ANDREW BOLAJI

AI Engineer | LLM Applications, Voice Agents & RAG | Full-Stack | Nine Years Production Software

Spring, Texas | (832) 602-3814 | andrewbolaji@gmail.com

linkedin.com/in/andrewbolaji • github.com/andrewbolaji • andrewbolaji.com

PROFESSIONAL SUMMARY

Nine years building and keeping production software running, from assembly-line control systems to enterprise financial platforms. That is what the AI work is built on. I assume the model can be fooled, so the things that stop it sit underneath it: access checked before retrieval, actions that wait for a person to approve them, and evals that fail the build instead of warning. Two AI systems are live in production and you can try both.

TECHNICAL SKILLS

AI & LLM: LLM application development (Anthropic / Claude, OpenAI) • Voice agents (Retell AI, VAPI, Twilio) • Retrieval-augmented generation (RAG) • Prompt engineering, function & tool calling • Agent orchestration (custom agent loop, tool registry, bounded steps, execution gating) • Embeddings & vector search (Qdrant, local BGE) • Prompt-injection defense • PII redaction • Red-team evaluation • Evaluation & guardrail design • AI-augmented development

Languages & Data: Python • SQL (PostgreSQL, MySQL, SQL Server) • TypeScript / JavaScript • Dart • Java • XML

Backend & Cloud: FastAPI • Supabase (Postgres, RLS, Edge Functions) • Firebase • Stripe • REST APIs • Webhooks • AWS (EKS, Redshift, S3, ECR, IAM) • Terraform • Kubernetes • Helm • Docker • dbt • Linux

Frameworks & App Dev: React 19 • Flutter • Node • Next.js • Tailwind • full-stack web and mobile (iOS)

Testing, CI/CD & Infra: Vitest • Flutter test • security-rules and RLS isolation testing against live databases • CI gates that block the build rather than warn • Eval harness design (golden datasets, adversarial suites, LLM-as-judge) • Docker multi-stage builds, pinned images • GitHub Actions, Jenkins • Adversarial security review

Enterprise Delivery: Tier-2 / Tier-3 escalation • Requirements elicitation & functional specs • Root-cause & impact analysis • UAT • SaaS implementation • Agile / Scrum

SELECTED PROJECTS

Designed, built, tested, and shipped end to end. Spec-driven, with adversarial security review at every milestone. Code and walkthroughs available on request.

- **Custos — Private AI Over Company Documents (live demo):** A private assistant that answers questions about a company's own handbooks, policies, and contracts, and cites the document behind every answer. Python and FastAPI, a local embedding model with self-hosted Qdrant so no document content leaves the customer's environment, Claude for generation, React and TypeScript streaming the reply. Access is checked at retrieval, so a file you are not cleared for is never a candidate for your answer, and personal data is masked whoever is asking. In red-team testing a poisoned document talked it into sending an email and it tried. The email never went, because tools with real-world effects cannot fire themselves. 230+ tests, 73 behavior evaluations, unauthorized action rate 0.

- **Reckon — Cloud Data Platform on AWS:** Three sources into a warehouse, modeled with dbt into Redshift, served through FastAPI to a React dashboard with an AI copilot that answers questions in plain English. Terraform manages 51 resources; the app ships to Kubernetes on EKS through Helm, with GitHub Actions building images into ECR. Data quality and freshness tests run inside the transformation, so a failed test stops the models below it from building. The dashboard goes stale loudly instead of wrong quietly. Deployed to real AWS and torn down again.

- **Porter — AI Voice Agent (live, callable at 515-608-4723):** Voice agent built on Retell AI with Claude and GPT for reasoning and Twilio for telephony. Holds context across a call, works out what is urgent, checks live availability, books the appointment, and sends an SMS confirmation. It will not book a slot it has not verified or quote a price it was not given, and when confidence drops it hands off to a human with a summary rather than inventing an answer. Call it and it answers.

- **Porter Evals — Eval Harness & CI Gate:** A containerized eval suite for Porter's reasoning layer, because a system prompt is the product and changing it can break behavior with nothing to catch it. 59 scripted call scenarios across eight categories, including guardrails, escalation, adversarial prompts, and brand safety. Deterministic checks where possible, an LLM judge only for open-ended answers with its reasoning logged. Wired into GitHub Actions so the build fails on any guardrail or adversarial violation. Back when the suite was 43 scenarios the first run scored 34, and three of those failures were bugs in my own scorer, including a check that flagged "we don't do landscaping" as an invented service because "don't do" contains "do". The agent was right and the test was wrong.

- **Junta — Two-Sided Marketplace with Escrow:** React 19, TypeScript, and Supabase (Postgres, row-level security, 37 migrations, 16 Edge Functions). Homeowners post one standardized scope, vetted contractors submit sealed bids, and funds are held through Stripe Connect destination charges with manual payouts, so the platform controls when money moves without ever holding it. Full job lifecycle through dispute resolution, location data stripped from photos in the browser, address withheld until a lead is paid, admin console for verification and audit. 528 tests, including 46 row-level-security isolation tests that run against the live database.

- **Roster — Compliance & Operations Platform (React 19, TypeScript, Supabase):** Operations software for regulated transitional-housing operators, built to a municipal ordinance with mandatory incident reporting and a five-year retention minimum. Hard DELETE is blocked in the database for residents, documents, and incidents; soft-delete with a retention window is the only path. Public form IPs are hashed, and every anonymous endpoint runs through a SECURITY DEFINER function that whitelists the exact fields it returns. 316 tests, 21 migrations. Running in production.

PROFESSIONAL EXPERIENCE

Co-Founder & Chief Technology Officer February 2026 – Present
AIntellect Solutions | *Des Moines, IA + Spring, TX*

Two-person firm building AI systems for small and mid-size businesses. I own everything technical: architecture, build, security review, and deployment.

- Shipped two products that are live and open to the public: an AI phone agent that answers and books jobs, and a private AI assistant that answers questions over a company's own documents
- Wrote the security model both products run on: access checked before retrieval, personal data masked by default, and an approval gate that holds even when the model itself has been compromised
- Scope client work, write the technical proposals, and set pricing against researched market rates

Senior Application Support Engineer October 2022 – Present
Vestmark | *Wakefield, MA (Remote)*

Technical liaison between enterprise financial-institution clients and engineering on the VestmarkONE wealth management platform, owning Tier-2 and Tier-3 escalations end to end.

- Debug production issues at the database and integration layer: SQL query optimization, XML configuration, REST API integrations, and SAML / SSO authentication flows on mission-critical financial systems
- Built Python automation that removed repetitive support work and cut time to first response
- Run root-cause and impact analysis across portfolio management and trading workflows, surfacing systemic issues and driving the permanent fix with engineering rather than the patch
- Reproduce client-reported defects and turn them into specifications and reproduction steps engineering acts on directly

Senior Software Support Engineer February 2021 – October 2022
Veracross | *Wakefield, MA (Remote)*

Backend support and solutions engineering for 300+ institutional customers across 20+ countries on an enterprise EdTech SaaS platform.

- Traced bugs through the full stack, from client-reported symptom to database root cause, and triaged into engineering by priority
- Designed and wrote custom SQL reporting and Python tooling to fulfill bulk data and reporting requests
- Troubleshot client API integrations end to end, isolating failures across authentication, payload, and data-mapping layers
- Installed and configured client database modules to customer specification

Software Support Engineer March 2019 – February 2021
SSI Schäfer | *Atlanta, GA*

- Maintained and debugged Java-based warehouse management systems processing millions of daily e-commerce transactions for enterprise clients including Walmart
- Wrote and ran SQL for diagnosis, data migrations, corrections, and optimizations across live databases
- Ran regression testing to verify builds and deployments before release

Software Developer September 2017 – March 2019
Honda Manufacturing of Indiana | *Greensburg, IN*

- Proposed, built, and debugged Java changes to an assembly-line control application running live manufacturing at a target of 1,000 cars per day
- Built and maintained real-time plant signboard systems, integrating databases with Moxa and DataPipe to show live line status across the floor
- Validated deployments and bug fixes by running automated simulations through Jenkins before release

EDUCATION

Bachelor of Science, Computer Science 2017
Georgia State University | *Atlanta, GA*