Founding AI Engineer

Job description · October 30, 2026

Build the agents and intelligence behind Helmsted, the AI-native operating system for wealth advisors of complex households.

Location	Remote (US) today; eventual relocation as the team centralizes
Type	Full-time
Reports to	Co-founders

About Helmsted

Advisors use Helmsted for financial planning, tax and estate analysis, meetings and their client portal, and our goal is to replace a firm's entire tech stack. We are a small team, live with pilot firms and shipping every day.

The role

You will work directly with our CTO on the core of the product: the agents that do the work and the intelligence behind them. You will experiment constantly with what the frontier labs release and turn the best of it into a Helmsted that reasons, plans and executes better, faster and more efficiently. The goal is an intelligence layer that lets advisors focus on clients, not the fragmented, chaotic tech stack they have lived with for thirty years.

What you'll do

- Make it low-lift for advisors to keep consolidating their stack into Helmsted (integrations, browser control, agents that finish real tasks), our top priority as work moves away from typing and clicking
- Build features end to end, from agent logic and tools to what advisors and their clients see
- Improve how Helmsted reasons through a household's planning, tax and estate questions, and the reports it writes
- Own quality with evals, guardrails and source tracing, so every output is accurate and ready for advisor review
- Own AI security: least-privilege agents, permission-aware retrieval and memory, and defenses against prompt injection and data leakage that you red-team yourself
- Make the model layer efficient: route each task to the cheapest model that does it well, and tune caching, batching, memory, retrieval and context for cost and latency
- Run a constant experiment loop: put promising new models and techniques through our evals within days of release, and ship what makes Helmsted better, faster or cheaper
- Work directly with advisors at our pilot firms and turn what you learn into product

What you bring

- A record of shipping LLM-powered products that real people use: agents, tool use, retrieval, evals
- Deep understanding of how AI systems leak and get attacked: prompt injection through emails, documents and web pages; data leaking between clients through retrieval, memory or caching; tool misuse; and which client data should ever reach a model provider
- Strong Python and TypeScript, and comfort across the stack (APIs, Postgres, React)
- Deep hands-on time with frontier model APIs, and judgment about which model fits which job
- An experimenter's discipline and an efficiency mindset: fast, measured trials on new models, and a record of cutting cost and latency without losing quality
- High agency: you find the problem, decide with imperfect information and ship
- Pro-AI: you can build through AI tools almost exclusively, and still ship without them when you have to

Highly preferred

- Fintech or wealthtech experience (advisor software, custodians, banking, payments)
- Working knowledge of the rules advisors operate under (SEC and FINRA oversight, fiduciary duty, Reg S-P privacy, books and records)
- Browser automation or computer-use agents
- Fine-tuning or serving open-weight models
- AI red-teaming or security work (OWASP Top 10 for LLM and Agentic Applications), or building inside a SOC 2 program
- Fast decision models (models that return a choice with a calibrated probability instead of text) used to route, triage or check agent actions

Why join

- Founding equity and real ownership of the product, the stack and how we build
- Start remote anywhere in the US
- Freedom to experiment, with every frontier model and AI tool you need covered
- Work directly with the founders and our pilot firms, and see advisors use what you ship

Pay and details

Base salary	$160,000–$200,000
Equity	Early-team grant, sized by experience
Benefits	Paid sick leave; health insurance and a 401(k) are under discussion but not offered yet
AI tools	All covered
Applications	Accepted on an ongoing basis
Background check	Final offers are contingent on a background check

How to apply

Skip the cover letter. Email careers@helmsted.ai with short answers to these; bullets are fine. We want to see how deep your thinking on AI systems goes. Use AI to help you answer, and tell us how you did.

1. Design a system that measurably improves every week from its own usage, without retraining a base model. What does it learn from, how do you stop it learning the wrong lessons, and how do you prove it got better?
2. An agent works with the same people for years. How would you design its memory so it learns what matters and forgets what it should?
3. One shared “brain” serves a whole organization, but each person, team and client may only see and do what they are permitted to. How would you keep that brain divisible and permissioned across retrieval, memory and agent actions, and prove nothing ever crossed a line?
4. Your agents must reason over millions of documents, records and conversations that change daily. What goes into context, what gets retrieved, summarized or fine-tuned, and how would you know when retrieval is quietly failing?
5. Every task could run on the best frontier model, but cost, latency and rate limits say otherwise. How would you route work across frontier and open models, caching and batch jobs, and how fast could you prove whether a new release should change that?
6. A client sends their advisor a PDF with hidden instructions telling the agent to email account statements to an outside address. Walk us through every layer that stops it, and how you would prove those layers keep working as models change.

Add a resume, GitHub or LinkedIn, plus a link to an agent or AI system you’ve shipped if you can. We move fast.

Helmsted, Inc. is an equal opportunity employer. We evaluate candidates on merit and do not discriminate on the basis of race, religion, color, national origin, sex, sexual orientation, gender identity, age, veteran status, or disability.