

Shreyaan Seth | Backend & AI Infrastructure Engineer

Amritsar, India (Open to Remote)

✉ shreyaans20@gmail.com

🌐 [linkedin.com/in/shreyaan-seth](https://www.linkedin.com/in/shreyaan-seth) • github.com/shreyaan

Summary

I am a product engineer who takes systems from zero to production, owning everything from architecture and backend systems to the infrastructure that runs them. I've built and shipped real-time voice platforms, multi-agent AI systems, and production-grade RAG pipelines —designing for low latency, reliability, and scale from day one. I operate everything I ship in production, currently running 10 concurrent services across a fleet of 6 VMs, handling deployment, monitoring, and ongoing system reliability.

Work Experience

Autonomix Solutions Inc

Remote

Engineering Lead —Backend & AI Infrastructure

Aug 2024–Present

Sole architect and primary builder of all backend systems, AI infrastructure, and DevOps across the organization's entire product portfolio. Designed, built, and maintained every system from scratch while leading a 5-person engineering team.

- **Voice AI Platform** : Architected and built a production-grade outbound and inbound calling platform from 0 to 1, featuring real-time WebRTC streaming, live transcription, AI agent personas, and multi-channel support (voice, WhatsApp, Slack, widget). Now being productized as a SaaS offering.
- **Multi-Agent Document Intelligence (Provisional Patent)**: Designed and built a multi-agent orchestration system for automated document generation in early 2025 —ahead of mainstream adoption of agentic AI patterns. System is the subject of a provisional patent filing.
- **Production RAG Infrastructure**: Architected end-to-end RAG pipelines featuring hybrid vector/keyword retrieval, multi-source document ingestion (Notion, Slack, Calendar, web scraping), fault-tolerant embeddings, and memory persistence layers serving enterprise clients.
- **Multi-Tenant SDK**: Designed and built a production SDK with strict environment isolation, RBAC, audit logging, usage quotas, and infrastructure cost controls —enabling onboarding of enterprise clients with zero cross-tenant data risk.
- **DevOps & Infrastructure Ownership**: Built and maintained CI/CD pipelines, deployment automation, and a fleet of 6 VMs hosting 10 concurrent production projects. Sole owner of all infrastructure from day one.
- **Full-Stack Java Delivery (Enterprise Client)**: Self-taught Spring Boot and Java in 3 weeks to lead delivery of a 4-month, two-phase enterprise project —shipped on time. Phase one covered core platform build; phase two delivered Payments and self-onboarding, Calendly integration, and Square integration.
- **Engineering Hiring**: Conducted senior technical interviews (System Design, DB internals, Message Queues) to scale the engineering team from 10 to 20+ members.

Tech Stack: TypeScript, Python, Hono, ElysiaJS, FastAPI, Next.js, PostgreSQL, Vector DBs, Redis, RabbitMQ, WebRTC, Multi-Tenant Architecture, AWS, GCP, Docker, GitHub Actions, CI/CD.

ByteLeap

Remote

Software Engineer

Feb 2024–Aug 2024

Owned end-to-end performance and infrastructure for a high-traffic production platform.

- **Backend Optimization**: Reduced critical API latency by 94% (7s → 400ms) through query profiling, indexing strategy overhaul, and infrastructure refactoring.
- **Frontend Performance**: Improved Next.js application performance by ~50% via bundle analysis, dynamic imports, and advanced caching layers.
- **Infrastructure Migration**: Migrated platform from Vercel to self-hosted cloud infrastructure, significantly cutting monthly costs while improving deployment control.

Tech Stack: Next.js, TypeScript, PostgreSQL, Self-hosted Infrastructure, Performance Engineering.

Co.Lab

Remote

Software Developer

Jul 2023–Nov 2023

Built ScamGuardian —a Chrome extension with scam-detection heuristics, URL scanning APIs, and user-facing educational UX flows.

Tech Stack: Chrome Extension APIs, React, Next.js, Web Security.

Side Projects

Diet Paglu

dietpaglu.com

Founder & Solo Engineer

2025–Present

Built an AI-powered nutrition tracking product end-to-end —solo. Users log meals via photo, voice, or WhatsApp ; the system uses conversational AI to detect food, estimate macros, and track progress automatically.

- Zero-friction UX: no manual entry, works natively over WhatsApp —targeting India’s mobile-first market.
- Built entire stack: AI food recognition pipeline, WhatsApp integration, backend, and frontend.

Tech Stack: Typescript, Bun, Hono, Next.js, LLMs, WhatsApp Business API, PostgreSQL.

Technical Skills

Languages: **Expert:** TypeScript, Python, JavaScript, Java | **Proficient:** SQL, Go

Backend: **Expert:** Hono, FastAPI, Node.js, Next.js, Bun | **Proficient:** Express, Spring Boot

Databases: **Expert:** PostgreSQL, Redis, MongoDB | **Proficient:** Vector DBs (Pinecone/Milvus), RabbitMQ

Cloud & DevOps: **Advanced:** AWS, GCP, Docker, GitHub Actions, CI/CD, Serverless (Lambda/Cloud Run), VM fleet management

AI & Systems: **Advanced:** RAG Architecture, Multi-Agent Orchestration, Real-time Voice (WebRTC), System Design, Multi-Tenant SaaS

Education

D.A.V College

Amritsar, India

Diploma in Digital Communication & Media

Jan 2019–Jan 2022