

Rikhi Singh

Full-Stack Engineer — AI/ML Specialist — Cloud Architecture

Waterloo, ON | +1 (647) 239-1336

rikhi2003@gmail.com | rikhi.dev | linkedin.com/in/rikhi-singh | github.com/rikhisingh | npmjs.com/~rikhi

PROFESSIONAL SUMMARY

Full-stack software engineer with **3+ years** shipping production systems serving **90K+ users**. Specialized in **AI/ML integration (RAG pipelines, LLMs)**, cloud-native architecture (TypeScript, Next.js, AWS), and high-performance APIs (<100ms latency). Track record of leading cross-functional teams to deliver **45% faster** while reducing costs by **\$20K+ annually** through strategic automation and optimization.

TECHNICAL SKILLS

Languages: TypeScript, JavaScript, Python, Java, C#, C++, SQL (PostgreSQL, MySQL, Oracle), PL/SQL

Frontend: React, Next.js 15, React Native, Tailwind CSS, Responsive Design, WCAG/AODA Accessibility

Backend: Node.js, Hono, Bun, tRPC, FastAPI, Express, ASP.NET MVC, Spring Boot

Databases: PostgreSQL (Drizzle ORM, Prisma), Oracle, MySQL, SQL Server, Redis, PostGIS, pg-vector

AI/ML: LangChain, GPT-4, Gemini, Morphik, YOLOv8, RAG, Vector Search (Pinecone, Qdrant), HNSW

DevOps: Docker, Kubernetes (EKS), Terraform, ArgoCD, CI/CD (GitHub Actions), BullMQ

Cloud & Tools: AWS (S3, SES, EC2), Vercel, Grafana, Prometheus, Splunk, Toad, SSIS, Git

PROFESSIONAL EXPERIENCE

- **Aphantasia Network** Kitchener, Ontario
Mar 2025 – Present
 - Architected and shipped production **turborepo monorepo** from scratch with 3 Next.js apps and 10 shared packages (**200K+ LOC**), engineering 6-module SaaS platform (CRM, LMS, events, directory, surveys) with unified authentication, database layer, and component library, **enabling 65% faster builds** and parallel team workflows.
 - Engineered high-performance **REST API with 200+ endpoints** serving <100ms p95 latency across 30 route handlers and 38-table PostgreSQL schema, implementing PostGIS geospatial indexing, pg-vector semantic search, and BullMQ job queues, processing **50K+ daily queries** for research recommendations and event coordination.
 - Built intelligent research pipeline using **Morphik for document parsing and RAG**, integrating vector embeddings and semantic search across research paper corpus to power AI-driven recommendations, **improving content discoverability by 40%** and enabling semantic similarity searches.
 - Automated end-to-end **Stripe integration managing 8 subscription products** with webhook-driven role synchronization and Redis caching, **eliminating 12+ hours/week of manual billing work** and enabling real-time capability-based access control across membership tiers.
 - Built CI/CD infrastructure with **GitHub Actions** automating deployments to Coolify, configuring 3 parallel workflows with turborepo caching, achieving **zero-downtime releases and 65% faster deployments**.
 - Implemented enterprise security with **10-role RBAC system governing 20+ capabilities**, Zod validation across all API endpoints, Better Auth with magic link and OAuth support, and comprehensive structured logging (Pino + CloudWatch), ensuring secure multi-tenant data isolation and **reducing security incidents by 40%**.
- **Eptics Inc.** Remote / Canada
Jul 2024 – Feb 2025
 - Full Stack Engineer (*Contract*)
 - Led **4-person cross-functional team** to architect and deliver AR-based laboratory training platform using **React Native, Unity WebGL, and FastAPI**, implementing real-time 3D rendering with **WebRTC** for collaborative sessions, achieving **60 FPS on mobile devices** and **reducing training costs by 40%** for 15+ enterprise clients.
 - Engineered **AI-powered crop recommendation engine** using **LangChain, GPT-4 Turbo, and computer vision (YOLOv8)**, processing **100K+ farm images and IoT sensor data** with vector search (Pinecone) for personalized strategies, achieving **89% prediction accuracy** and driving **35% increase in customer retention**.
 - Built scalable e-commerce platform using **Next.js 14, tRPC, and Prisma ORM**, implementing **edge caching with Vercel** and Stripe Connect for multi-vendor payments, supporting **200+ farm vendors** with **99.9% uptime** and sub-200ms page load times.
 - Implemented comprehensive observability stack with **Grafana, Prometheus, Splunk, and LogRhythm SIEM**, configuring **40+ custom dashboards and ML-based anomaly detection**, reducing false positive alerts by **45%** and cutting incident response time from **2 hours to 12 minutes**.

- Deployed **containerized microservices** using **Docker, Kubernetes (EKS), and Terraform IaC**, implementing GitOps workflows with ArgoCD, auto-scaling policies, and blue-green deployments that improved deployment reliability by **75%** and reduced infrastructure costs by **\$8K/month**.

- **University of Guelph** Guelph, Ontario
App Developer (Contract) Sep 2023 – Jul 2024
- Engineered **enterprise web applications** using **ASP.NET MVC (C#), Oracle Database (PL/SQL), and JavaScript**, implementing role-based access control and RESTful APIs serving **15K+ users**, improving page load times by **42%** and reducing server costs by **\$6K/year**.
- Built **automated data migration pipeline** using **Python (Pandas), Toad for Oracle, and SSIS**, orchestrating **5M+ record transfers across Oracle, MySQL, and SQL Server** with **zero data loss**, reducing manual migration time from **3 weeks to 2 days**.
- Developed **internal admin tools suite** with **Java Spring Boot and React**, automating **12+ daily workflows**, **eliminating 60% of manual errors** and saving **25 hours/week** across 8 administrative teams.
- Spearheaded **security remediation project**, patching critical ColdFusion vulnerability, conducting **OWASP audits across 40+ applications**, implementing **automated scanning (Snyk)**, achieving **100% compliance** and preventing breach affecting **50K+ student records**.

SELECTED PROJECTS

- **Better Tables (Open Source)**: Architected and published **type-safe table library for modern web frameworks** solving complex relationship filtering and type inference challenges from database to UI layer. Actively maintained with **4K+ NPM downloads, 70+ resolved issues, and 850+ monthly downloads** since launch. Implements zero-config API using TypeScript generics and Drizzle ORM schema introspection.
Tech Stack: Next.js 15, TypeScript, Drizzle ORM, PostgreSQL, Tailwind CSS 4, Better Auth, AWS, Docker.
- **BeaverMade.ca (Viral Launch)**: Engineered and launched **AI-powered product recognition platform** enabling Canadians to identify local products via image search. Built **native iOS (Swift) and Android (Kotlin) apps** with computer vision pipeline (Gemini Vision API) and **vector search across 2M+ product embeddings** (Pinecone + HNSW). Achieved **10K+ users in 48 hours**, featured by **CBC and Toronto Sun**, processing 50K+ image searches in first week with <500ms latency.
Tech Stack: Swift, Kotlin, Next.js, TypeScript, Gemini, OpenAI, Pinecone, Qdrant, AWS.
- **Better Auth Contributor (Open Source)**: Core contributor to **authentication framework with 25K+ GitHub stars**, owning Stripe subscription and billing integration. Architected **fault-tolerant webhook processing system** with idempotency guarantees, implemented token lifecycle management, and ensured backward compatibility across 5 major releases. Contributions used by **1K+ production applications**.
Tech Stack: TypeScript, Node.js, JWT, Stripe API, PostgreSQL, Redis, Bun.

EDUCATION

- **Congestoga College** Waterloo, ON
Diploma in Computer Programming with High Distinction Sep 2023 – Apr 2025
- **University of Michigan** Ann Arbor, MI
Computer Science & Software Engineering Studies 2021 – 2022