

ADRIAN KIMANI

Full Stack Engineer

Nairobi, Kenya · **Open to remote and relocation** · UTC+3 (full overlap with CET, mornings with US Eastern)

adriankimani5@gmail.com · +254 727 323 382 · adriankimani.dev · github.com/saintAdrian7 · linkedin.com/in/adrian-kimani-840468308

SUMMARY

Full stack engineer with 2+ years shipping production web applications in legal tech and AI-powered education. Builds end to end on Next.js, React, and TypeScript over NestJS, PostgreSQL/Supabase, and AWS — authentication, Stripe billing, real-time messaging, and LLM/RAG integrations built with LangChain. Delivered 10 client products from zero to launch at Tison.io.

TECHNICAL SKILLS

Languages: TypeScript, JavaScript, Python, SQL

Frontend: React, Next.js, Tailwind CSS

Backend: NestJS, Node.js, Express, REST APIs, PostgreSQL, Supabase, MongoDB

AI & LLM Systems: LLM integration (OpenAI, Groq, Gemini), LangChain, LangSmith, RAG pipelines, vector search (pgvector, ChromaDB), prompt engineering, multi-step LLM orchestration

Cloud & DevOps: AWS (Cognito, Lambda, S3, DynamoDB), Azure, Docker, CI/CD, Git

Integrations: Stripe, SendBird, Puppeteer, Leaflet.js

EXPERIENCE

Full Stack Engineer, Tison.io

Aug 2024 – Present

Remote · Consultancy arm of DirectEd Development, building products for early-stage startups

- Shipped 10 client websites and web applications from zero to launch on teams of 1–5, owning React/Next.js UI through NestJS APIs and PostgreSQL schema design.
- Built internal workflow automations for the company's own delivery operations, replacing recurring manual processes with scheduled jobs and third-party integrations.
- Sustained 1,021 GitHub contributions over the last year across 5 client repositories, working to reviewed-PR and branch-protection standards.

Full Stack Engineer, LegalRescue.ai

Aug 2024 – Oct 2025

Remote · Client engagement via Tison.io · legalrescue.ai — AI legal-tech platform matching attorneys to clients, built 0 to 1 pre-launch

- Built the platform's authentication from scratch on AWS Cognito across 3 providers (email, Google, Microsoft OAuth) with email verification, password reset, and JWT session management.
- Shipped the subscription system end to end on Stripe: dynamic pricing driven by county/ZIP selection, practice area, and firm type, plus discounts, trials, referral tracking, and webhook lifecycle handling.
- Designed the 8-stage case-management workflow with automated conflict-of-interest checks, staged client-information disclosure, and validation logic.
- Built a pricing calculator on Leaflet.js with live validation, covering all 3,100+ US counties and 41,000+ ZIP codes.
- Built per-case real-time messaging on SendBird with multi-case channel management and file sharing, responsive across web and mobile.
- Integrated EngageBay CRM with progressive waitlist syncing and custom tagging, carrying 30+ pre-launch signups into the sales pipeline automatically.

Backend Engineer, Cleverguide.ai

Oct 2025 – Feb 2026

Remote · Client engagement via Tison.io · cleverguide.ai — AI education-consulting platform, built 0 to 1

- Owned the NestJS backend as one of 7 engineers, authoring 170 of 193 commits (88%) and 25 merged PRs, and shipping 21 REST endpoints across 5 controllers with Swagger docs, JWT auth, and 3-tier rate limiting (10/s, 120/min, 400/15min).
- Established the backend test suite from zero — 192 unit and 21 end-to-end tests across 15 files — behind a committed 80% coverage gate enforced in CI on every PR, reaching 87.7% line coverage on instrumented services.
- Built a Puppeteer PDF generator rendering a 12-section branded A4 report from 7 concurrent Supabase queries, hardened for containerized headless Chrome with a 90-second render timeout.
- Designed a rules-based article recommendation engine (1,132 LOC across 5 modules) with a 4-stage fuzzy institution matcher — exact, Jaccard token overlap at 0.6, acronym, and keyword containment.
- Integrated the Python RAG recommendation service over HTTP, owning orchestration across 5 sequential AI stages with timeout, retry/backoff, and persistence handling.

PROJECTS

Study RAG Assistant — Grounded Document Q&A

github.com/saintAdrian7/RAG · Python, OpenAI, Groq, Google Gemini, YAML-driven prompts

- Built a retrieval pipeline answering strictly from indexed documents, quoting the source span for every claim and refusing when the corpus lacks coverage.
- Implemented multi-turn memory with query rewriting, and abstracted providers across OpenAI, Groq, and Gemini with key-based fallback and YAML-tunable prompts.

EDUCATION

Software Engineering Certificate, DirectEd Development

Remote accelerator training African engineers for global teams

2024