

Gilmar Vitor Manoel da Silva

I build end-to-end digital products — web, mobile and applied AI in production.

✉ contato@gilmar.dev

☎ +55 81 98661-7887

🌐 @gilmarvitor

🌐 in/gilmarvitor



Scan for WhatsApp

PROFESSIONAL SUMMARY

Fullstack software engineer with 19+ years of experience designing, building and shipping web, mobile and AI platforms to production. Applied-AI specialist — RAG, autonomous agents and vector databases — with a track record of real-time systems serving thousands of concurrent users. I work end to end: architecture, backend, frontend, mobile, cloud and DevOps.

19+

years building software

100+

projects delivered

12

programming languages

9+

industries served

CORE COMPETENCIES

- ▶ Applied AI engineering (RAG, agents, embeddings)
- ▶ Scalable systems architecture
- ▶ Real-time applications (WebSockets, streaming)
- ▶ Native and cross-platform mobile
- ▶ Cloud & DevOps (IaC, containers, CI/CD)
- ▶ IoT and industrial telemetry

AI SPECIALTY

I design and ship systems with RAG (retrieval-augmented generation), autonomous agents, document ingestion and structuring, embeddings and vector databases, automations and real-time voice interfaces — always focused on reliability, cost and scale. I've delivered AI platforms used daily by thousands of people.

RAG · Autonomous agents · Embeddings · Vector databases · Fine-tuning · LLM Ops · Real-time voice · Smart automations

EXPERIENCE & PROJECTS

Live streaming with AI · Platforms · AI

Live events and education platform scaling to thousands of concurrent viewers, with real-time chat, polls and assessments.

WebSockets · Realtime · Python

Native and cross-platform apps · Mobile

Apps published on both stores — native in Kotlin and Swift and cross-platform with React Native/Expo and Flutter.

Kotlin · Swift · Expo · Flutter

Industrial telemetry · IoT · Industry

Data collection and telemetry over LoRaWAN and MQTT, integration with PLCs and Modbus and real-time operations dashboards.

LoRaWAN · MQTT · Modbus

Multi-tenant platforms · SaaS · ERP/CRM

Multi-tenant management systems for regulated sectors, with APIs, billing and end-to-end integrations.

[Laravel](#) · [Multi-tenant](#) · [APIs](#)

Ingestion and RAG · AI · Documents

Document ingestion and structuring pipelines with embeddings and semantic search for reliable answers.

[RAG](#) · [Embeddings](#) · [Vector DB](#)

IoT and observability · Infra · DevOps

Infrastructure as code with Terraform, containers and in-house monitoring and error-reporting SDKs.

[Terraform](#) · [Docker](#) · [Observability](#)

TECHNICAL SKILLS

Languages: TypeScript, JavaScript, Python, PHP, C#, Java, Kotlin, Swift, Dart, Go, Rust, C++

Mobile: Kotlin (Android), Swift (iOS), React Native, Expo, Flutter

Frontend: React, Vue, Angular, Next.js, Astro, Tailwind CSS, Inertia

Backend: Laravel, Yii, FastAPI, Django, ASP.NET, Spring Boot, Node.js, Express

AI & LLMs: RAG, AI Agents, Embeddings, Vector DBs, pgvector, Qdrant, Fine-tuning, PyTorch, Hugging Face, LangChain, LangGraph, LlamaIndex, Ollama, Prompt Engineering, n8n / Make

Cloud: AWS, Azure, Google Cloud, Cloud GPU/ML, Serverless, Terraform

DevOps: Docker, Kubernetes, CI/CD, Prometheus, Grafana, Ansible, Git

Data: PostgreSQL, MySQL, MongoDB, Redis, SQLite, Supabase

IoT & Industry: LoRaWAN, ChirpStack, MQTT, EMQX, Modbus, PLC / HMI, Node-RED, Raspberry Pi, Arduino

Integrations: REST, GraphQL, Webhooks, RabbitMQ, Kafka, Microservices, ETL/ELT

CONTINUOUS LEARNING

Hands-on and always on the frontier: I track the state of the art in generative AI — RAG, agents and LLMs — and bring every advance into production.

LANGUAGES

Portuguese — native · English — technical reading & writing

GVMS TECH LTDA · CNPJ 62.776.795/0001-81 · Recife/PE

[gilmar.dev](#) · Active since 2007 · Updated 2026