

Rodrigo Afonso dos Santos

Backend & AI Engineer

Email: rodrigo@rasantos.dev Web: rasantos.dev LinkedIn: linkedin.com/in/rasantos93 GitHub: github.com/radsdev93

X: [@rasantosdev93](https://twitter.com/rasantosdev93) | Remote (Brazil, UTC-3) • Open to English-speaking markets

SUMMARY

Backend & AI Engineer with 8+ years designing and scaling distributed systems, high-traffic microservices, and AI-powered automation. Deep expertise in Node.js/TypeScript, event-driven pipelines (AWS SQS/SNS), search/caching (Elasticsearch, Redis), and cloud (AWS, Azure AKS). Proven impact: avg latency ~600 ms → 194 ms (-68%), p99 3.6 s → 446 ms (-88%), 5xx 4% → 0.1% (-97.5%), and up to -80% LLM API cost via caching + prompt strategy.

CORE COMPETENCIES

Languages & Frameworks: Node.js (22+), TypeScript/JavaScript, Express (5), BullMQ; Go (collab), PHP/Laravel (collab)

Cloud & DevOps: AWS (S3, SQS, SNS), Azure (AKS), Docker, Kubernetes, GitOps (FluxCD), NGINX Ingress, CI/CD

Data & Search: MySQL 8, PostgreSQL, Elasticsearch 7.x, Redis

Messaging & Pipelines: Event-driven architectures, queue workers, streaming validators

Observability & QA: Prometheus, Grafana, Elastic APM, New Relic, structured logging, OpenAPI/Swagger

AI/LLM: LangChain, OpenAI API, RAG pipelines, prompt engineering, inference/cost optimization

PROFESSIONAL EXPERIENCE

Backend Engineer — Enterprise Omnichannel Platform (Linx)

Feb 2022 – Present | Remote (Brazil)

- Owned and evolved OMS microservices (freight quotation, routing logic, custom freight rules) within SCRUM squads.
- Freight Quotation CRUD microservice:** carved out the CRUD from legacy business rules; 13 entities, 40+ REST endpoints; Node.js 22+, Express 5, MySQL 8, Redis, Elasticsearch, AWS S3/SQS/SNS; automated OpenAPI 3.0 docs.
- Performance program:** Node clustering + dependency slimming (replacing heavy libs with native/built-ins) guided by CPU profiling and autocannon → avg latency ~600 ms → 194 ms (-68%), p99 3.6 s → 446 ms (-88%), 5xx 4% → 0.1% (-97.5%); throughput uplift on hot paths (~+40%).
- Designed event-driven pipelines (SQS workers, SNS notifications) and S3 streaming validators for bulk operations; added Elastic/Redis caching strategies and selective invalidation.
- Cross-team contributor:** PR reviews and features across Go-based Freight Cost API (log/trace refactor, minute-level quotation granularity), PHP Inventory Management API (VTEX stock-query endpoint, Silex→Laravel migration support), and Azure AKS platform (production deployments with runbooks, rollback planning, observability via FluxCD/Prometheus/Grafana).

Freelance Developer — Self-Employed

2016 – Present | Remote

- B2B backend & AI consulting focused on distributed systems, data APIs, and intelligent automation.
- Brazilian Companies Data API:** multi-DB (PostgreSQL/MySQL) with Redis caching and WHOIS/Hunter enrichment → ~95% fewer repeat queries; sub-second cached responses.
- AI-Powered Legal Document Automation:** LangChain + OpenAI RAG pipeline with multi-stage prompts → up to -80% API cost; ~10× faster preparation.

- **WhatsApp → AI Assistant Gateway:** event-driven gateway (BullMQ/Redis/Auth0) → ~95% auth overhead; stateless, horizontally scalable.
- **Customer Reputation Insights Platform:** AI-driven text classification/sentiment for complaint data → ~90% less manual analysis; 10k+ documents processed.

Backend Developer — Screencorp

Nov 2020 – Feb 2022 | Remote (São Paulo, Brazil)

- Led features for a corporate communications/TV CMS platform; owned production triage and delivery.
- Extracted push notifications into a dedicated Node.js microservice with queueing to improve stability and throughput; implemented WebSocket real-time features.

Technical Lead — Global MKT (Contract)

Aug 2016 – Sep 2019 | Hybrid

- Led a 4–6 developer team building e-commerce/marketing solutions (LP generator with CRM/payment integrations).
- Managed servers, DNS/CDN, SSL, and analytics; replaced paid tools with OSS → approx. **US\$56k/year** cost savings.

SELECTED PROJECTS & CASE STUDIES

- **Freight Quotation CRUD Service (Enterprise OMS, contributor):** latency –68%, p99 –88%, 5xx –97.5%; 13 entities, 40+ endpoints; async pipelines; OpenAPI.
- **Brazilian Companies Data API (owner):** multi-DB, Redis caching, enrichment; ~95% cache hit; sub-second responses.
- **AI Legal Document Automation (owner):** RAG + prompt system; –80% cost; 10× speed.
- **WhatsApp → AI Assistant Gateway (owner):** –95% auth overhead; event-driven, resilient queues.
- **Customer Reputation Insights Platform (owner):** +90% analysis efficiency; 10k+ docs.

CERTIFICATIONS

LPI LE-1: Linux Essentials — Issued Nov 2021

LANGUAGES

Portuguese: Native/Bilingual **English:** Full Professional Proficiency