

SAM ADRIAN

AI Systems Engineer

CONTACT

Email itsmesamuel.2007@gmail.com
GitHub github.com/x-TheFox
LinkedIn [linkedin.com/in/xthefox](https://www.linkedin.com/in/xthefox)
Portfolio samadrian.pro
Location Chennai, India

TECHNICAL SKILLS

TypeScript / Next.js
React / Tailwind CSS
Python / FastAPI
C++ / Systems Prog.

LLM Integration

Docker / CI-CD

PostgreSQL / MongoDB

Cloud (GCP / Vercel)

AI / ML STACK

OpenAI / Groq / Google GenAI
LangChain / Llama-Index
HuggingFace Transformers
Pinecone (Vector DB)
LiveKit (STT/LLM/TTS)
K-Means / Prompt Engineering

INFRASTRUCTURE

Docker / Docker Compose
GitHub Actions (CI/CD)
Cloudflare Workers
Vercel / Render / Netlify
Linux Server Admin

EDUCATION

B.E. Computer Science
Panimalar Engineering College
Anna University, Chennai

PLATFORMS

HuggingFace: EdenCorp
Kaggle: xthefox

Profile Summary *Overview*

Self-taught engineer with 5+ years of building across the stack — from low-level C++ distributed systems to AI-augmented web platforms. Founded EdenCORP, an AI integration startup shipping production tools for education and enterprise. Built and deployed a Raft-consensus distributed cache in C++ with full CI/CD, sanitizers, and codecov. Designed persona-adaptive portfolio systems using K-means clustering + GPT classification. Integrating LLMs (OpenAI, Groq, Llama) into real products — not demos.

Professional Experience *Career*

EdenCORP **Founder & Lead Engineer** 2025 — Present

- Founded AI integration and product agency designing automation infrastructure for education, logistics, and enterprise clients. Built ORCHARD, a cross-industry AI platform targeting 10+ verticals, and PRISM, a campus digital infrastructure system with elastic auto-scaling.
- Architected EdenCORP website (Next.js 16, Tailwind CSS 4, Vercel) with custom particle animations, 3D effects, and premium dark-mode design language. Published EdenCorp/tts text-to-speech model on HuggingFace.
- Managing end-to-end product lifecycle: system design, AI pipeline integration (Groq, OpenAI, Google GenAI), deployment, and infrastructure. Solo operator maintaining codebase, CI/CD, and production uptime.

Key Projects *Technical Work*

distributed-cache **Systems Engineer** Dec 2025 — Apr 2026

- Built a production-grade distributed in-memory cache in C++17 implementing Raft consensus for leader election and log replication, consistent hashing for data sharding, and a Reactor-pattern networking layer with epoll (Linux) and kqueue (macOS) for non-blocking I/O.
- Implemented RESP (Redis-compatible), Memcached text protocol, and HTTP protocol parsers supporting pipelining and request multiplexing. Designed a striped locking LRU/LFU eviction cache with configurable policies.
- Built full cluster orchestration: membership service with heartbeat-based failure detection, automatic shard rebalancing, key routing with MOVED redirection, and Write-Ahead Logging with snapshot-based recovery.
- Shipped with 21 unit/integration test files (GoogleTest), GitHub Actions CI matrix (Ubuntu/macOS, ASan/UBSan, cppcheck, clang-format, Codecov), multi-stage Docker builds, and 3-node docker-compose cluster demo.

The-Adaptive-Developer **Full-Stack Engineer** Jan 2026 — Apr 2026

- Engineered an AI-augmented portfolio (Next.js 16, React 19, TypeScript 5) that detects visitor personas via K-means behavioral clustering + OpenRouter GPT classification, dynamically adapting the entire UI across 6 persona types (Recruiter, CTO, Engineer, Designer, Gamer, Curious).
- Integrated Groq-powered streaming chatbot (Llama 3.1 8B at 840 tok/sec), Notion SDK as headless CMS, and an AI interviewer that processes structured intake form responses. Implemented GDPR-compliant edge behavior logging with 7-day retention.
- Tech stack: Neon PostgreSQL (Drizzle ORM), Three.js / React Three Fiber for 3D, Framer Motion, Vercel AI SDK, and deployed on Vercel at samadrian.pro.

Resume Optimizer *AI Systems Engineer* Jan 2026

- Built a multi-agent AI system that autonomously analyzes job descriptions and generates optimized resumes using deep web research (Firecrawl, Brave, DuckDuckGo)
- Designed a “Success Profile” generation pipeline to model ideal candidate attributes and perform structured gap analysis (0–100 scoring)
- Implemented in-place PDF/DOCX editing with formatting preservation + OCR fallback for complex resumes
- Deployed and adopted as a college-wide resume optimization tool, used by students across multiple departments and integrated into soft skills training workflows
- Generated automated outputs including optimized resumes, cover letters, and interview preparation material

Etro-Bot *AI/ML Engineer* Jan 2023 — Jan 2026

- Architected a modular ChatGPT-powered Discord bot (Python, pycord) with 9 service cogs: GPT-3 conversations, DALL-E image generation, web search, audio transcription, DeepL translation, content moderation, document indexing, and prompt optimization.
- Implemented a RAG pipeline using Pinecone vector embeddings for semantic document retrieval, with queue-based message processing for concurrent request handling. Dockerized with docker-compose for production deployment.
- Integrated HuggingFace Transformers for language detection, YouTube transcript extraction for content analysis, and role-based access control with persistent SQLite conversation history.

Learnzy *Full-Stack Engineer* Feb 2026 — Mar 2026

- Built a full-stack adaptive learning platform (Next.js 16, Express.js, MongoDB, Firebase Auth) designed for students with ADHD and dyslexia, featuring AI-powered quizzes that dynamically adjust difficulty based on student performance.
- Implemented dyslexia-friendly reading tools (customizable fonts, color overlays, AI text simplification via Groq Llama 3.3 70B), text-to-speech via browser APIs, and ADHD-aware UX with focus timers and minimal-distraction interface.
- Deployed with Vercel (frontend) and Render (backend), with Google OAuth and email authentication via Firebase.

LiveKIT Voice Agent *Real-Time AI Engineer* Apr 2026

- Built a Discord voice channel bot connecting to LiveKit for real-time AI voice interaction, integrating Deepgram STT, OpenAI LLM, Cartesia TTS, and Silero VAD in a streaming voice pipeline with < 500ms end-to-end latency.

Apeiron *Algorithm Engineer* Apr 2026

- Implemented a Gomoku AI engine in Python using minimax search with alpha-beta pruning on an infinite sparse board representation. Designed a pattern-based positional evaluator with tactical shortcuts for immediate win/block detection. Comprehensive pytest test suite.

Achievements

IEEEExtreme 19.0 — Ranked #59 globally (Top 0.3%) among 19,000+ participants

Built 30+ production-grade repositories across C++, Python, TypeScript, and Java

Selected Impact

- Ranked **Top 0.3% globally (59/19,000+)** in IEEEExtreme 19.0
- Built and deployed **30+ production-grade systems** across AI, distributed systems, and backend infrastructure
- Developed AI tools adopted by **college-wide student body across multiple departments**
- Designed distributed systems using **RAFT consensus and consistent hashing**, comparable to production-grade architectures

