

Konstantin Gredeskoul

AI-First Staff / Principal Software Engineer, Tech Lead
Startup ex-CTO (4x), Conference Speaker, Open Source Contributor

Staff/Principal engineer who ships AI-assisted systems into production: retrieval/embedding pipelines, eval harnesses, and automation around real business workflows. Strong bias toward reliability, observability, and safety-by-design so the AI doesn't become a new failure mode. Two decades of experience building and scaling distributed platforms. Strong foundation in systems architecture, building distributed systems with a focus on business value, reliability, fault-tolerance, auto-recovery, all the while applying advanced techniques such as pair-programming, test-driven development, and more recently — multi-agent code generation. I thrive in product environments that value ownership, technical rigor, and long-term system quality over demos and rapid-fire MVPs.

Creative, methodical problem-solver with a strong foundation in mathematics and statistics, comfortable operating across levels of abstraction from system architecture to low-level debugging. I prioritize clean, maintainable code and robust systems delivered iteratively, with high test coverage, clear documentation, and practical design. I deliver code that works, requires little to no maintenance, and flexible enough to support future changes.

I focus on reducing operational risk through gradual rollouts, canary deployments, A/B testing, and strong observability, while improving developer experience by streamlining onboarding, CI, and deployment workflows. After years of building and scaling large distributed systems, I sought hands-on production experience with applied AI, shipping real systems quickly. **I'm at my best in product environments where I can own systems long-term and continuously raise the bar on quality.**

CAREER HIGHLIGHTS

Transformational technical leader with two decades of building distributed scalable transactional systems serving tens of millions of users. 4x CTO, led complete system rewrite (90K LOC Java → 10K LOC Rails in 2 months), and, immediately after that scaled the platform 100x to 5K req/sec with just six engineers pairing full-time, while maintaining 99.99% uptime four years in a row. At another company, prevented a potential extended PostgreSQL downtime due to transaction wraparound on a 27TB PostgreSQL table. Expert Rubyist, with over 50 Ruby Gems on rubygems.org and a **cumulative 200M downloads**.

Proven track record solving complex problems and leading by example: As a Systems Architect, chose the technologies and architected one of the largest distributed email delivery platforms of its time (½M msgs/day) before the Cloud even existed. At Homebase, within the first two weeks completely fixed daily PostgreSQL crashes during peak hours. At HealthSherpa achieved 100% uptime during critical ACA open enrollment period (handling 10x traffic jump overnight). Built a fully fault-tolerant system supporting 25 million active users with near-zero pagerduty alerts. Introduced transformative engineering practices — full-time pair programming, TDD, stand-ups, retrospectives at Wanelo — creating self-managing, ultra-efficient, and high-performing team, loosing just one engineer in 3.5 years. Conducted over 500 interviews, including bootcamp graduates, hiring top talent into the team, some of whom were recognized in Forbes Under 30 honorees.

Industry thought leader with an eye for a talent:

Co-created original open source Ruby rules for Bazel mono-repos, inspiring current industry standard. Enthusiastic self-learner, mentor, and speaker. Volunteered on multiple occasions to teach Rails (RailsBridge), was invited to speak at several local and international conferences, (e.g. RubyConf Australia, Silicon Valley PostgreSQL, Ruby Mexico, etc). Scaling on top of PostgreSQL presentation has ½M cumulative views on SlideShare. Judged and advised LATAM startups at a “Geek Camp” — startup incubator in Santiago, Chile, run by Astral Capital.

Monash University, Australia

Bachelor of Science w/Honors
Mathematics & Statistics Summa
Cum Laude, Top 1%



[Published Graduate Thesis](#)

Kharkov State University, Ukraine

Advanced Mathematics,
Calculus & Analysis



Ukraine National Mathematics Olympiad Winner

SKILL PROFICIENCY

ALL 10 POINTS IS THE EXPERT LEVEL

Systems Thinking & Architecture



Scaling to Manage Spikes in Concurrency



PostgreSQL, NoSQL, Sharding, SQL



Ruby, Ruby on Rails, E-Commerce



Python 3, Pydantic, AI frameworks



C (99) / C++ (2011)



JavaScript, ECMA'19, NodeJS



TypeScript, Node, JS (ECMA'19)



Unixes, BASH/ZSH, Docker, IPCs



DevOps, Cloud, Security, Terraform, CDK



AI Systems, Agents, Tooling



OpenAI, Anthropic, Gemini APIs



ReactJS, JSX



PROFESSIONAL EXPERIENCE

Fractional.ai, Software Engineer

09/2025 – 02/2026

- Designed and shipped a production recommendation service for an e-commerce platform, implemented in Python with PostgreSQL and pgvector, blending semantic similarity with historical sales and ranking signals.
- Built an automated, zero-downtime data refresh pipeline that regenerates embeddings and derived database objects from nightly PostgreSQL backups using parallel schemas, followed by a sub-100 ms multi-schema swap with full rollback capability.
- Implemented end-to-end evaluation and reliability tooling using Pydantic, Logfire, and Weave, ensuring recommendation quality and operational stability in production.
- Delivered an internal AI automation integrating Slack and Notion to extract actionable todos from team conversations, reducing manual tracking and improving follow-through.

Paternity Leave of Absense, two months

06/2025 – 08/2025

Academia.edu, Staff Software Engineer

08/2024 – 06/2025

Halved CI build times (15→7 min, hundreds of runs daily) by introducing Knapsack and automating CI/CD image builds with Terraform and AWS CodeBuild. Built a one-command dev environment, cutting onboarding from weeks to hours. Led a Datadog metrics overhaul, authoring an open-source Ruby gem for custom metric validation, and delivered internal training on Ruby concurrency. Redesigned and shipped the card expiration warning email system.

HealthSherpa, Principal Software Engineer

02/2022 – 03/2024

Removed backend single points of failure, achieving HealthSherpa's first 100% uptime during ACA open enrollment under 10x traffic spikes. Automated onboarding (1 week → 1 hour), unified observability with Datadog, and delivered internal talks on concurrency, PostgreSQL, and scalability.

FOSSA, Staff Software Engineer

09/2020 – 09/2021

Averted catastrophic XID wraparound on a 27TB PostgreSQL DB by migrating to multi-AZ RDS and completing emergency vacuum under 2-month deadline. Cut AWS costs, compacted DB schema, and added unique job support to Faktory's NodeJS client—mastering TypeScript/Kubernetes to deliver production-grade features fast.

Coinbase, Staff Software Engineer (L6)

06/2019 – 02/2020

Unified Ruby & Go into a Bazel mono-repo, creating the first open-source [rules_ruby](#) for Bazel, to enable hermetic Ruby builds/tests. Pioneered Bazel support for Ruby using Starlark, laying the groundwork for the current [Bazel/Ruby](#) integration that currently is part of Bazel Contrib.

Homebase, Infrastructure Architect

08/2017 – 01/2019

Stepped in during critical instability caused by daily peak traffic, leading cross-functional efforts that restored 99.99% uptime through strategic DB read/write splitting (Makara) and system hardening. Spearheaded full AWS infrastructure rebuild using Chef and Docker, cutting cloud spend by 50%. Drove major Ruby performance gains (3x memory reduction, optimized GC/concurrency), introduced zero-downtime deploys, and accelerated CI by 2-3x. Empowered a high-performing team and built systems that let them sleep soundly again, and could be scaled another 10x without rewrites.

12/2011 – 01/2016

Wanelo.com, CTO & VPE

Joined as employee #3 and built the Engineering team from scratch to 25 employees. Collaborated with the founding engineers in rewriting the entire app in Rails. With 6 engineers, we rebuilt it in 2 months, creating 10x less code than Java, and 99% test coverage. Over the following six month, [traffic jumped 100x](#).

Established CI/CD, TDD, continuous deployment, and pairing culture. Conducted over 1K interviews, and had just one person leave in 3.5 years. Invented a new distributed design pattern based on HAProxy for routing and auto-recovery—featured in my RubyConf talk "[DevOps Without the Ops](#)".

The application reached top-10 Apple Store rankings multiple times, won [TechCrunch Disrupt](#), and handled 10K requests/second through heavily optimized multi-threaded Puma processes running on Joyent Cloud.

OPEN SOURCE CONTRIBUTIONS

This sidebar features several open source projects that demonstrate a wide range of skills, while at the same time showing the code I wrote in the past, proving the ability to deliver and finish a project, whether it's an MVP, an online free tool, or a game.

[gomoku-ansi-c \(3 ★\)](#)

ANSI Terminal game with a pretty strong AI player that you can build & play on a Mac OS.

[simple-feed \(335 ★\)](#)

Ruby gem that uses Redis shards to implement O(1) read time for a social feed of any kind.

The next set of gems are simply either popular, or unusual, or interesting in their own right.

[puma-daemon \(39 ★\)](#)[cmake-project-template \(945 ★\)](#)[bashmatic \(170 ★\)](#)[sym \(138 ★\)](#)<https://makeabox.io>

Free online tool for creating vector drawings for laser-cutting 3D enclosures

<https://bashmatic.dev>

Project site for the open source Bashmatic framework.

CONSULTING

Between 2015 and 2018 I ran a consultancy through the newly founded company ReinventONE: <https://reinvent.one>.

The contracts typically covered urgent issues such as scalability, lowering latency of the web pages or API responses, adding observability, vertical and horizontal sharding, multi-AZ migrations, event-driven data modeling, etc.

The clients included:

- [Flow Investor Portal](#)
- [HIRED.com](#)
- [JoinHomebase.com](#)
- [Reflektive, Inc](#)
- [Returnly.com](#)
- [GoShippo.com](#)
- Butter.com

SPOKEN LANGUAGES

- ENGLISH
- RUSSIAN
- SPANISH (Beginner)

REFERENCES

Available on request