

RENAN DE MORAES

Analytics Engineer | Data Engineer | AI Automation Developer

Contact: - Porto Feliz, São Paulo, Brazil -+55 (15) 99136-7797 -
renan.de.moraes777@gmail.com

Linkedin Profile - <https://www.linkedin.com/in/renan-moraes-data-ai-engineer/>

Github Profile - <https://github.com/mpraes>

PROFESSIONAL SUMMARY

Senior Data Engineer with over 5 years of experience in data architecture and a strong focus on Microsoft Fabric (Certified) and Azure Data Engineering. Proven track record in designing and deploying AI-ready data foundations using Medallion architectures (Lakehouse). Expert in optimizing Fabric Capacity (F2-SKUs), implementing robust CI/CD pipelines, and ensuring data integrity for LLM-based automation. Experienced in Python, PySpark, and SQL, with a unique background in logistics and procurement that drives cost-aware and business-centric engineering decisions.

KEY EXPERTISE: MICROSOFT FABRIC & AI DATA ENGINEERING

- Fabric Ecosystem: Hands-on delivery of Lakehouse designs, OneLake integration, and performance tuning for Spark Notebooks.
- AI-Ready Infrastructure: Building curated data layers specifically for Agentic AI and automation consumption, ensuring high-fidelity data lineage and schema enforcement (Pydantic).
- Cost Management: Strategic implementation of Fabric capacities, prioritizing high-efficiency Spark jobs over Dataflows to minimize CU (Capacity Unit) consumption.
- DevOps & Governance: Version control with Git, automated deployment using Fabric Deployment Pipelines (Dev/Prod), and enterprise-grade data security.

PROFESSIONAL EXPERIENCE

SCORAS | AI Automation Developer & Data Engineering Specialist *August 2024 – Present*

- **Agentic AI Architecture:** Designed a full-stack AI ecosystem for a major transportation conglomerate using LangGraph, orchestrating complex workflows (ticketing/password resets) via REST APIs and Databricks SQL Execution API.
- **AI Chatbot Refactoring (Text-to-SQL):** Optimized a Real Estate WhatsApp chatbot by implementing a LangChain strategy that converts natural language into precise PostgreSQL queries using system-prompted views, delivering real-time property data to users.
- **Cloud Migration (Azure to AWS):** Led the migration of a backend and PostgreSQL infrastructure from Azure Container Apps to AWS, ensuring service continuity and performance optimization.
- **NoSQL to SQL ETL (AWS Glue):** Developed ETL pipelines using AWS Glue and Python to normalize DynamoDB data into Aurora PostgreSQL, managed via Terraform CI/CD and strict Pull Request governance.

NINECON | Senior Data Engineer *January 2024 – Present*

- **Microsoft Fabric Implementation (Greenfield):** Architected a complete Microsoft Fabric Lakehouse (Medallion architecture), managing 30+ production tables. Solved complex connectivity gaps by configuring dedicated VMs and On-premises Data Gateways to integrate legacy Oracle EBS data into the Fabric ecosystem.
- **High-Volume Data Migration (SAP HANA to Oracle):** Developed a custom Python migration framework utilizing multiprocessing and threading to automate the transfer of massive datasets (up to 100M+ records per table) from SAP HANA to Oracle, ensuring high throughput and data consistency.
- **Automated Monitoring Dashboards:** Engineered a reporting tool in Python that generates dynamic HTML dashboards for weekly stakeholder updates, providing visibility into migration logs, volume metrics, and system performance.
- **Legacy Data Warehouse Development:** Designed and built a comprehensive Oracle Data Warehouse from the ground up (70+ tables). Developed complex ETL logic using IBM DataStage, advanced SQL (MERGE), and Views, replacing legacy processes with a structured Kimball-based architecture.
- **Capacity Optimization:** Optimized Spark workloads to maintain production stability within Microsoft Fabric F2 Capacity, prioritizing cost-effective engineering by replacing resource-heavy Dataflows with high-performance PySpark notebooks.

BLUER TECNOLOGIA | Senior Data Engineer & Tech Lead *February 2024 – Present*

- **Fabric Strategic Migration (Tech Lead):** Spearheaded the architectural transition of SAP ECC flows to Microsoft Fabric. Established standardized framework for Bronze and Silver layers (Notebooks/Pipelines) that became the client's engineering standard.
- **Real-Time Data Engineering:** Engineered a high-frequency ingestion pipeline using Python and Jenkins to consume vehicle tracking APIs every 5 seconds, storing data in SQL Server via optimized Upsert/Merge logic.

- SAP RFC Integration: Developed robust ETL processes using pyrfc to extract procurement data from SAP RFC directly into Delta Tables, using metadata-driven notebooks for error tracking and automated alerts (Teams/Outlook).
- Advanced RPA & Document Engineering: Architected Python RPA solutions to automate data retrieval from public agency APIs (SIGOR) and SFTP servers, delivering automated HTML status reports.

AB INBEV (via Peapply Consulting) | Senior Analytics Engineer *March 2024 – September 2024*

- Performance Optimization for Big Data: Refactored and optimized high-volume data models within Databricks (PySpark & SQL), processing tens of millions of records regarding competitor pricing and revenue metrics for the global pricing sector.
- CI/CD Pipeline Implementation: Established automated deployment workflows to promote Databricks notebooks and data assets from development to production environments, ensuring code stability and version control.
- Complex Data Modeling: Engineered robust data architectures to support stylized, high-performance Power BI dashboards, utilizing Azure Analysis Services (AAS) for semantic layer modeling and granular pricing analysis.
- Revenue Intelligence: Transformed raw competitor data into actionable insights, reducing notebook execution time and improving the refresh rate of critical revenue management reports.

IBM (Client: Vale) | Data & BI Engineer *February 2022 – February 2024*

- Enterprise Asset Monitoring: Engineered data solutions for a large-scale corporate asset monitoring portal, tracking real-time metrics (OEE, availability) for mining equipment such as trucks, locomotives, and conveyor belts.
- Semantic Layer Architecture: Developed and managed complex semantic models in Azure Analysis Services (AAS), implementing advanced DAX calculations, relationship optimization, and data partitioning to ensure high-performance reporting.
- Cloud ETL Orchestration: Designed and maintained Azure Data Factory (ADF) pipelines to ingest and process data from diverse sources, including SAP HANA, Oracle, SQL Server, and unstructured files, into a Synapse Analytics Data Warehouse.
- Dimensional Modeling: Architected robust Star Schema and Snowflake structures, ensuring data integrity and scalability for global maintenance KPIs.
- Agile Leadership: Led internal improvement projects within an agile framework using Azure DevOps, bridging the gap between raw data engineering and executive BI layouts.

GRUPO PETRÓPOLIS | BI Analyst

December 2020 - February 2022 - Developed data flows with RPA (Python) and dashboards for Logistics department - Created and maintained PL/SQL procedures, views, and packages for generating critical distribution information - Developed "Logistics Daily Bulletin" in Python (HTML/CSS) using Numpy, Pandas, and Dash libraries - Implemented ETL/RPA processes for

automated SAP integration - Supported over 250 resellers and 8 factories with solutions that prevented million-dollar losses

INTERNATIONAL FLAVORS & FRAGRANCES

October 2016 - July 2020

Raw Material Planner (October 2019 - July 2020) - Daily MRP analysis for national and international raw material purchasing - Strategic purchasing management and follow up

IM&S Buyer (October 2016 - September 2019) - Buyer of indirect materials and services for Frutarom do Brasil - Purchasing management for three branches in São Paulo State

LORENZON MANUTENÇÃO INDUSTRIAL | Buyer

December 2008 - October 2016 - Responsible for strategic parts purchasing that directly impacted monthly billing goals - Complete warehouse control: internal customer service, position mapping, inventory participation - Direct report to Commercial Manager

EDUCATION

FACULDADE PORTO DAS MONÇÕES

MBA in Logistics Engineering | 2014 - 2017

FACULDADE PORTO DAS MONÇÕES

Bachelor's Degree in Logistics | 2011 - 2013

ALURA

Data Analytics Training (300+ hours) | 2021 - 2022

SENAC SOROCABA

Purchasing Negotiation Course | 2017

CERTIFICATIONS

Cloud Certifications:

- Microsoft Azure Fundamentals (AZ-900)
- Microsoft Azure Data Fundamentals (DP-900)
- Microsoft Fabric Data Engineering Associate (DP-700)
- Oracle Cloud Infrastructure Associate
- Oracle AI Agent Associate

Technical Certifications:

There is a lot of diverse certifications what can't put here, since the resume would be too big. Check on my linkedin profile.

TECHNICAL SKILLS

- **Platforms:** Microsoft Fabric, Databricks, Azure Data Factory, AWS Glue, IBM DataStage, Jenkins, Docker.
- **Languages:** Python (Advanced), SQL, PySpark, TypeScript.
- **Databases:** Oracle, SAP HANA/ECC, PostgreSQL, DynamoDB, SQL Server, Aurora.

LANGUAGES

- Portuguese: Native
- English: Full Professional Proficiency
- Spanish: Professional Working Proficiency

FEATURED PROJECTS

QuickELT - Data Engineering Template

GitHub: <https://github.com/mpraes/quickelt> | Stars: 33 | Forks: 5

Open source template for data engineering projects using functional programming as the main paradigm. Project created to accelerate data pipeline development with best practices from the start.

Technologies: Python, Pandas, Polars, DuckDB, AWS S3, Azure, SQLAlchemy, PostgreSQL, MySQL, Oracle, Pytest

Key Achievements: - Development of reusable templates for data ingestion from multiple sources - Implementation of modular architecture with bronze/silver/gold separation - Creation of AWS S3 utilities and database integration tools - Automated testing structure and bilingual logger - 21 active commits with features for web scraping, APIs, and SharePoint

Gás Mais Barato - LPG Price Intelligence Platform

GitHub: https://github.com/mpraes/webapp_gas_mais_barato |

Deploy: gasmaisbarato.site

High-performance web dashboard that processes official government data (ANP) to compare cooking gas prices across Brazil. The application functions as a Business Intelligence tool, calculating real-time market metrics to highlight consumer savings potential.

Stack Technologies: Next.js 14.2 (App Router), TypeScript 5.4 (Strict Mode), Tailwind CSS

DeepSQL.pro site - Learning SQL free directly from projects

Learn SQL the way professionals use it: solving real analytical challenges. Work with authentic business datasets, uncover insights, and improve your skills with instant feedback, guided hints, and immersive scenarios inspired by real corporate challenges.