

Senior Software Engineer, Applied AI

Reports to CTO · Rosslyn, VA (Hybrid with BU Travel)

About Graham Holdings Company

Graham Holdings Company (NYSE: GHC) is a diversified holding company whose operations include education, media, healthcare, automotive, manufacturing, and restaurants. Our corporate team is small and purposefully generalist — we provide oversight, strategic advisory services and additional capacity to a portfolio of businesses that vary widely in size, industry, and operational maturity. Everyone on the corporate team does real work alongside their advisory role.

The corporate Applied AI Team, led by our Chief Technology Officer, was recently established to help our business units navigate AI enablement with credibility and practical impact. This is not a center of excellence in the traditional sense — it is a small, high-leverage team that works directly in and with our portfolio companies.

The Role

This is a rare opportunity to see firsthand how AI creates real impact across vastly different industries and at companies ranging from regional to international operations.

This is a builder's role. We are looking for a senior engineer who can architect and ship production AI systems — someone equally comfortable writing Python to integrate a vendor API, designing an agent framework that non-engineers can extend on their own, deploying a workflow to the cloud, and sitting with a finance or operations team to understand where AI can and can't make impact.

You will partner closely with product and project management counterparts on the Applied AI team. Together you will deliver Gen AI solutions across multiple business unit engagements at different stages — some early in their AI journey, others further along — while also building the shared infrastructure and composable components that accelerate the entire portfolio.

You will operate as a high-leverage individual contributor with the autonomy to select tools and platforms. The role requires strong engineering judgment, an instinct for where AI actually creates leverage versus where it doesn't, and the communication skills to work alongside non-technical stakeholders as easily as you architect a system.

Success requires exceptional emotional intelligence and the ability to build trust quickly across very different organizational cultures, leadership styles, and technical environments. Technical fluency matters, but people skills are equally important.

What You'll Do

AI Project Execution at Business Units

- Architect, build, and deploy end-to-end Gen AI solutions across multiple portfolio companies — from prototype through production — in partnership with Applied AI team product management and business unit (BU) leadership and technical teams.

- Translate decomposed business workflows into technical designs that fit each BU's actual operating environment, data, and constraints.
- Act as a hands-on technical resource for BU teams during project engagements: pairing on architecture, unblocking integrations, and lending engineering depth where they need it.

Platform & Composable Infrastructure

- Build shared infrastructure that makes AI adoption fast and safe across the portfolio: reusable agent frameworks, skill libraries, evaluation pipelines, integration templates, and self-serve building blocks that non-technical teammates can compose into their own workflows.
- Establish engineering best practices for prompt design, tool/function calling, agent architectures, retrieval, evaluation, and quality assurance.

Cross-Functional Execution with the Applied AI Team

- Pair with product and project management counterparts on the Applied AI team, working alongside BU counterparts for delivery and hand-off.
- Act as the translation layer between business needs and technical execution, working across BU leadership, BU engineering, and the corporate Applied AI team.
- Provide candid, grounded technical assessments of AI vendors, model and tooling tradeoffs, and build-vs-buy decisions.

Production Operations & Quality

- Guide the reliability, observability, and security posture of the AI systems you ship into production at BUs and at corporate.
- Build evaluation harnesses and monitoring so that quality testing, prompt drift, and failures surface quickly.
- Comfortable picking up the data and integration work that shipping AI requires — SQL, REST/GraphQL APIs, and simple data pipelines.

Cross-Portfolio Collaboration

- Produce clear technical documentation, usage guides, and training materials so non-technical teammates and BU engineers can confidently extend the tools you build.
- Contribute to corporate AI initiatives, internal tooling, and the broader AI policy / governance work as capacity and priorities allow.
- Represent Graham Holdings credibly with AI vendors, partners, and the wider practitioner community.

What We're Looking For

Required

- Senior software engineering background, with demonstrated experience building, shipping, and operating production systems.
- Hands-on experience building applications using modern LLM platforms and APIs, spanning agent architectures, context engineering, and evaluation of model outputs.

- Strong product instinct and stakeholder skills: you can run a requirements conversation with an accounting or HR team as comfortably as you can architect a system, and you build for the end-user experience, not just whether the code runs.
- Comfort with ambiguity and incomplete information; you make good decisions without waiting for perfect data.
- 6+ years of professional engineering experience, though we care more about depth and judgment than the exact number.

Preferred

- Experience as the first or early engineer on an internal platform or core startup team.
- Exposure to multi-business-unit environments (holding company, conglomerate, consulting, or PE-backed portfolio) and the organizational dynamics of mid-market businesses with varying technical maturity.
- Active use of AI coding tools (Claude Code, Codex, Cursor, or similar) and a developed point of view on where they help and where they don't.
- Fluency in a mainstream programming language like Python, TypeScript or Java.
- Comfortable working across SQL, JSON, messy CSVs, and unstructured documents (PDFs, contracts, emails); can clean, restructure, and integrate data across warehouses and operational systems.
- Strong cloud infrastructure skills: you can take a working prototype and deploy it to AWS, Azure or GCP with proper monitoring and security.
- Experience designing platforms, internal tools, or self-service systems intended for non-developer users.
- Experience in customer-facing technical roles where requirements were never fully specified up front.
- Bachelor's degree or higher in Computer Science, Engineering, or equivalent practical experience

Logistics & Expectations

- Location: This role is based in our Rosslyn, VA corporate headquarters. We expect regular in-office presence several days a week. Remote work is possible some of the week but not the norm.
- Travel: Expect periodic travel to business unit locations during active BU engagements, typically lighter than the product management roles but meaningful — likely a few days per month on average, varying by project phase. Our portfolio companies are primarily U.S.-based.
- Compensation: \$150,000 – \$250,000 base and a comprehensive benefits package. Leveling will reflect depth of experience.

Growth & Scope

We are candid about the fact that this role and the team around it are still taking shape. As a first engineering hire for the team, the immediate focus is on shipping Gen AI solutions at BUs and building the shared AI infrastructure that supports them. Over time, the role has the potential to expand into broader corporate AI architecture, engineering leadership, vendor and

partner relationships, and team-building — depending on the skills, interests, and track record of the person we hire. If you want a role with a fixed job description that won't change, this is not it.

People and Culture

Graham Holdings is committed to being an exceptional place to work — one built on inclusion, fairness, and mutual respect. We hold ourselves to the highest ethical and professional standards and are dedicated to hiring and developing talented people from all backgrounds.

Benefits

We offer a comprehensive benefits package designed to support you at every stage of your career and life. Graham Holdings offers standout retirement benefits — a 401(k) plan plus a pension with company pension contributions starting at 8% and growing to 10% with tenure. Beyond that: 4 weeks of vacation to start, paid transportation, family leave, short- and long-term disability, and \$5,000 in adoption assistance.

We also offer \$5,250/year in tuition assistance plus tuition discounts, a choice of 3 medical plans, 2 dental plans, vision coverage, life and accident insurance, FSA/HSA accounts, a legal plan, and employee discounts.

For consideration, please send resume and salary requirements to careers@ghco.com, reference Senior Software Engineer Applied AI in the subject line.

Graham Holdings Company is an Equal Opportunity Employer committed to diversity and inclusion in the workplace. All qualified applicants will receive consideration for employment without regard to race, color, religion, age, sex, sexual orientation, gender identity/assignment, disability, genetic information, marital status, national origin, pregnancy, or maternity status, protected veteran status, or any other legally protected basis, in accordance with federal, state, and local applicable law. We will ensure that individuals with disabilities are provided reasonable accommodation to participate in the process. Please contact us to request accommodation.

The wage range for this role considers various factors used in making compensation decisions, including but not limited to skill sets, experience, training, licensure, certifications, and other business and organizational needs. It is not typical for an individual to be hired at or near the top of the range for their role, and compensation decisions are dependent on the specifics of each individual situation. A good-faith reasonable estimate of the current range is \$150,000-\$250,000.