

Vova Zakharov

Fractional, hands-on CTO for AI-native delivery

vzakharov@gmail.com · vovazakharov.com

Profile

158 days, **250,000 lines of production TypeScript**, none of it hand-written, up to **20 agents working at once**, and a **production deploy every 2.4 days** — on a live product that stayed in production for its users the whole way through. The engineering platform I built to get there outlived my involvement: three engineers run it today, and it travels to their new projects as a checked-in boilerplate.

Most of my work has revolved around LLMs and image generation ever since 2020, following years of messing with smaller neural networks. This means I **don't buy or sell buzzwords**, I understand things **a level or two deeper than most**, and I've developed an **intuition for what models can and cannot do**. What I'm not is a compliance function or a headcount manager: this is founder-and-small-team scale, where I run the delivery rather than the org chart.

What I Offer

Engagements

- **Idea → production:** your idea in the hands of real users in one to two months, built the way it will be maintained rather than as a demo to throw away.
- **Architecture and foundations:** the system drawn and the foundation stood up — framework and platform chosen, the structure and its conventions in place — for a team that will build the features itself.
- **Your team on agent rails:** the architecture, the staged pipeline and the pre-push gate that let a crowd of agents ship without drifting — set up on your codebase, then handed over as checked-in procedure rather than as advice. It has been handed over once already, and the engineers who took it carried it to projects of their own.
- **Rescue and diagnosis:** the prototype that got vibe-coded into a corner, and the agent that doesn't do what you expected — usually the same purchase.
- **Standing technical judgment:** build-vs-buy calls, and a straight answer on what the models still cannot do.

The Engineering System

- Feature-sliced architecture held by linters, so a crowd of agents working at once cannot drift the structure.
- A staged pipeline — plan, implement, review, finalize — each stage its own session, with reviewable artifacts between them.

- Branch hygiene held in-session: a gate of concurrent checks before every push, the production build among them, so no agent contaminates main with code that breaks it.
- CI/CD that fits your workflow — ship every commit off main, staging-based release cycles, or release trains.

AI Expertise

I tend to stick to a more "economical" paradigm of AI usage: smaller prompts, more aware of context limitations, less "let's feed it all and it'll somehow work!" My expertise predates the ChatGPT hype cycle, which means I understand not just the APIs, but the underlying patterns and limitations that most "AI engineers" are still discovering.

Working Style

I'm the kind of developer you can generally leave unattended, as long as the overall vector is clear. I have enough life expertise to figure stuff on my own, enough imagination to fill the missing pieces, and a "treat any job as if it was your own brainchild" mentality.

In other words: a surprisingly low-maintenance LLM tinkerer with a TypeScript kink, ready to prototype, debug, or sanity-check your AI-infused ambitions.

Experience

Fractional CTO

Playgram (March – August 2026)

Read the case study: yovazakharov.com/case-studies/playgram

Rebuilt a live, feature-rich AI chat product — multiple model providers, realtime team chats, image and file libraries, memory and knowledge management, voice input — from Bubble, a no-code builder, into a production Next.js 16 codebase, while it stayed in production for its users throughout.

Delivered in 158 days — parallel AI sessions, up to 20 agents working at once:

- **1,395 commits on main** and 250,000 lines of production TypeScript, none of it hand-written.
- **48 releases and 18 hotfixes:** a production deploy every 2.4 days.
- **Sub-second cold loads**, down from multi-second.
- **Feature-sliced architecture held by linters**, so a crowd of agents working at once could not drift the structure.
- **A repeatable agent pipeline:** plan, implement, review, finalize — each stage its own session, with a reviewable artifact between them.

Next.js 16, Railway + Supabase, feature-sliced design

Senior Fullstack Developer

October 2025 – February 2026

AI-powered English learning application for kids, combining generative AI with time-tested pedagogical methodology to make language acquisition engaging and effective.

- Interactive game-based learning experiences powered by LLM-generated content.
- Adaptive difficulty and personalized learning paths based on individual progress.
- Integration of proven teaching methods with modern AI capabilities.

Next.js, OpenAI API, custom game engine

Senior Fullstack Developer

Orcool (June – August 2025)

Short-term project shaping an AI-agent-based review intelligence tool for brand marketing.

Built the system end-to-end, covering:

- Automated review collection from diverse public sources.
- AI-driven summarization with structured outputs (SWOT, competitive positioning, feature insights).
- Early productization steps toward a scalable SaaS-style platform.

Next.js/NestJS, Cloudflare Workers, Firebase, custom LLM orchestration framework

Fullstack Developer

randddb.com / DDB (2023 – 2025)

Experimental AI platform at one of the world's largest ad firms. Built, back to front, a suite of interlinked AI tools:

- **Chatbot:** ChatGPT-style assistant for corporate needs — secure, document-aware, image-savvy, with personas and real-time multi-user support.
- **Robowriter:** Build-your-own-agent tool for content generation. Functions as both user-facing product and meta-backend for LLM ops.
- **Picmaker:** Replicate-powered image generation/training tool with UX optimized for rapid iteration and intuitive pinning.
- **Client-specific apps:** Including YesYouCannes, which builds Cannes award decks from idea to imagery in minutes.
- **Spaces:** White-label mini-platforms for agency clients (e.g. randddb.com/s/ikea).

Django + PostgreSQL, Vue + TypeScript

Demo available on request.

Developer – Independent Projects

2020 – Present

Selected highlights:

- **Jukebox Web UI:** Gradio-based UI running inside Google Colab for OpenAI's music model.
 - **almostmagic:** TypeScript package for instant LLM integration – one line, zero ceremony.
 - **JobGenie:** AI assistant for job seekers that helps create roles as much as apply for them.
 - **Flows:** A Django app for orchestrating async task pipelines with minimal ceremony.
 - **Unfindables:** An experiment in surfacing valuable but hidden web content via absurd search queries.
-

Prototyper – Experience & Innovation

Voicemod (2023)

Prototyped bleeding-edge tools for audio + AI experimentation.

- Built a YAML-powered Discord bot framework for zero-code bot creation.
 - Developed a local Python web API with a ChatGPT-style interface.
-

Tech Stack

Backend

- **Django / FastAPI** – If you rely on your own ML models or otherwise require computations/libs that the JavaScript world hasn't yet caught up with
- **Next.js API** – If you're mostly using other people's web APIs and don't want to go all fancy and domain-driven
- **NestJS** – If you do want to go all fancy and domain-driven

Frontend

- **Next.js / React** – If you want to stay on the safe and "we can definitely hire another dev to work with it later" side
- **Nuxt / Vue** – If you're feeling all progressive and SFC-ish

Serverless / Other

- **Cloudflare Workers** – Because why not rely on something that can make the entire Internet stop a few times a year?
 - **Google Colab** — Had my fair share of fails and successes building and testing models
 - **Gradio** – Quick but powerful UIs for ML experiments
-

Education

Moscow Institute of Physics and Technology

Master's in Applied Math & Physics (2000 – 2006)

Contact

vzakharov@gmail.com · github.com/vzakharov · linkedin.com/in/vovahimself · x.com/vovahimself

Full portfolio: vovazakharov.com