Rafael Mendes

Senior SQL Developer

Senior SQL Developer with 9 years of experience at JPMorgan Chase and Databricks, architecting high-performance T-SQL and Snowflake data pipelines for capital markets and analytics platforms. Optimized a 2.8 TB derivatives pricing database achieving 91% query latency reduction, led a team of 5 developers, and hold Microsoft Azure Data Engineer Associate and Databricks Certified Data Engineer Professional credentials.

rafael.mendes@email.com | (555) 916-4287 | New York, NY | linkedin.com/in/johndoe

EXPERIENCE

JPMorgan Chase — New York, NY

Senior SQL Developer — March 2021 - Present

- Architect and maintain T-SQL data pipelines supporting the derivatives pricing system, handling 18M trade records daily with sub-second query SLA for 500 front-office traders
- Achieved 91% reduction in P95 query latency on 2.8 TB SQL Server database through columnstore index redesign, statistics maintenance automation, and TempDB optimization
- Lead team of 5 SQL developers; conduct weekly code reviews, defined department-wide T-SQL coding standards reducing production incidents by 40%
- Designed partitioned tables and sliding-window partition scheme for 7-year trade history, reducing month-end archival time from 6 hours to 22 minutes
- Collaborated with quant team to migrate 60 legacy stored procedures from SQL Server to Snowflake, enabling self-service SQL analytics for 200 research analysts

Databricks — San Francisco, CA

SQL Developer III — August 2017 - February 2021

- Developed Databricks SQL transformation layers for customer analytics platform processing 5B events/day; reduced pipeline run time by 55% through partition pruning and Z-ordering
- Built dbt project with 180 models, 400+ tests, and automated CI/CD deployment via GitHub Actions, cutting data freshness SLA from 4 hours to 45 minutes
- Wrote Delta Lake MERGE statements replacing slow UPSERT patterns; reduced daily ETL cost by $38K (22%) on 200-node cluster

Citigroup — New York, NY

SQL Developer II — June 2015 - July 2017

- Maintained Oracle 12c and SQL Server 2014 databases for retail banking reporting; processed 2M accounts across 8 product lines
- Optimized 45 legacy PL/SQL packages, reducing batch window from 9 hours to 3.5 hours

PROJECTS

Lead Developer

Database Optimization Project — 2020-2021

Optimized database performance for a large e-commerce application

SQL Server | Azure | Python

CERTIFICATES

Microsoft Certified: Azure Data Engineer Associate (DP-203) — September 2023

Microsoft

Validates expertise in designing and implementing data solutions on Microsoft Azure, including data pipelines and SQL analytics workloads.

Databricks Certified Data Engineer Professional — April 2022

Databricks

Advanced certification validating ability to build, optimize, and secure data pipelines using Apache Spark and Databricks.

dbt Certified Developer — November 2020

dbt Labs

Validates expertise in building, testing, and deploying dbt transformation models in production cloud data warehouse environments.

SKILLS

T-SQL / SQL Server 2022
Snowflake SQL / Delta Lake
dbt (Advanced)
Databricks SQL
Query Optimization & Execution Plans
Columnstore & Partition Design
PL/SQL / Oracle Database
Python (PySpark, pandas)

EDUCATION

Rutgers University

New Brunswick, NJ

B.S.

May 2015

Computer Engineering

GPA: 3.7

LANGUAGES

English ●●●●●