

Vishal Jadeja

+91 7359907544 | vishaljadeja.work@gmail.com | linkedin.com/in/vishal-jadeja | github.com/vishal-jadeja | leetcode.com/u/vishaljadeja

TECHNICAL SKILLS

Languages: JavaScript (ES6+), TypeScript, Python, SQL, C++

Frontend: React.js, Next.js, Redux Toolkit (RTK Query), WebSocket UI, HTML5, CSS3

AI/ML & Tooling: LLM API integration (Anthropic Claude, OpenAI, Gemini, Groq, Cerebras), RAG pipelines, LLM fallback chains, pgvector, MongoDB Atlas \$vectorSearch

Backend: Node.js, Express.js, RESTful API Design, WebSockets, Socket.io, JWT, OAuth 2.0 (PKCE)

Databases: PostgreSQL (indexing, partitioning, RLS), MongoDB (aggregation, compound indexing), MySQL, Supabase

Infrastructure: Redis, AWS SQS, BullMQ, Docker, GitHub Actions, CI/CD, Vercel

Security: RBAC, HMAC-SHA256, XSS & NoSQL injection prevention, Rate limiting

CS Fundamentals: Data Structures & Algorithms, OOP, Distributed Systems, System Design, Design Patterns

PROFESSIONAL EXPERIENCE

Full Stack Developer

Aug 2024 – Present

Glitchover (Diasimos Private Limited) | B2C Gaming Platform

Onsite, Ahmedabad

- Architected a real-time tournament system serving **5,000+ concurrent users** — migrated from REST polling to WebSocket-based architecture, cutting API load by **60%** and achieving **sub-200ms** latency at peak traffic.
- Architected a resilient **3-provider LLM fallback chain** (Gemini → Cerebras → Groq) for a real-time AI support bot — quota-aware skipping, per-provider JSON-repair retries, and graceful human-agent handoff on exhaustion, eliminating single-provider dependency for uninterrupted support coverage.
- Built a **vector search RAG pipeline** on MongoDB Atlas (\$vectorSearch) with **1,024-d L2-normalized** embeddings and a **0.55** confidence gate — low-confidence queries auto-escalate to human agents, grounding all LLM responses and eliminating hallucinated support answers.
- Architected an **RTK Query cache-aside layer** across **35+ code-split endpoint files** with 40 tag types and targeted cache invalidation — implemented optimistic mutations with automatic rollback across 8+ tournament flows for instant, race-free UI feedback.
- Reduced MongoDB query latency by **45%** for high-frequency leaderboard reads through schema redesign, compound indexing, and aggregation pipeline optimization — identified and eliminated $O(n)$ bottlenecks under tournament load spikes.
- Engineered an 11-stage Express middleware pipeline implementing defense-in-depth — CORS, NoSQL injection prevention, XSS sanitization, two-tier rate limiting (**100** req/min global, 10 req/s on payment endpoints), JWT auth, and RBAC across **50+ API routes**.
- Designed and deployed **2 independent microservices** (payment gateway, Discord bot) — applied Factory Pattern to decouple provider-specific integration logic, enabling new API providers to be onboarded with **zero refactoring** and cutting future integration effort by **~60%**.

PROJECTS

Cadenz – Personal Branding & Knowledge Platform

Next.js, TypeScript, React, PostgreSQL, pgvector, Trigger.dev, Upstash Redis, OAuth 2.0

- Built a **unified AI API adapter layer** supporting 4 LLM providers (Anthropic Claude, OpenAI, Gemini, Groq) — abstracted provider-specific schemas behind a single interface using OOP Adapter Pattern; parallel execution via Promise.all with streaming; **zero consumer-side changes** required to onboard a new provider.
- Engineered a fault-tolerant async job system using Trigger.dev — all AI generation, publish, and email jobs queued with **retry logic (3 attempts)**; API routes never block on external calls; system degrades gracefully under provider outages.
- Implemented a semantic search pipeline using **pgvector** — embedding jobs run asynchronously at low priority; vector queries scoped strictly to the requesting user's embeddings to prevent cross-tenant data leakage.
- Designed and built the full Next.js frontend — multi-page app with streaming AI responses rendered in real time, a provider-agnostic content generation UI, and semantic note search backed by **pgvector** embeddings.

EDUCATION

L. J. Institute of Engineering and Technology

Ahmedabad, India

Bachelor of Engineering in Computer Science and Engineering | CGPA: 8.48/10.0

Nov 2021 – July 2025