

TANISHQ DASARI

Backend Engineer

Solapur, India | tanishqdasari2004@gmail.com | github.com/Tanishq1030 | linkedin.com/in/tanishq-dasari-15697231a

EXPERIENCE

FINOS Labs — Open Source Contributor

Feb 2026 – Present

- Shipped a merged 18-file, 1,659-line pull request to *dtcch-2026-stochastic-blade-runners* (FINOS Labs, DTCC Innovate Hackathon), adding CALM architecture discovery tools (`list_calm_architectures`, `get_calm_architecture`, `search_calm_architectures`) and automated threat-modeling utilities that map architecture components to governance mitigations.
- Fixed Windows-specific MCP server pipe I/O deadlocks by replacing blocking `select.select` calls with a threaded queue pattern, and added one-command pytest tooling for the team.
- Filed a feature proposal to *finos/legend-studio* (109-star FINOS repo) for real-time streaming data visualization via Perspective, including a working prototype, as part of a GSoC 2026 proposal.

Google Developer Groups — AI/ML Lead

Sept 2024 – Aug 2025

- Organized 10+ workshops on LLM systems, cloud infrastructure, and backend AI pipelines for 200+ students.

PROJECTS

Anchor — Static Analysis & Governance CLI for AI Codebases

github.com/AnimusLab/Anchor · pypi.org/project/anchor-audit · [Zenodo preprint: 10.5281/zenodo.19734724](https://zenodo.org/record/105281/zenodo.19734724)

AI codebases deployed in production have no commit-time audit trail — ungoverned LLM API calls, unencrypted vector stores, and unsandboxed subprocesses ship with no regulatory mapping. Anchor catches these at scan time and maps each finding to EU AI Act Art. 12, NIST AI RMF, and OWASP LLM Top 10.

- Validated against real codebases — 11 unencrypted vector-store writes in FINOS's *architecture-as-code* (684 files), 12 findings in HuggingFace Hub, and governance-drift findings in Django, LangChain, CrewAI, and OpenSpiel; each rule carries citation IDs mapping to EU AI Act Art. 12, NIST AI RMF, and OWASP LLM Top 10 — published as a Zenodo preprint.
- Ran Anchor against its own codebase as a precision test: an initial scan returned 14 false positives caused by regex patterns matching help text and detection-pattern strings rather than real code; root-caused the issue and shipped a "balanced line prefix" pattern-matching fix that reduced false positives to zero across the same scan.
- Set up CI workflows (GitHub Actions), a test suite (pytest), and a tamper-evident governance lockfile (SHA-256-sealed rule config) so rule changes can't silently bypass review.

Nexus — AI-Powered Microblogging Platform

github.com/ronibhakta1/NEXXUS · nexus.vercel.app

Social platforms moderate reactively — harmful content spreads before it is flagged. Nexus moderates at write-time, analyzing sentiment before a post goes live and prompting the user to rephrase rather than removing content after the fact.

- Built a multi-service backend: TypeScript/Hono on Cloudflare Workers with Prisma + PostgreSQL for the core API, and a Python/FastAPI AI layer integrating Weaviate vector storage, VADER sentiment analysis, Azure OpenAI embeddings, and a LangChain suggestion chain — deployed live with a React frontend on Vercel.
- Integrated Gemini 1.5 Pro for AI post refinement with dual API key failover (automatic retry on 429/503), Weaviate connection retry logic with graceful degradation, and JWT-authenticated edge routing on Cloudflare Workers.

TECHNICAL SKILLS

Languages: Python, Rust, Java, C++, TypeScript, SQL

Backend: FastAPI, PostgreSQL, Redis, Docker, WebSockets, CI/CD (GitHub Actions), Linux

Static Analysis & Tooling: Tree-sitter (AST parsing), WASM/WasmEdge, regex-based pattern matching, pytest

AI / ML: LangChain, RAG pipelines, vector search, PyTorch basics

Cloud: Google Cloud Platform, Azure (OpenAI, Weaviate), Cloudflare Workers, Vercel

EDUCATION

Dr. Babasaheb Ambedkar Technological University

Expected Aug 2026

B.Tech, Computer Engineering