

Jakub Szwajka

Senior Full Stack Engineer

Wrocław, Poland | Remote (EU/US timezones)

📞 (+48) 696 657 022

✉️ szwajkajakub@gmail.com

🔗 [LinkedIn](#) | [GitHub](#) | [Blog](#)

Summary

Senior Full Stack Engineer with 6+ years building 0-to-1 products for US startups. I care about the "why" as much as the "how" - understanding the business context before writing code, and building systems that grow with the product.

Backend, frontend, infrastructure, architecture, e2e testing - I've done all of it, often on the same project.

Best fit: Growing products where I can shape architecture alongside the business - greenfield builds or systems constantly building new capabilities. Not maintenance.

Stack: Python (FastAPI), TypeScript (React), AWS (Fargate, CDK, Kinesis), PostgreSQL, Docker

AI/Tools: LLMs, RAG, MCP, Claude Code, Agentic workflows

Languages: Polish (native), English (fluent - daily work language with US clients)

Experience

Sofomo | Senior Full Stack Engineer (Jun 2022 – Present | Remote, US clients)

Embedded with US-based startups as a contractor through Sofomo. With each project taking on increasing ownership and technical leadership across different clients, teams, and tech stacks. Beyond client work: mentoring junior engineers, technical recruiting, architecture consulting for new projects, and creating technical documentation.

IoT Platform (100K+ connected devices) - Data Platform Engineer / Architecture Consultant

- Helped build the data team responsible for real-time processing and analytics services.
- Built data warehouse and streaming pipelines (AWS Fargate, AWS Kinesis Streams, Kinesis Firehose).
- Conducted architecture review of external team's backend rewrite; provided technical recommendations that informed project direction.
- Scale: 100,000+ connected IoT devices, high-throughput data ingestion.

Real-time Captioning Marketplace - Tech Lead / Fullstack Engineer

- Building 0-to-1 marketplace for booking live captioners for meetings and events.
- End-to-end ownership: backend (Python), admin panel (React), full infrastructure.
- Stripe integration for payments and bi-weekly contractor payouts.
- External service integrations for real-time captioning delivery.
- Greenfield product, currently in market launch phase.

EdTech Platform (SAT/ACT Test Prep) - Lead Backend Engineer

- Built entire backend from scratch using Next.js.
- Designed test generation system serving 100+ concurrent test-takers in US schools.
- Contributed to frontend when needed.

AI-powered Enterprise Platform (\$350M+ funding) - Software Engineer

- Owned a domain service within a microservices architecture (Kubernetes).
- Worked in a small, autonomous team responsible for end-to-end delivery.

TietoEvry | Software Engineer (Aug 2021 – Jun 2022 | Wrocław, Poland)

- Implemented features for fleet management web application (Django).
- Refactored legacy components to improve development velocity.
- Active participant in feature design and code review processes.

BTech | Python Developer (Jun 2019 – Aug 2021 | Wrocław, Poland)

- Built web platform for managing GCP notebook instances (Django, React).
- Set up CI/CD pipeline (Docker, GitHub Actions) - to improve CI process and release frequency.

AI & LLM Projects

Conversational AI Agent with RAG & Long-term Memory (2024)

- Built AI assistant with persistent memory using **vector embeddings** and semantic retrieval.
- Implemented full **RAG pipeline**: vector database for storage, embeddings for semantic search.
- Designed **memory distillation system** - LLM extracts facts from conversations, embeds them, stores for future context retrieval.
- Multi-source orchestration: Slack (real-time chat), Airtable (structured data), vector search.

Multi-Agent System via Prompt Engineering (ongoing, in daily use)

- Built **4 specialized AI agents** (productivity, writing, finance, strategy) on a single platform through prompt engineering.
- Designed agent architectures: distinct personas, capabilities, memory systems.
- Agents collaborate and reference each other's context within conversations.

AI-Augmented Development Workflows

- Daily use of **Claude Code** for repository-aware development - custom instructions tuned to project standards.
- Built custom **MCP server** (Model Context Protocol) - custom search tool for LLM to extend the regular UI search in client app.
- Building and iterating on prompts and agent configurations for consistent, high-quality code generation.

What I Bring

Full-stack ownership - I take full ownership. When I take on a project, I'm happy to own it end-to-end: backend, frontend, infrastructure, deployment.

Architecture with business context - I start by understanding the business domain - identifying natural boundaries and reusable system capabilities. The goal: build systems that evolve with the business, not against it.

Async-first, remote-native - 3+ years working with US clients across timezones. Clear written communication, proactive updates, and documentation that doesn't require a call to understand.

AI-augmented development - I build and configure AI agents tailored to specific repositories, tuned to follow project standards and deliver precise work. This has a direct impact on my development speed.

Education & Certifications

WWSIS Horyzont - B.Eng. Computer Science (2018–2021)

Thesis: Vendor Compliance Monitor for VAT Whitelist Verification in SAP ERP

Professional Development:

- **Developer Jutra** (*in progress*) - Full-stack development with AI tooling: TypeScript, React/Next.js, AI-powered coding workflows.
- **Domain Drivers** (2025) - Comprehensive Domain-Driven Design program (strategic & tactical patterns, Event Storming, bounded contexts).
- **Legacy Fighter** (2024/2025) - Refactoring and legacy code modernization: techniques for safely evolving complex systems.
- **DNA: Designing Next-Gen Architecture** (2023/2024) - Software architecture program covering DDD, Event Sourcing, microservices, and legacy modernization.
- **AI Devs** (2023/2024, 2 editions) - Practical AI/LLM course for developers: building agents, prompt engineering, production AI workflows.