

Graham Anderson

Buffalo, NY (EST)

graham@grahama.co · grahama.co · linkedin.com/in/grahamanderson · github.com/grahama1970

I build agent systems that leave evidence of what they did.

Principal AI Engineer · AI Architect · Machine Learning Engineer

Agentic AI · LLM/RAG · Knowledge Graphs · Defense & Aerospace · Founder, grahamaco

U.S. citizen. Extensive experience delivering under export-controlled (ITAR) constraints.

This is the two-page version. The longer one — full project detail and the live capability inventory — is at grahama.co/resume.

ABOUT

Principal AI Engineer and AI Architect building verifiable agent systems, knowledge-graph platforms, and AI infrastructure for regulated technical environments. Most of my recent work sits in export-controlled aerospace and software-assurance settings, so I keep public claims focused on architecture, code, evaluation, interfaces, and technical briefing.

My path has been unusual: interactive production, data science, and now defense R&D. Earlier on, that included commercial and interactive work around familiar names like Sony, Adidas, X-Games, Disney, Microsoft, and Toyota; more recently it has meant aerospace, software assurance, and agent systems. I work best when the problem crosses disciplines and still needs one person to take responsibility for the complete result. Open to scoped consulting and selected Principal, Staff, Architect, LLM Platform, and Security/Compliance AI roles.

SELECTED IMPACT

- DARPA ARCOS / ACERT — technical lead for knowledge-graph and LLM-assisted certification of mission-critical software, briefing prime-contractor teams and program collaborators while remaining hands-on with architecture and code.
- Tau and agent-skills — built receipt-gated agent orchestration plus 340+ reusable agent capabilities and 90+ bounded worker roles, with deterministic checks on roughly 85% of the public capability set.
- Regulated graph-memory and compliance systems — built private RAG/GraphRAG systems over NIST, CWE, MITRE ATT&CK, D3FEND, SPARTA, and requirements-to-proof corpora; public details stay bounded because much client work is export-controlled.

EXPERIENCE

Founder & Principal AI Engineer / Architect | grahamaco | Buffalo–Niagara Falls Area · Remote

Feb 2025 - Present Independent AI engineering practice taking short, scoped engagements for aerospace primes, federally funded laboratories, and defense contractors. Agentic-pipeline work is active daily; client work is export-controlled (ITAR) and names are withheld; publicly releasable engineering is at github.com/grahama1970.

- Built and maintained tau, agent-skills, scillm, extraction, evidence, and evaluation systems for agent work that must produce receipts, checks, and explicit limits instead of only plausible outputs.
- Develop a heavily diverged fork of pdf_oxide (origin: yfedoseev/pdf_oxide, MIT/Apache-2.0; independent since Mar 2026): 430 commits, ~137K lines added across Rust-core changes, Python pipeline/plugin work, layout/table extraction, PDF-cloning fixtures, extraction calibration, and NIST document-validation tooling.
- Built private ArangoDB memory/compliance systems with hybrid BM25, vector, and graph recall over large evidence corpora; public claims are bounded to non-ITAR architecture, scale, and method.
- Delivered scoped client engagements end to end under ITAR, including a React/TypeScript/D3 dataset explorer over a security-control knowledge graph — graph relationships, integrity and coverage checks, and quality gates feeding downstream ingest and evaluation.

Lead Research Scientist, Agentic Formal Methods | grahamaco (independent practice) | Buffalo, NY

Jan 2024 - Feb 2025

- Designed verifiable agentic systems: a process-driven autoformalization workflow translating engineering requirements into Lean 4 proofs for safety-critical workflows.
- Built probabilistic-deterministic self-correcting loops where LLM generation is validated by compiler/prover feedback.
- Consulted in aerospace/defense settings under export-controlled / ITAR constraints.
- Conference speaking: presented agentic cybersecurity research at venues across the country, including to the Air Force Research Laboratory; received an AFRL "Hacker" challenge coin. Authored the talks, decks, and technical writing myself.

Principal Data Scientist & ACERT Technical Lead | CS Group (DARPA ARCOS)

Sep 2020 - Dec 2023

- Joined CS Group specifically for the 4-year DARPA ARCOS program (Automated Rapid Certification of Software), on the prime-contractor team, working alongside Honeywell, Lockheed Martin, MIT, GE Research, SRI, and other program collaborators.
- Led design and delivery of ACERT (Automated Certification of Requirements Tool): architected the ArangoDB knowledge-graph schema and LLM-assisted reasoning pipeline for multi-hop compliance verification of mission-critical software against complex certification standards.
- Extraction pipeline and datalake ran on AWS GovCloud, processing Boeing engineering documents for program partners.
- Briefed the program and its collaborators at reviews and conferences nationally; split time roughly 50/50 between executive/technical briefings and hands-on architecture and code, authoring the decks and demos myself.

Data Scientist | grahamaco (independent practice) | NYC · Remote

Sep 2011 - Sep 2020

- Freelance data science practice across gaming, manufacturing, entertainment, education, health, and e-commerce; ITAR-rated work included.
- Clients included Toyota, Sony, Fox, Boehringer Ingelheim, UCLA Med, Dartmouth (edu-tech prediction), and Domain Industries (manufacturing intelligence).

Earlier: Interactive Executive Producer & Composer | Los Angeles, CA

2005 - 2011

- Director of Interactive Services at Dentsu America (LA division; 50% internal labor reduction, budgets \$10K–\$1M, teams of 3–50); Executive Producer, God of War: Ascension campaign for Sony (Webby-recognized, 80+ person productions); commercial composer for Adidas, Pepsi, and X-Games.

PUBLIC WORK (non-ITAR) — github.com/grahama1970

Client work is mostly export-controlled, so here is the public, verifiable side:

- [agent-skills](#) — public working record: 340+ reusable agent skills, 90+ worker roles, ~85% with sanity gates. Public repo, private runtime.
- [tau](#) — receipt-gated multi-agent harness. "Agents hallucinate. Tau contains them."
- [pdf_oxide](#) — heavily diverged fork of yfedoseev's Rust PDF toolkit (430 commits, ~137K lines added: Rust-core changes, Python pipeline, PDF cloning, NIST validation).
- [scillm](#) — LLM gateway/proxy and LLMops layer: provider routing and fallback across hosted and local models, batch inference pools, structured outputs with repair, and streaming transport for agent runtimes. Runs on Docker/Linux.
- [grahama.co](#) — this site, built and designed by me: Next.js static export, self-hosted variable type, a live d3-force capability graph, and generated surfaces that fail the build if any count drifts from the repo.
- [extractor](#), [anvil](#), [fetcher](#), [chatterbox](#) voice-agent fork — supporting cast, all public.

EDUCATION

- Metis Data Science Fellowship | 2016
- Trinity University | BS, Finance, Marketing, and Economics

CORE COMPETENCIES

- Evals & Quality: LLM Evaluation, Agentic Evaluation Harnesses, Adversarial/Blind Testing, Regression Gates, Ground-Truth Fixtures
- Observability & LLMops: AI Observability, Drift Detection, LLMops
- Agentic Orchestration: Multi-Agent Systems, AI Agents, DAG Contracts, State Tracking, Tool Calling, Model Context Protocol (MCP), Bounded Agent Roles, Prompt Engineering
- LLM Platform: Large Language Models (LLM), Generative AI, LLM Gateway/Router, Provider Fallback, Structured Outputs, Batch Inference, Guardrails
- Retrieval & Knowledge: Retrieval-Augmented Generation (RAG), GraphRAG, Knowledge Graphs, ArangoDB, Vector Databases, Hybrid BM25 + Vector Search
- Verification & Compliance: Formal Verification, Lean 4, NIST 800-53/171, MITRE ATT&CK, AI Governance
- Document AI & Interfaces: PDF Extraction, Layout & Table Extraction, Extraction Calibration, React, TypeScript, D3, Design Systems
- Briefing & Communication: Conference Speaking, Executive & Technical Briefings, Technical Writing
- ML & Platform: Machine Learning, Model Fine-Tuning, Classifier Training, Python, Rust, Docker, Linux