Aditya Dutt Pandey

AI & Agents Engineer

adityapandeyadp@gmail.com · +91-9415379379
lab.adpandey.dev · github.com/Adityaadpandey
linkedin.com/in/adpandeyadp

I build AI agents that hold up in production. tunyt runs events with a team of agents at **485K requests a week**, and **zerd** (zerd.si), my Rust coding agent, won't mark a task done until it has proved it.

Education

Lovely Professional University · B.Tech, Computer Science (AI & ML) · CGPA 8.39 — 2024 – 2028

Experience

tunyt · Co-Founder & CTO · tunyt.com — Dec 2025 – Present

Agentic events platform: one AI agent per part of an event. TypeScript, PostgreSQL, Redis.

- Built the first version solo (agents, real-time backend, UI). Now own the architecture and hiring, and lead a team of five across mobile, web, AI infrastructure and design.
- Designed one agent per event function (booking, ticketing, vendors, run of show), with analytics and AI jobs on background workers, off the request path.
- 10K+ sign-ups in month one with no marketing; ~3,000 paid ticket buyers; repeat hosts include Coding Blocks LPU, Trendians and Seen Kya Hai.
- Scaled weekly requests **7×** (69K → 485K, Jun–Aug 2026) while cutting latency **110 → 59 ms** and errors **2.84% → 0.32%**, by instrumenting every route first and fixing what the data showed.

Pokus · Software Development Intern — Jan 2026 – Feb 2026

- Built production backend pipelines for an AI agent that books, orders food and makes calls for users, and worked on its learning loop so behaviour changes shipped faster.

Lovely Professional University · Research Intern — Jun 2026 – Jul 2026

- Two patent applications filed (hardware-anchored attestation for AI agent decisions; cognitive-emotional adaptive intervention) and two research papers co-authored.

Projects

zerd, a coding agent that proves its work · zerd.si · GitHub — 2026

- Live at **zerd.si** as a single Rust binary with one-line installers for macOS, Linux and Windows. I use it daily on tunyt; it also did Veris's contract migration to Base.
- Before calling a task done it runs the repo's checks on changed files, reruns them on the pre-change tree to separate old failures from new ones, has a second model review the diff while a verifier tries to disprove each finding, and checks every cited file:line on disk.
- Routes each request to a cheap or strong model without an extra model call and fails over across providers on quota errors. Safety is enforced in code: confined file tools, a shell path scan and a macOS sandbox.
- Tested against real PostgreSQL and Redis, not mocks: on a 15-scenario payments suite (worker kill -9, 25 concurrent-claim races) one run surfaced 12 real bugs.

Speculative RAG, voice Q&A in four Indian languages · GitHub — 2026

- Answers spoken questions in Hindi, English, Tamil and Urdu in 20–30 ms on a 512 MB server with no LLM on the request path (Sarvam STT → BM25 on Qdrant → extractive answer).
- On a 100-query Hindi eval (MSMARCO-XI, half unanswerable), an async LLM recovery chain lifts correctness 57% → 80% and cuts false refusals 30% → 12%, but raises false confidence 34% → 66%, so it is reported apart from the fast path.

More at lab.adpandey.dev: PHANTOM (multi-agent RL for AI security); Bubble (macOS app for coding agents); Veris (hardware-signed photos).

Achievements & Community

- **Google Solution Challenge 2026:** Top 14 of 85,000+ registered teams, with Veris.
- **11+ hackathon wins**, including LPU InnoTek 9th Innovation Expo (1st, Student Incubation, Apr 2026), Xpecto'26 IIT Mandi (1st, Brand Wars, Mar 2026), Syntechs 2026 by CodeChef × Kreo (Jan 2026) and IIT BHU.
- **Teaching:** backend and DevOps to 100+ students at Learner's Arc; guest lectures on system design; open-source mentor, GSSoC 2025.
- **Writing:** 24 engineering write-ups on scaling, incidents and agents at lab.adpandey.dev/writing.

Skills

Strongest	Python, TypeScript, LLM agents, RAG, PostgreSQL
Also shipped	Rust, FastAPI, Qdrant, BM25, model routing, Redis, Docker
Familiar	PyTorch, LangGraph, multi-agent RL (GRPO)