GRAHAM HOLDINGS

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

We are looking for a software engineer who wants to build and ship production AI systems — someone comfortable writing Python to integrate a vendor API, extending the agent frameworks and building blocks that non-engineers use, deploying a workflow to the cloud, and sitting with a finance or operations team to understand how they actually work and where AI can help.

You will work closely with senior engineers and product 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 contributing to the shared infrastructure and composable components that accelerate the entire portfolio.

This is a hands-on individual contributor role with real ownership and room to grow. You will own well-scoped projects and components end to end, with mentorship and design guidance from senior engineers. The role calls for strong engineering fundamentals, a bias toward shipping, curiosity about where AI actually creates leverage versus where it doesn't, and the communication skills to work directly with non-technical stakeholders.

What You'll Do

AI Project Execution at Business Units

- Build and deploy end-to-end Gen AI solutions across multiple portfolio companies — from prototype through production — working with senior engineers, Applied AI team product management, and business unit (BU) leadership and technical teams.
- Learn how BU teams actually work — not just how they describe it — and help design solutions that fit their real stack, data, and constraints.

- Serve as a hands-on technical resource for BU teams during project engagements: building integrations, debugging issues, and pairing with BU engineers to get solutions into production.
- Help migrate successful local prototypes — laptop scripts, ad-hoc workflows, manual spreadsheet processes — onto reliable, scheduled, or always-on hosted infrastructure with proper monitoring and observability.

Platform & Composable Infrastructure

- Contribute to the shared infrastructure that makes AI adoption fast and safe across the portfolio: reusable agent components, skill libraries, evaluation pipelines, integration templates, and self-serve building blocks that non-technical teammates can compose into their own workflows.
- Apply and help refine the team's engineering practices for prompt design, tool/function calling, agent architectures, retrieval, evaluation, and quality assurance.

Cross-Functional Execution with the Applied AI Team

- Work in close partnership with senior engineers and with product and project management on the Applied AI team, and with BU leadership and engineering through delivery and hand-off.
- Translate business needs into technical requirements and working software, and communicate progress, tradeoffs, and limitations clearly to non-technical stakeholders.
- Evaluate AI vendors, models, and tools hands-on, and share clear, candid findings that inform build-vs-buy decisions.

Production Operations & Quality

- Own the reliability, observability, and security of the components you ship into production at BUs and at corporate, following the team's standards.
- Build and maintain evaluation harnesses and monitoring so that quality issues, prompt drift, and failures surface quickly.
- Pick up the data and integration work that shipping AI requires — SQL, REST/GraphQL APIs, and simple data pipelines.

What We're Looking For

Required

- 2–5 years of professional software engineering experience, though we care more about ability, curiosity, and trajectory than the exact number.
- Demonstrated experience building and shipping software to production — at work, or through substantial projects you have taken from idea to real users.
- Fluency in a mainstream programming language like Python, TypeScript or Java.
- Hands-on experience building applications with modern LLM platforms and APIs — for example, tool/function calling, agents, retrieval, context engineering, or evaluating model outputs.
- Strong communication skills and product curiosity: you are comfortable sitting down with an accounting or HR team to understand their work, and turning what you learn into something you can build.

Preferred

- Experience on a small, fast-moving team — an internal platform team, an early-stage startup, or similar — where you wore multiple hats.
- Interest in or exposure to multi-business-unit environments (holding company, conglomerate, consulting, or PE-backed portfolio) and mid-market businesses with varying technical maturity.
- Active use of AI coding tools (Claude Code, Codex, Cursor, or similar) and a point of view on where they help and where they don't.
- Comfortable working with SQL, JSON, messy CSVs, and unstructured documents (PDFs, contracts, emails), and cleaning and restructuring data so it can be put to use.
- Experience deploying applications to AWS, Azure, or GCP, with a working understanding of monitoring and security basics.
- Experience building internal tools or self-service systems for non-developer users.
- Experience in a client-facing or forward-deployed role — solutions engineering, technical consulting, or similar — where requirements weren't 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 — likely a few days per month on average, varying by project phase. Our portfolio companies are primarily U.S.-based.

Growth & Scope

We are candid about the fact that this role and the team around it are still taking shape. As an early engineering hire on the team, your immediate focus will be shipping Gen AI solutions at BUs and contributing to the shared AI infrastructure that supports them, working alongside and learning from senior engineers. Over time, the role offers a clear path toward senior engineering responsibilities — owning larger systems and architecture decisions, leading BU engagements, and mentoring others — depending on your skills, interests, and track record. 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 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 $125,000-$175,000.