

Anubhav Joshi

Jaipur, India | magicalfizz@gmail.com | github.com/anubhav-qt | linkedin.com/in/anubhav-qt

Backend Engineer

SUMMARY

CS graduate (2026) with hands-on experience designing and shipping backend systems in production: async pipelines, microservices, and databases under real concurrency and rate-limit constraints. Sole architect of an async CQRS pipeline serving 11,000+ items from PostgreSQL in under 50ms, and built a 28-module NestJS marketplace backend handling payments, invoicing, and fulfillment. Comfortable owning a system end to end: API design, data modeling, and containerized deployment.

TECHNICAL SKILLS

Languages: Python, TypeScript

AI & Frameworks: FastAPI, PyTorch, Google ADK, LangGraph, MCP

Data & Persistence: PostgreSQL, Neon, Redis, pgvector / Pinecone

Cloud & DevOps: Google Cloud, Docker, Vercel, Render

UI/UX: Next.js, React Native

EXPERIENCE

Co-Founder & CTO - Anchorate

Jan 2026 - Present

- Designed and built Anchor8, a policy/orchestration layer sitting between AI agents and their tools: PII redaction, prompt-injection detection, and anomaly detection over an append-only audit log using pgvector similarity search, packaged as a PyPI SDK with a 3-line LangChain integration.
- Built a DualJuror escalation system where two LLM agents argue opposing risk positions and vote on ambiguous tool calls, backed by fast heuristic and regex-based guards for low-latency checks.
- Built Cargonto, a multimodal pipeline turning invoice photos and packing lists into structured freight paperwork; owned AWS/GCP infrastructure, Docker, and CI/CD for both products in production.

SDE Intern - Blinkadz

Feb 2025 - Apr 2025

- Built a 12-agent orchestration pipeline on Google ADK and FastAPI that cut video ad creation time by 95%, designing clean agent-to-agent handoffs across the workflow.
- Automated LinkedIn ad campaign setup end-to-end through the LinkedIn Marketing API, handling rate limiting and retry logic for reliable service-to-service integration; wrote Playwright suites reaching 90% coverage.
- Benchmarked 20 LLMs across production use cases to select models for avatar and video generation, balancing cost, latency, and output quality.

PROJECTS

Spoyn - CQRS Pipeline & Quota Governor

- Sole architect (55 ADRs) of an async generation pipeline built with FastAPI that keeps expensive LLM calls entirely off the read path, serving 11,499+ pre-computed cards from PostgreSQL in under 50ms.
- Designed a 2D quota governor (API key x model) with per-key fallback ladders and per-cell serialization across 16 free-tier keys, sustaining 74.6 cards/minute peak generation throughput without tripping rate limits.

Paribelle - Multi-Vendor Marketplace Platform

- Built a ~28-module NestJS/TypeORM microservices backend (vendor KYC, GST/HSN-compliant invoicing, Razorpay payments, bookings, promotions, search, wallet ledger) for a multi-vendor marketplace.
- Fixed a stock overselling race condition by making every reservation a single conditional UPDATE in the same transaction as the sales-count increment, closing the gap between a stock check and a decrement.

Trotter - Deterministic Quantitative Stock Analysis Engine

- Built a deterministic TypeScript core computing momentum, valuation, volume, FinBERT sentiment, and volatility scores directly from market data across three time horizons, no LLM in the calculation path.
- Layered Gemini 2.5 Flash on top for narrative generation and Tavily-grounded target estimates, capped to a small bounded shift so LLM output can never override the underlying deterministic score.

EDUCATION

B.Tech, Computer Science (Artificial Intelligence and Machine Learning) - Manipal University Jaipur

Aug 2022 - Jul 2026 | CGPA: 9.28/10

Dean's List and Student Excellence Award, all 8 semesters. Vice Chair, IEEE GRSS student chapter.