

JOSEPH DAUDU

PRODUCT ENGINEER | AI & SOFTWARE SYSTEMS ARCHITECT | DESIGN ENGINEER

Abuja, Nigeria | Blessedgloves94@gmail.com | 09121980784 | [daudu.framer.website](#) | [GitHub](#) | [LinkedIn](#)

PROFILE

Product Engineer with 3+ years building AI-powered tools, distributed backend systems, and scalable web applications. Designed and shipped production-grade platforms integrating LLMs (OpenAI, Claude, Gemini), RAG pipelines, Web3 payment infrastructure, and microservices architectures. Currently serving as IT Manager at Subtle Innovative Services while completing a Computer Engineering degree at the University of Benin.

WORK EXPERIENCE

IT Manager | Subtle Innovative Services (Educational Consultancy) — Lagos Jan 2024 – Present

- Designed, built, and launched the company's primary client-facing website end-to-end, covering architecture, development, and ongoing maintenance. The platform supports student discovery, service enquiries, and lead generation across six international study destinations.
- Engineered the front-end and back-end infrastructure using React, TypeScript, and Node.js/Express with PostgreSQL data models, ensuring reliable performance for student and client traffic.
- Maintained 99.8% platform uptime by architecting cloud infrastructure on AWS using Docker, Nginx, and automated CI/CD pipelines via GitHub Actions.
- Designed and shipped the UI prototype for the Subtle Travel mobile app (iOS & Android), a cross-platform hospitality product covering resort discovery, flight/hotel booking, and in-app travel management.
- Built internal automation tooling and Python analytics scripts that improved operational efficiency and informed data-driven product decisions across the consultancy's service lines.
- Achieved a 95% client satisfaction rate as measured by post-engagement NPS surveys, by managing end-to-end technology strategy, infrastructure, and platform maintenance.

Software Engineering Intern | Khemsafe Limited — Abuja 2024 – 2025

- Reduced critical bug resolution time by 35% as measured by internal incident tracking logs, by contributing to backend logic implementation and systematic code optimisation across production workflows.
- Improved system reliability across 3 core internal services as measured by automated regression test pass rates, by writing structured debugging scripts and implementing Python-based test coverage.
- Accelerated feature delivery by 20% as measured by sprint completion rate, by applying data structures and algorithm optimisations that reduced query complexity in the primary database layer.

Contract Software Developer | Emad Services — Abuja Jan 2023 – Mar 2023

- Designed, built, and maintained Emad's client-facing website, handling front-end development, back-end API integration, and ongoing performance improvements for their service portfolio.
- Reduced application load time by 30% as measured by Lighthouse performance scores, by refactoring backend data pipelines and optimising MySQL queries using Python profiling tools.
- Saved the engineering team 8 hours per week by building Python automation scripts that streamlined repetitive build and deployment workflows.

UI/UX Designer & Developer | Kelly Stores — Abuja Dec 2024 – Present

- Designed and built a full e-commerce website for Kelly Stores, an online clothing brand, covering UX research, Figma wireframing, front-end development, and production deployment.
- Increased e-commerce conversion rates by 15% as measured by Google Analytics checkout funnel data, by redesigning the user journey with evidence-based UX patterns and structured information hierarchy.

- Improved user retention by 20% as measured by 30-day cohort analysis, by conducting A/B tests and using Python (pandas, NumPy) for statistical significance testing on user journey drop-off data.

PROJECTS

Gitty — AI Codebase Intelligence Tool — *Full-Stack, AI Engineering | In Progress*

- Building a Next.js 14 developer tool that combines a codebase map, AI-powered chat with repo-aware RAG retrieval, and an automated PR code review engine, designed to give developers deep contextual understanding of any codebase.
- Integrating Semantic Chunk Compression (SCC) as Gitty's core RAG retrieval service, replacing naive file-chunking with AST-aware semantic units and canonical deduplication, enabling retrieval that reasons at the function and class level rather than arbitrary text windows.
- Architecting a durable review dispatch system (ReviewRunDispatcher / ReviewExecutor) with context window estimation and chunk trimming to prevent token overflow on large pull requests.
- Implementing MeiliSearch and pgvector hybrid retrieval with an LLM pre-check classifier that skips RAG on conversational queries, reducing latency and cost on non-retrieval turns.

Stack: Next.js 14, TypeScript, PostgreSQL, Prisma, Clerk, OpenAI API, MeiliSearch, pgvector, Redis, pnpm

Semantic Chunk Compression (SCC) — AI Infrastructure / Systems Design — *Personal R&D | In Progress*

- Designing and building a code intelligence engine that parses repositories into AST-aware semantic units using Tree-sitter, deduplicates them via canonical fingerprinting, and stores a single compressed chunk per unique piece of logic, eliminating redundant embeddings and enabling semantic-level version diffs.
- Achieving a 38.5% embedding reduction on a representative mid-size repo (18,904 chunks across 4,281 files → 11,622 unique chunks after deduplication), with near-zero recomputation cost for repeated retrievals.
- Designing a three-phase build roadmap: Phase 1 (current): RAG-focused chunk store; Phase 2: semantic diff and function-level history on top of Git; Phase 3: native chunk manifest store replacing Git's file snapshot model.

Stack: Node.js, TypeScript, Tree-sitter, PostgreSQL, pgvector, Redis, Express

Dev Help — AI Developer Assistant Platform — *Full-Stack, AI Engineering*

- Extended LibreChat's open-source foundation into a specialised AI developer assistant, adding a custom RAG pipeline, multi-model support (OpenAI, Claude, Gemini, and local models via Ollama), and real-time code execution across 20+ languages.
- Built a custom embedding pipeline using Google text-embedding-004 stored in pgVector, reducing RAG retrieval latency to under 100ms and enabling semantic document search not present in the base platform.
- Scaled the platform to 200 concurrent users (validated by 2,066 automated test cases) by expanding the backend to an 8-microservice architecture with MongoDB, PostgreSQL, Redis, and MeiliSearch.

Research paper submitted for publication: "DevHelp: Design and Implementation of an AI-Powered Chatbot Using LLM-Based NLP and Remote Code Execution APIs" — University of Benin

Stack: Node.js, React 18, TypeScript, OpenAI API, Claude API, Gemini API, LangChain, pgVector, Redis, MongoDB, Docker, AWS S3, MeiliSearch

Nebula — Multi-Chain Crypto Payment Infrastructure — *Full-Stack, Web3, Systems Engineering | In Progress*

- Building a production-ready payment API across 5 blockchain networks (Ethereum, Polygon, Base, Arbitrum, Sepolia) with 15–30s live detection latency via custom Ethers.js workers polling every 15 seconds.
- Delivering API response times under 200ms and ~80% cache hit rates via Redis-cached analytics workers and 5 Bull queue background processors running 24/7.
- Engineering a full SaaS platform spanning 46 REST API endpoints and 16 MongoDB collections, built entirely solo.

Stack: Next.js 14, Node.js, Express, MongoDB, Redis, Ethers.js, Bull, HMAC-SHA256, bcrypt, JWT, Vercel

EDUCATION

B.Eng. Computer Engineering

Completed

University of Benin — Relevant coursework: Data Structures & Algorithms, Web Development, User-Centered Design, Systems Engineering

Senior Secondary School Certificate (SSCE)

Completed

Lapan Secondary School

CERTIFICATIONS

Building with the Claude API — Anthropic Academy	2026
Full Stack Web Development Certificate — Udemy	2024
Front-End Web Development Certification — freeCodeCamp	2023
Google UX Design Certificate — Coursera	2023