

Sergio Brito

Agentic AI Solutions Architect | LLM & Agent Workflows · Data Engineering · Client-Facing Delivery | n8n / Make · Python · BigQuery | 11+ Years

Campo Grande, Brazil (Remote, US/LATAM hours) · sergio.ssbfgmail.com · +55 67 98155-8888 · sergio-brito.dev · linkedin.com/in/sergio-brito145

PROFESSIONAL SUMMARY

Solutions-minded engineer with 11+ years across software, data, and applied AI, who sits naturally at the intersection of client strategy, data engineering, and agentic systems. I lead discovery conversations, frame ROI for non-technical stakeholders, and then roll up my sleeves to design and ship end-to-end agentic architectures: data sources → automation / LLM → actions and dashboards. Heavy hands-on background in both low-code orchestration (n8n, Make) and custom-code agent systems (multi-agent orchestration, RAG, multi-LLM routing) — so I meet clients where they are and know exactly when to graduate from a no-code workflow to production code. Consultancy and client-services experience; daily Python, BigQuery, and AWS/GCP.

CORE SKILLS

Agentic AI & Automation: Agent orchestration, tool calling, RAG, multi-LLM routing, prompt engineering; n8n (extensive), Make, Flowise, custom orchestration; LLM apps for enrichment, content generation, segmentation

Data Engineering: BigQuery (extensive), SQL, data modeling, data warehousing; ETL with Apache Airflow; data migration at scale; pipeline design, validation, and data-quality checks

AI / ML & LLM: OpenAI (GPT, Vision, Whisper), Claude, Gemini, OpenRouter, LiteLLM; RAG + pgvector; embeddings, structured outputs, evaluation harnesses

Client-Facing: Discovery sessions, ROI framing, stakeholder management, technical advising, solution documentation; consultancy / client-services delivery

Backend / Languages: Python (scripting, data, ML deployment), Ruby on Rails, Node.js / TypeScript, REST APIs

Cloud / Ops: AWS (S3, RDS, ECS, Lambda, Glue-adjacent ETL), GCP (BigQuery), DigitalOcean, Docker, CI/CD, orchestration patterns (Airflow / workflow engines)

PROFESSIONAL EXPERIENCE

MONUMENT — August 2025 to May 2026

Senior Data Migration Engineer (US Client, Contract)

US self-storage SaaS startup. Client-facing data migration across dozens of enterprise customers, hundreds of facilities, thousands of units.

- Owned end-to-end migration of messy, heterogeneous data into a unified schema; financial calculations reconciled to the cent before cutover.
- Worked directly with client stakeholders to scope each migration, manage expectations, and unblock data discrepancies.
- Claude Code as primary development driver for ingestion, transformation, and validation tooling.

Stack: Python, TypeScript, Node.js, MySQL, AWS, Claude Code

VIK — May 2023 to July 2025

Principal Software Engineer / AI Platform Lead

Enterprise B2B healthtech. Designed the company's AI platform from zero and the analytics layer on top of it.

- Designed AI platform from zero: agents, function calling, RAG, multi-LLM orchestration via OpenRouter and LiteLLM.
- Built BigQuery + Looker analytics layer; interactive dashboards on Rails + React.
- Shipped 'Vik Nutri' conversational AI to 1,000+ users (tool use, structured outputs, per-user memory).
- Led company-wide adoption of AI-augmented development; defined prompting and review conventions.

Stack: Python, Ruby on Rails, React, OpenAI, Claude, Gemini, OpenRouter, LiteLLM, BigQuery, Looker, AWS

META | YOUSE (Caixa Digital Insurance) — November 2022 to May 2023

Senior Software Engineer (Consultant)

Senior consultant at Brazil's largest digital insurer. Client-facing delivery within a 50+ microservices event-driven architecture.

- Led legacy data migration (CSV/JSON to unified schemas) critical for risk assessment.
- Designed Node.js/Ruby microservices; built React analytics dashboards; AWS deployments with CI/CD.

Stack: Node.js, Ruby, React, AWS, Kafka, data migration

ADMINER — September 2021 to November 2022

Software Engineer

E-commerce analytics platform at scale (millions of requests/day).

- Analytics platform with async processing; complex aggregations across PostgreSQL, MongoDB, Elasticsearch.
- 'Adminer Analytics' Chrome extension (50,000+ downloads); CI/CD reducing deploy errors 90%.

Stack: Node.js, Ruby, Angular, PostgreSQL, MongoDB, Elasticsearch

MJV | BRADESCO INSURANCE — November 2020 to November 2021

Software Developer (Data / Backend)

Critical async batch and data-integration systems for one of Brazil's largest banks.

- Built ETL pipelines with Apache Airflow for cross-system data integration.
- Async batch processing up to 1M records/day; PostgreSQL tuning cut query time 60%.

Stack: Java, Spring Boot, Apache Airflow, PostgreSQL

PSG | DETRAN (Government) — January 2017 to November 2020

Software Developer

Full lifecycle of a high-availability government data system at multi-million-record scale.

- Owned a high-availability infractions system integrated with multiple government APIs; 99.9% uptime.
- Relational data modeling and internal tooling in PostgreSQL.

Stack: Java, Spring Boot, PostgreSQL, Heroku

SELECTED AGENTIC & AUTOMATION WORK (built solo)

Production agent systems: multi-agent orchestrators with tool calling, tier-based safety, multi-LLM routing, and eval/observability layers (Rails 8 + Claude/GPT/Gemini). **Multimodal:** a sports-analytics pipeline (FFmpeg + Whisper + Gemini Vision) running real-time tool calls over WebSocket. **Low-code era:** before moving to custom code, built extensive agentic and automation workflows on n8n and Make (lead routing, content generation, data sync, notifications) for real businesses — including ongoing Instagram/ads automation for a family bakery. Details: sergio-brito.dev

EDUCATION & LANGUAGES

B.Sc. Software Engineering — Federal University of Mato Grosso do Sul (UFMS)

Languages: Portuguese (Native) · English (C1 Advanced) · Spanish (Advanced)