

Koushik Dhatkshinamoorthy

Senior Data Engineer

Coimbatore, India | 8825727459 | koushikdugk007@gmail.com | www.linkedin.com/in/koushikmoorthy | github.com/koushik-moorthy

PROFESSIONAL SUMMARY

Senior Data Engineer with 5+ years delivering and modernizing Azure and AWS data platforms. Builds resilient SQL, Python, and SSIS pipelines, strengthens production data operations, and turns complex datasets into Power BI and SSRS reporting solutions. Hands-on experience spans Databricks, Spark, Microsoft Fabric, Unity Catalog, HL7 FHIR, semantic models, incremental-load optimization, incident recovery, technical review, and mentoring junior engineers.

TECHNICAL SKILLS

- **Languages and processing:** SQL, Python, PySpark, Spark SQL, C#/.NET, Pandas
- **Data engineering:** ETL/ELT, SSIS, batch ingestion, incremental and watermark-based loads, stored procedures, data modeling, data-quality validation, keyed upserts
- **Azure:** Microsoft Fabric Data Pipelines, Lakehouse, Fabric Data Warehouse, Azure Data Factory, Azure Databricks, Apache Spark, Unity Catalog, Delta tables, Azure DevOps, Dynamics 365 Synapse Link
- **AWS:** Amazon ECS, AWS Lambda, Amazon S3, Amazon MSK/Kafka, Amazon SQS/SNS, AWS X-Ray
- **Databases and reporting:** SQL Server, PostgreSQL, Azure SQL Managed Instance, Amazon Redshift, Amazon DocumentDB, Amazon DynamoDB, Redis, Azure Analysis Services, Power BI, SSRS, Kibana
- **Delivery and operations:** SQL Server Agent, Flyway, Git feature branches, Development/QA/Stage promotion, Visual Studio package deployment

PROFESSIONAL EXPERIENCE

Presidio Solutions Private Limited

Senior Engineer

Data Engineer

Associate Engineer

Software Engineer Intern

Coimbatore, India

September 2023 - Present

November 2021 - September 2023

July 2021 - November 2021

September 2020 - July 2021

Pilot Association Operations and Analytics Data Platform | SQL Database Administrator assignment

May 2024 - Present

- Redesigned incremental SSIS extraction using a previously undocumented source timestamp, substantially reducing processing volume and runtime while preserving data completeness through reconciliation.
- Engineered and deployed production ingestion enhancements for new source tables across PostgreSQL, FTP, CRM, and frequently updated data sources by extending SSIS packages, database schemas, stored procedures, and SQL Server Agent schedules.
- Strengthened platform resilience by diagnosing schema, datatype, network, and source-data failures and executing controlled, upsert-based recovery without creating duplicate records.
- Drove recovery of delayed or missing airline FTP feeds through direct source coordination, controlled reprocessing, and clear restoration updates for downstream stakeholders.
- Modernized reporting delivery by migrating a business team's complete paginated-report suite from SSRS to Power BI, including schedules and user-level access configuration.
- Delivered operational reports and automated notifications across multiple internal teams, and raised delivery quality through peer review of SQL and reporting changes before promotion.

Microsoft Fabric Modernization Pilot | Proof of concept

May 2025 - August 2025

- Co-designed a Microsoft Fabric medallion architecture with data architects and delivered the complete proof-of-concept flow from airline PostgreSQL ingestion to business reporting.
- Engineered watermark-driven Fabric Data Pipelines, dedicated Bronze and Silver Lakehouses, PySpark and Spark SQL transformations, and Gold fact and dimension models in Fabric Data Warehouse.
- Demonstrated the modernized reporting flow to stakeholders while clearly retaining the initiative as a proof of concept rather than a production deployment.

Modern Data Platform Modernization | Databricks proof of concept

October 2023 - April 2024

- Delivered Spark and Python transformation components for Microsoft Dynamics 365 and Bloomberg source flows, including Unity Catalog definitions and governed Delta table outputs.
- Standardized Databricks job packaging with Asset Bundles and established an Azure DevOps CI/CD pipeline for repeatable source and bundle deployment; the initiative concluded at the proof-of-concept stage.

Insights | Enterprise Food-Industry Analytics Platform

March 2023 - October 2023

- Stabilized and enhanced an inherited Azure Data Factory estate spanning more than 20 source systems, Azure SQL Managed Instance, Azure Analysis Services models, and Power BI reporting.
- Delivered end-to-end onboarding of a SQL-based source across ADF ingestion, SQL transformation and loading, and analytical cube integration.
- Engineered complex stored procedures and optimized SQL queries powering business transformations and reporting workloads.

Microsoft Sentinel Security Analytics in Power BI

January 2023

- Transformed raw Microsoft Sentinel data into curated analytical datasets and interactive Power BI security reports.

Healthcare Data and Document Workflows

February 2022 - January 2023

- **Healthcare FHIR Data Processing Platform:** Engineered Python bulk-load scripts for CSV-to-DocumentDB ingestion and .NET components that applied supplied mapping logic to produce HL7 FHIR outputs in Amazon S3.
- Extended AWS Lambda consumer and ECS task logic across an Amazon MSK/Kafka, Redis, DynamoDB, S3, SQS, and SNS workflow; improved operational visibility through monitoring dashboards while the wider team retained architecture and mapping-rule ownership.
- **EHR document extraction and normalization:** Implemented API and document-processing components for pre-signed S3 transfers, ZIP extraction, MIME detection, Base64 conversion, integrity validation, and multi-format handling.
- **EHR CDC-to-FHIR transformation:** Developed transformation queries and mapping implementations that converted EHR change-data-capture records into FHIR-formatted data.

Healthcare Enterprise Data Hub

May 2021 - February 2022

- Optimized Redshift ingestion through targeted transformation queries, reusable UDFs for varied data formats, and bulk-load consistency validation.
- Built AWS Data Wrangler and Pandas workflows to read S3 CSV files, transform records, and load Amazon Redshift.
- Improved application observability by implementing AWS X-Ray tracing and Kibana dashboards for real-time monitoring.

EDUCATION

Bachelor of Technology (B.Tech) in Information Technology

Kongu Engineering College | 2017 - 2021 | CGPA: 8.1/10

CERTIFICATIONS

- Microsoft Certified: Fabric Data Engineer Associate
- Microsoft Certified: Azure AI Engineer Associate
- NVIDIA Certified Associate: Generative AI LLMs
- Microsoft Certified: Power BI Data Analyst Associate
- Databricks Certified Data Engineer Associate
- Microsoft Certified: Azure Data Engineer Associate
- HashiCorp Certified: Terraform Associate
- AWS Certified Solutions Architect - Associate
- AWS Certified Cloud Practitioner
- Microsoft Certified: Azure Fundamentals