teammates replicate common transformation tasks across project workstreams.

EDUCATION

Thakur College of Engineering and Technology, Mumbai - BE, Information Technology (GPA 9.00/10) **2018 - 2022**

SKILLS & TOOLS

Data Engineering: Azure Databricks, Apache Spark, Unity Catalog, Delta Lake, Databricks Jobs, materialized views, Azure Data Factory, ADLS, Azure Synapse, Snowflake, medallion architecture, dimensional modeling (star schema), ETL, CDC and full-load ingestion, stored procedures, data quality and validation

Languages: Python, PySpark, SQL (Spark SQL, T-SQL), JavaScript, C#

Sources & BI: SQL Server, Oracle, BigQuery, Dynamics 365, Power BI, Excel, Tableau, Qlik Sense

AI: Agentic AI, Databricks Agent Bricks, Mosaic AI Agent Framework, Databricks Genie, Databricks Apps, Unity Catalog functions, Model Serving, tool calling, prompt engineering, LLM evaluation, Snowflake Cortex, Streamlit

Certifications: Microsoft Certified: Azure Data Engineer Associate (DP-203, 2024), Azure Fundamentals (AZ-900, 2023), Rice University: Introduction to Scripting in Python Specialization (2021)