NIRMAL KUMAR P

Data Engineer | Microsoft Fabric & Databricks | SQL, PySpark, Power BI

Chennai, India | +91 75988 44460 | nirmalkcse@gmail.com | linkedin.com/in/nirmalkumar6 | github.com/nirmalkcse | IEEE: ieeexplore.ieee.org/document/10544873

PROFESSIONAL SUMMARY

Data Engineer specializing in Microsoft Fabric and Databricks lakehouse delivery with SQL, PySpark, CDC/SCD2, validation, BI migration, and Power BI.
Impact: 3 telecom source systems, 18-24 daily incremental tables, 15+ reports validated, 90% fewer discrepancies, 6 Tableau reports migrated, and a Databricks-hosted MDM/entity-resolution product.

TECHNICAL SKILLS

Platforms & Storage: Microsoft Fabric (Lakehouse, Warehouse, Data Factory pipelines, OneLake), Databricks (Apps, SQL Warehouse, Jobs API, Secrets API, Unity Catalog), Delta Lake, Power BI

Data Engineering: SQL, Spark SQL, PySpark, Metadata-driven Ingestion, Medallion Architecture, Incremental/CDC pipelines, SCD Type 2, Schema Enforcement, Data Quality

Analytics & BI: Semantic Modeling, DAX, Power Query, Power BI Report Builder, KPI modeling, Tableau-to-Power BI migration, MicroStrategy validation

Application & Workflow: TypeScript, Next.js, React, Zod, TanStack Query, REST API routes, Git/GitHub, GitHub Copilot, Copilot Studio, SSMS

PROFESSIONAL EXPERIENCE

iLink Systems - Software Engineer Trainee, Data Engineering & Product Development | Sep 2024 - Present

iCORE Enterprise MDM / Entity Resolution Platform | Jan 2026 - Present

- Own end-to-end development of a Databricks-hosted MDM/entity-resolution product using Next.js 15, TypeScript, Databricks Apps, SQL Warehouse, Jobs/Secrets APIs, Unity Catalog, and Delta control tables.
- Built a Master Entity model spanning connectors, source mappings, entity-scoped pipeline launch, review queues, golden records, resolution templates, and governance workflows.
- Implemented Databricks APIs for SQL execution, job submission, catalog access, credentials, RBAC, audit logging, and pipeline observability.
- Developed stewardship features for match explanations, score breakdowns, triage workflows, approval/SLA tracking, and evidence-pack foundations.

Tier-1 European Telecom Provider - Fabric Data Migration PoC | Oct 2025 - Dec 2025

- Engineered metadata-driven Fabric ingestion for 3 source systems - GRASP SFTP, AMDOCS Oracle, and AIDA SFTP - supporting 18-24 daily incremental source tables.
- Built Bronze/Silver/Gold processing with Fabric pipelines, PySpark, SQL notebooks, Delta tables, hash-based CDC, and SCD Type 2 historization.
- Processed Gzip text/CSV, pipe-delimited CSV, and Binary JSON sources with schema correction, JSON flattening, UTF-8 standardization, and table normalization.
- Refactored Teradata/BTEQ logic into Fabric SQL/PySpark and created SA/SM transformations for Gross Adds, Churn, and telecom reporting.

Data Engineering, BI Migration & Analytics Delivery | Sep 2024 - Sep 2025

- Validated 15+ MicroStrategy-to-Fabric reports with SQL regression checks, Gold-to-BI comparisons, and Report Builder validation; reduced discrepancies by 90%.
- Migrated 6 Tableau reports to Power BI, rebuilding semantic models, relationships, hierarchies, DAX measures, and layouts while preserving business logic.
- Built 2 Fabric/Power BI PoCs, including an insurance data-agent solution and a FabCon25 Europe 2025 demo recognized by senior company and Microsoft leadership.

CERTIFICATIONS

Microsoft Certified: Fabric Data Engineer Associate (DP-700) | Microsoft Certified: Fabric Analytics Engineer Associate (DP-600) | Databricks Certified Data Engineer Associate | Databricks UC Upgrade Delivery Specialization | Microsoft GitHub Copilot

EDUCATION & RESEARCH

Bachelor of Engineering in Computer Science and Engineering, Sathyabama Institute of Science and Technology, 2024 | CGPA: 9.11/10
IEEE research article: Analysis of Semiconductor Chip Performance using Medallion Architecture; applied Microsoft Fabric and Power BI to Bronze/Silver/Gold analytics design.