DAN GILLES

dangilles@gmail.com | +1 872-259-4604 | github.com/Audacity88 | linkedin.com/in/dfgilles
U.S. citizen | English (native), Mandarin Chinese (professionally fluent)

SUMMARY

AI-native full-stack engineer who builds directly with the people who use the software. At Superbuilders I took products from an ambiguous brief to a working pilot and worked hands-on with the teams who ran my workflows. I also build agent infrastructure: Eyrie, Commander, and code ownership in ZeroClaw, a 33K-star open-source agent runtime. Coding agents write much of my code; I specify the problem, review every change, test it live, and decide what ships.

EXPERIENCE

AI Engineer (Contract), Superbuilders | *Austin, TX (remote)* Apr 2025 – Oct 2026

- **Zero to one:** took Civics Debate Academy, a multiplayer AI-judged debate platform, from concept to a pilot with 39 students across multiple schools: wrote the PRD, built and demoed the MVP to the Bill of Rights Institute, and led the build-out.
- **Full stack, shipped fast:** built the real-time multiplayer stack in React, TypeScript, PostgreSQL and Socket.io (lobbies, matchmaking, live state sync, disconnect and timeout recovery), delivered in a six-week sprint; AI judge (GPT-4 rubric scoring) with response caching to control cost and calibration checks to find where its scores were unreliable.
- **End users:** worked directly with the curriculum and product teams who used my agentic workflows; turned their needs into requirements, explained tradeoffs (cost, reliability, readiness for students), and decided with them what was ready to pilot.
- **Workflows:** built a multi-stage AI video pipeline (Civics Time Machine) with quality gates that regenerated defective output instead of shipping it; evidence-comparison interaction for Witness, a Grade 7 U.S. History app (React, TypeScript, Vite, PostgreSQL).

Code Owner and Core Contributor, ZeroClaw | *Open-source Rust agent runtime, ~33K GitHub stars* Mar 2026 – Present

- **Scale:** 752 merged pull requests since March 2026; listed in CODEOWNERS for the runtime, tools, gateway, config, memory and cost code, and a repository-wide default reviewer. I specify, review and live-test changes that agents help write.
- **Debugging:** while live-testing mid-turn steering, found that messages waited for the agent's whole turn instead of its next step; the cause was confirmed in source and fixed (#11725, merged). Closed a race that silently dropped late steering messages (#11617).
- **Cost and performance:** traced lost prompt-cache hits to three causes and shipped one fix per cause; a controlled A/B on the same 30 prompts raised hit rate from 81.8% to 93.1% and cut effective input cost ~51%, and later caught a regression on main.
- **Community:** member of GitHub's Maintainer Community (2026).

Independent AI Consultant | *AI agent assistant setup for small-business clients; first paid engagement Nov 2026* 2026 – Present

PROJECTS

Eyrie: cross-runtime agent coordinator | *Go, React, TypeScript, SQLite, SSE* 2026

- Coordinator that lets agents on different runtimes (ZeroClaw, Codex App Server) and machines work in one project without a human relaying messages; v0 prototype validated across a laptop and a Mac Mini over SSH.
- Cut scope from a broad durable-task platform to an agent-first v0: one bounded request in, collaborators consulted, a checked result out plus only the decisions that need a human. Worked through identity, cancellation, busy-session handling and per-runtime approval deadlines; fixed a concurrent-turn bug found in review with a server-side busy check.

Commander: Git-backed multi-agent system | *Node.js, Ruby, YAML, Git* 2026

- Supervisor/captain/specialist hierarchy of named agents passing work through file-based inboxes, with validator scripts, generated status pages, and pushes gated on human approval of the exact snapshot. Lesson that led to Eyrie: inspectable file-backed state worked, but I was still the relay between agents.

EARLIER

Software Engineering Fellow (Cohort 1), Gauntlet AI | *Austin, TX* Dec 2024 – Mar 2025

- 12-week, 1,000-hour AI-first engineering program: scoped, built and demoed a new application from scratch every week.

Graduate Student Researcher, University of Chicago Knowledge Lab Nov 2023 – May 2024

- Produced technical visualizations of neural networks and training dynamics for a forthcoming deep learning textbook.

English and STEM Educator, Beijing, China 2012 – 2021

- Taught English and STEM from university courses to AP sciences, trained new teachers, and co-founded a children's English program. Also worked as a Chinese-English translator and editor.

EDUCATION

M.A., Computational Social Science (NLP/ML), University of Chicago 2024
M.A., Eastern Classics, St. John's College, Santa Fe 2022
B.A., Liberal Arts (Great Books), St. John's College 2013

SKILLS

Stack: TypeScript, React, Vite, Node.js, PostgreSQL, SQLite, Socket.io, SSE, Go, Rust (Tokio), Python, SQL, Git
AI systems: LLM provider integration, tool calling, streaming fallback, prompt caching and cost analysis, evals
AI-assisted development: Claude Code, Codex, Cursor; I own requirements, review, testing and correctness of what ships