

DAWN GABRIEL

Brooklyn, NY | dgabriel@gmail.com | 617-871-9686 | linkedin.com/in/dawn-gabriel-3948946 | dawngabriel.com

Engineering Leader | Enterprise AI Deployment & Agentic Systems

Summary

Engineering leader and hands-on builder with 20+ years shipping enterprise software and 2+ years managing engineering teams of up to 20, spanning consulting, higher education, and independent product development. Owns customer-facing delivery end to end, from technical strategy and security standards through implementation and ongoing operations, in security- and compliance-constrained environments across public sector, finance, and healthcare. Builds and evaluates agentic systems firsthand: directed production LLM deployments with automated adversarial-evaluation release gates; created Flaught, an open-source adversarial AI code-review agent; delivers client platforms at 4x prior speed using AI-generated PRDs, test plans, and requirements traceability. High agency; known for translating between engineers and executives and for bringing agentic solutions into every step of the development process.

Skills

AI & Agentic Systems: LLM application development, adversarial evaluation, release-gating pipelines, agent development (Claude, OpenAI, Microsoft Copilot), prompt engineering, model-agnostic architecture, spec-driven development, AI-assisted delivery (Claude Code, Codex, Copilot 365)

Languages & Frameworks: Python, TypeScript/JavaScript, Java, Scala; Spring Boot, Flask, React, Angular, htmlx

Cloud & Infrastructure: AWS, Azure, Google Cloud; Terraform/IaC, Docker, CI/CD (GitHub Actions), Cloudflare Workers, Databricks (MLOps), Microsoft Power Platform

Data: PostgreSQL, MongoDB, DynamoDB, MySQL

Leadership & Delivery: customer engagement ownership, cross-functional team leadership (up to 20), technical strategy and security/compliance standards, stakeholder and executive communication, pre-sales scoping, hiring and mentoring; public-sector compliance environments (FedRAMP, GSA contracting)

Experience

EY (Ernst & Young) | Manager, Digital Engineering

Apr 2024 – Present | New York, NY

Lead AI adoption and engineering delivery for city and state government clients, from technical strategy and security standards through implementation and ongoing operations, across cross-functional teams of up to 20 (engineering, business analysis, QA).

Production AI Systems & Agentic Delivery

- Defined the release-gating standard for production LLM features: no model or prompt change ships without passing an automated adversarial evaluation pipeline (Azure, Python/Flask, OpenAI); directed migration of the supporting infrastructure to Terraform IaC.
- Leading AI-assisted, spec-driven delivery of a statewide disaster-recovery grants platform (Azure, Microsoft Power Platform) with a 20-person cross-functional team, using LLMs to generate PRDs, test plans, tickets, and requirements traceability, achieving 4x delivery speed compared with similar engagements.
- Built an agent that generates daily executive delivery dashboards, replacing 30–60 minutes of manual reporting per update.
- Scoped and directed a Microsoft Copilot agent for the NYC Department of Education Risk team to normalize reporting over standardized document sets, with projected savings of 60+ hours per quarter over manual review.
- Led 8 engineers building AI Workbench, an internal platform that prioritizes and scopes AI consulting engagements (Spring Boot, PostgreSQL, Azure, Angular).

Platform Delivery & Client Ownership

- Led delivery of a statewide disaster-recovery grants platform across forensics, case-management, and customer-service workstreams; the system has processed 1,300+ applications and \$88.45M in approved funding and is now in operations and maintenance.
- Directed an 18-person engineering team for NYC's Office of Technology & Innovation across the Small Business Services portal (AWS, Spring Boot, React) and an Adobe AEM CMS migration; set technical strategy and security standards.
- Tech lead for a 6-engineer Databricks MLOps transformation kickoff for Vizient Healthcare (Python, Azure).
- Translate EY's spec-driven AI delivery methodology into go-to-market materials for public-sector pursuits; grow delivery teams through candidate interviews, campus recruiting outreach, and mid-career hiring.

Independent AI Product Development | Founder & Hands-On Builder

2023 – Present | Brooklyn, NY

Own products end to end (problem definition and personas through PRDs, technical design, implementation, and launch) as a

standing agentic-development practice.

- **Flaught** (open source, flaught.github.io): adversarial CI code-review agent built to catch what LLM-assisted development loops miss: vacuous tests, out-of-scope changes, findings that vary run to run. Model-agnostic swappable-LLM architecture; deterministic tools (Semgrep, linters, vulnerability scanners) as non-negotiable pipeline gates; persisted, explicit rationale required to dismiss flagged findings. Adopted into an internal EY AI development toolkit after advocating for its inclusion. (TypeScript)
- **Cosmic JEV** (dawngabriel.com/cosmic-jev): deadpan haruspicy oracle with serious plumbing; computes natal and transit charts from real ephemeris data (astronomy-engine) and routes questions through TypeSafe's Jev System One model (typed probabilistic answers, no generated text) via a spend-capped Cloudflare Worker proxy, plus a Manifest V3 extension scoring event listings across four sites (TypeScript, JavaScript, Cloudflare Workers).
- **MysteryMixClub**: streaming-service-agnostic music-discovery PWA taken from concept to beta across web and iOS; iOS port completed in days using Claude Code and Codex (Python, React, PostgreSQL).
- **tktk**: multi-tenant micro-LMS for writing workshops, built directly with poet customers on near-zero budgets (Cloudflare Workers, Hono, D1/Drizzle, htmx).

Nuvalence (acquired by EY) | Senior Software Engineer / Technical Lead

Jan 2023 – Apr 2024 | New York, NY

- Technical lead for NY State DMV online license and permit renewal (Java, Spring Boot, Google Cloud, Angular); led 5 engineers through post-acquisition knowledge transfer.
- Technical lead for Ford Credit customer-identity work (Java, Spring Boot, Google Cloud); gathered stakeholder requirements, set engineering priorities, and secured team adoption of a two-week release cycle.

Harvard Business School | Principal Software Engineer / Technical Team Lead

Jan 2008 – Jan 2023 | Cambridge, MA

- Partnered directly with faculty to define product scope and prioritize iterations, translating research concepts into interactive learning experiences used across MBA and Executive Education programs; several products commercially licensed via Harvard Business Publishing.
- Primary or sole developer of production simulations: transformed R/Python research models into production applications (Scala/Play, TypeScript, Java/Spring, Angular, SQL databases, MongoDB).
- Led cross-functional front-end, DevOps, and data-science delivery against fixed, immovable academic deadlines; managed deployments to on-premises Linux servers before pivoting exclusively to AWS; provided live-event support for running classroom simulations.
- Created the development department's code-review and "Dev Complete" standards; supervised and reviewed code from globally distributed development teams; mentored and managed Year Up interns and junior developers; served 24-hour on-call rotations.

Early Career

- 8th employee at Spinners, Inc., the first startup of future Wayfair founders Niraj Shah and Steve Conine.
- Software engineer, Stumpworld Systems: e-commerce platforms for clients including Phish, Kiss, and Aerosmith.
- Senior engineering roles at American Well (telehealth infrastructure), Aspen Technology (enterprise inventory systems), and Virtual Text (early digital publishing tools).

Presentations

- "The Stars Recuse Themselves: Real Ephemeris, Typed Answers, and a Chrome Extension That RSVPs for Me" (live technical demo), *Before LinkedIn finds out: Jev Demo Day*, AI Tinkerers NYC, Oct 2026 (upcoming).
- "Leveraging Generative AI to Negotiate Software Development Priorities" (lightning talk), Grace Hopper Celebration, Sep 2023.
- Spec-driven AI development methodology: recurring internal presentations and evangelism to EY partners, principals, and managing directors (2024–present).