

SIDDHARTHA NAGULA

682-458-5291 | siddharthanagula3@gmail.com | [LinkedIn/siddharthanagula](https://www.linkedin.com/in/siddharthanagula) | [GitHub/siddharthanagula3](https://github.com/siddharthanagula3) | [Leetcode](https://leetcode.com/siddharthanagula3)

SUMMARY

Applied AI Engineer and AI-native product builder with an MS in Computer Science, focused on shipping LLM applications, AI agents, local-first workflows, and privacy-aware product demos. Founder and AI Product Builder at AGI Automation LLC, with hands-on work across AGI Workforce, AGI Agent Automation, and HxF AI involving agent orchestration, BYOK provider routing, RAG, memory, frontend UX, evaluations, and guardrails.

SKILLS

- **AI Engineering:** Generative AI, Large Language Models (LLMs), AI agents, agentic workflows, multi-agent orchestration- , tool calling, tool orchestration, prompt engineering, Retrieval-Augmented Generation (RAG), vector search, embeddings, context engineering, memory, LLM evaluations, guardrails, responsible AI, local LLMs, on-device inference
- **Implementation and Product Building:** proof of concept development, prototype-to-demo delivery, product demos, workflow discovery, solution design, API integrations, implementation planning, technical storytelling, release-readiness, privacy-aware AI, trust-boundary design
- **Programming and Platforms:** TypeScript, JavaScript, React, Next.js, Vite, Tailwind CSS, Radix UI, React Native, Expo, Swift, SwiftUI, iOS, Tauri, Rust exposure, Python, FastAPI, Express, PostgreSQL, SQLite, SQLCipher, Supabase, Neon, Clerk, Stripe, Sentry, Upstash Redis, Monaco Editor, Sandpack
- **APIs, Cloud, and Data:** REST APIs, API gateways, WebSocket, Server-Sent Events, JSON-RPC, Supabase Edge Functions, Netlify Functions, Row Level Security, Google Cloud Platform, AWS, Docker, Git, GitHub Actions
- **AI Providers and Tools:** OpenAI API, Anthropic Claude API, Google Gemini, Vertex AI, Ollama, LM Studio, xAI, Perplexity, DeepSeek, OpenRouter, Model Context Protocol (MCP), Claude Code, BYOK provider workflows, multi-provider routing
- **Testing, Security, and Quality:** Playwright, Vitest, Jest, ESLint, Sentry monitoring, HMAC signing, JWT auth, rate limits, origin controls, secret redaction, provider labels, consent flows, payload previews

WORK EXPERIENCE

AGI Automation LLC

Aug 2025 - Present

Founder and AI Product Builder

Austin, TX

- Built AGI Workforce, a multi-surface AI agent platform across desktop, web, mobile, command-line interface, Chrome extension, VS Code extension, shared provider packages, API gateway services, local LLM runtime, and provider adapters.
- Designed and implemented agentic AI workflows for local LLM chat, on-device model execution, provider routing, streaming responses, browser context capture, desktop automation, artifact handling, and privacy-aware send previews.
- Built Local Mode, Bring Your Own Key (BYOK), and Managed Cloud execution paths to support privacy-sensitive AI workflows, user-owned provider keys, enterprise laptop deployment positioning, and future managed compute.
- Built AGI Agent Automation, a Vite, React, TypeScript, Supabase, Netlify, and Stripe SaaS prototype for an AI employee marketplace, multi-agent chat, Vibe-style coding workspace, LLM provider proxies, billing, and workflow automation.
- Implemented Supabase auth and database patterns, Row Level Security migrations, Supabase Realtime chat concepts, Netlify serverless LLM proxies, Stripe checkout and webhook code, Sentry monitoring, and Playwright and Vitest test infrastructure.
- Built HxF AI, a native Swift and SwiftUI iOS health and fitness AI app with HealthKit, StoreKit 2 subscriptions, widgets, App Intents, secure storage, backend SSE streaming, and OpenAI, Claude, and Gemini integrations.
- Implemented customer-facing technical concepts including local-to-BYOK handoff, secret redaction, payload preview, provider labels, explicit consent messaging, privacy controls, and trust-boundary explanations.
- Used AI-native development workflows to ship across multiple product surfaces while managing architecture, testing, release-readiness, documentation, product tradeoffs, and implementation risks.

Compass Group USA

Jan 2024 - May 2025

Starbucks Shift Supervisor

Arlington, TX

- Supervised customer-facing store operations while completing a Master of Science in Computer Science, coordinating team execution, training, inventory, customer service, and peak-hour workflows.
- Improved shift productivity through standardized communication, task ownership, and operational handoffs across a 35-person team.
- Used transaction and customer-flow patterns to improve staffing decisions, throughput, and reporting visibility.
- Built Excel and Python-assisted reporting improvements to reduce manual reporting time and improve operational tracking.

Black Iron Fitness Studio

Jan 2020 - Jul 2023

Owner and Managing Partner

Hyderabad, India

- Managed customer acquisition, renewals, staff coordination, vendor relationships, marketing campaigns, member retention, and daily gym operations.
- Used customer segmentation, campaign analysis, churn tracking, and renewal workflows to improve retention and guide sales decisions.
- Implemented customer relationship management workflows to improve payment tracking, renewal follow-up, and operational consistency.

PROJECTS

AGI Workforce, Multi-Surface AI Agent Platform

- Built a multi-surface AI platform covering desktop, mobile, web, command-line interface, Chrome, VS Code, API gateway, local LLM runtime, and provider adapters.
- Designed local-first privacy, BYOK provider usage, managed-cloud waitlist, enterprise laptop deployment positioning, browser automation, desktop automation, and AI workflow orchestration.
- Technologies: TypeScript, React, Next.js, Vite, Tailwind CSS, Radix UI, Zustand, TanStack React Query, Zod, React Native, Expo, Tauri, Rust exposure, PostgreSQL, SQLite, SQLCipher, Supabase, Neon, Clerk, Stripe, WebSocket, Server-Sent Events, JSON-RPC, Playwright, Chrome Extension APIs, VS Code Extension APIs, OpenAI API, Anthropic Claude API, Google Gemini, Ollama, LM Studio.

AGI Agent Automation, AI Employee Marketplace and Multi-Agent SaaS Prototype

- Built a web SaaS prototype for AI employee marketplace, multi-agent chat, prompt-based agent catalog, Vibe-style coding workspace, document export, media generation, billing, and provider proxy workflows.
- Implemented marketplace, chat, mission control, support, billing, Supabase schema, Row Level Security, Netlify serverless functions, Sentry monitoring, and test coverage.
- Technologies: TypeScript, React, Vite, Supabase, PostgreSQL, Row Level Security, Supabase Realtime, Netlify Functions, Stripe, OpenAI API, Anthropic Claude API, Google Gemini, Perplexity, xAI Grok, Sentry, Upstash Redis, Sandpack, Monaco Editor, React Router, Recharts, Playwright, Vitest.

HxF AI, Native iOS Health and Fitness AI App

- Built a native iOS wellness app connecting Apple Health data, nutrition and workout tracking, AI chat, plan generation, food photo analysis, subscriptions, privacy controls, and legal/support pages.
- Implemented HealthKit consent, StoreKit 2 subscriptions, secure storage, backend authentication, SSE streaming chat, health data export and deletion, medical disclaimers, usage limits, and App Store readiness workflows.
- Technologies: Swift, SwiftUI, HealthKit, StoreKit 2, WidgetKit, App Intents, Vision, ARKit, AVFoundation, Speech, Keychain and secure storage concepts, Express, Supabase, PostgreSQL, OpenAI API, Anthropic Claude API, Google Gemini, Server-Sent Events, JWT, HMAC signing, Helmet, Joi, Fastlane, Vercel.

EDUCATION

University of Texas at Arlington

Master of Science, Computer Science

Vardhaman College of Engineering

Bachelor of Technology, Computer Science and Engineering

CERTIFICATIONS

- Google Associate Cloud Engineer: Google, May 2023 - May 2026
- Building with the Claude API: Anthropic, Issued Mar 2026
- Claude Code in Action: Anthropic, Issued Mar 2026
- Introduction to Model Context Protocol: Anthropic, Issued Mar 2026
- Model Context Protocol: Advanced Topics: Anthropic, Issued Mar 2026
- Introduction to Agent Skills: Anthropic, Issued Mar 2026
- Introduction to Subagents: Anthropic, Issued Mar 2026
- Building LLM Applications With Prompt Engineering: NVIDIA, Issued Jun 2025
- Gemini for Application Developers: Google, Issued Jun 2025
- Prompt Design in Vertex AI: Google, Issued Jun 2025
- Vector Search and Embeddings: Google, Issued Jun 2025
- Inspect Rich Documents with Gemini Multimodality and Multimodal RAG: Google, Issued Jun 2025
- Machine Learning Operations (MLOps) for Generative AI: Google, Issued Jun 2025
- Responsible AI for Developers: Privacy and Safety: Google, Issued Jun 2025
- Responsible AI for Developers: Interpretability and Transparency: Google, Issued Jun 2025
- Responsible AI for Developers: Fairness and Bias: Google, Issued Jun 2025
- AWS Cloud Practitioner: Amazon Web Services, Issued Jun 2025