Vaibhav Saini

vaibhavsaini2409@gmail.com | LinkedIn | GitHub | Portfolio | Delhi, India

Profile

Computer Science and Engineering student focused on building reliable full-stack applications, backend systems, and AI-powered products. Interested in system design, core computer science concepts, and maintainable architecture.

Education

G.B. Pant DSEU Campus-I, Delhi, India — Aug 2024 – Jun 2028
Bachelor of Technology in Computer Science and Engineering
Relevant Coursework: Data Structures and Algorithms, OOP, DBMS, Operating Systems, Computer Networks

Technical Skills

Languages: TypeScript, JavaScript, Python, SQL (PostgreSQL), Dart
Frontend: Next.js (App Router), React 19, Tailwind CSS, Vite
Backend: Node.js, Express.js, FastAPI, tRPC, NestJS, REST APIs, Server-Sent Events (SSE)
Databases & ORMs: PostgreSQL (Raw SQL), Drizzle ORM, Prisma ORM, Supabase, SQLite, Redis
AI & Systems: Inngest pipelines, process sandboxing, OpenAI, Claude, and Gemini APIs
Tools & Testing: Git, Docker, Vitest, Cypress E2E, Postman, Swagger, GitHub Actions CI/CD

Projects

Shipshape – Autonomous Next.js Application Builder & Sandbox Runtime — GitHub
Next.js 16, React 19, TypeScript, Turbopack, Prisma ORM, SQLite, SSE, Vitest

- Architected a five-stage synthesis pipeline spanning intent planning, structured specifications, multi-file synthesis, sandboxed compilation, and unified diffs; generates 12–18 Next.js App Router files per project.
- Built a bounded three-attempt self-repair loop that uses TypeScript compiler diagnostics to target patch generation, achieving a 94.2% first-run compile pass rate.
- Implemented an isolated process sandbox manager with automatic port allocation, process lifecycle isolation, and secret redaction before terminal logs are persisted.
- Added immutable project versioning and line-level diffs for deterministic rollbacks, plus standalone ZIP source export without vendor lock-in.

Fellow – Multi-Tenant Collaborative Workspace Platform — GitHub
Next.js, TypeScript, tRPC, Drizzle ORM, PostgreSQL, Supabase Auth, DnD Kit

- Designed a multi-tenant PostgreSQL schema with organization-scoped RBAC and org-scoped data access to enforce tenant isolation.
- Implemented tRPC middleware for end-to-end compile-time contract safety between client and server, reducing runtime schema drift.
- Used fractional indexing for Kanban task reordering without rewriting adjacent rows, with automated rebalance thresholds.
- Built secure invitation workflows using organization-scoped invite links and atomic role provisioning when users join.

Memora – AI Knowledge Workspace & Action Management Platform — GitHub
Next.js 14, TypeScript, PostgreSQL, Drizzle ORM, Tiptap, Gemini API, Tailwind CSS, Zod

- Engineered a distraction-free rich-text editor with a debounced background persistence pipeline and optimistic concurrency control (OCC), eliminating stale-write data loss across concurrent edits.
- Implemented an asynchronous AI processing layer (Gemini API with strict Zod validation) to generate reviewable title suggestions, executive summaries, and action items linked to note revisions.
- Built PostgreSQL native full-text search with GIN indexing, `ts_rank_cd` relevance ranking, and contextual snippet extraction scoped strictly to user tenant boundaries.
- Designed a zero-knowledge guest sharing engine using SHA-256 hashed access tokens with expiration, instant revocation, and scoped view-only permissions.

Certifications

Google Cloud Skills Boost – Vertex AI: Prompt Design — Jul 2025