SecureLLMs.org

AI Data Security for Asset Managers

What You Need to Know Before Your Next ChatGPT Session

The Question You Should Be Asking

Every time you paste a client portfolio into ChatGPT or ask Claude to analyze a trading strategy, that data leaves your control. It travels across the public internet to servers you do not own, where it may be stored, analyzed, and potentially used to train the next generation of AI models. If your competitors are using the same tools, your proprietary insights could be shaping the very AI they rely on.

How AI Actually Works (The 60-Second Version)

A Large Language Model (LLM) like ChatGPT or Claude is essentially a pattern recognition system built from billions of text examples. Think of it like this: if you read every book, article, and website ever written, you would start to predict what words come next in any sentence. That is what these models do, just at a superhuman scale.

Training is when the AI learns these patterns from massive datasets. *Inference* is when you ask it a question and it generates a response based on those learned patterns.

Here is what matters for security: when you use a public AI service, your prompts and the data you share become inference data. The provider can see it, store it, and in many cases, use it to improve their models. Your proprietary information becomes part of their training pipeline.

The Self-Hosting Solution

Self-hosting means running AI models on infrastructure you control, rather than sending your data to OpenAI, Anthropic, or other public providers. It is the difference between using Gmail (where Google processes your emails) and running your own email server (where only you have access).

With self-hosted AI, your prompts, documents, and analysis never leave your environment. The models run inside your firewall. No data transmission to third parties. No multi-tenant servers. No risk of your information being used to train public models.

Why Public AI Is a Compliance Problem

For asset management firms, the risks of public AI extend beyond general privacy concerns:

Regulatory Exposure: SEC and FINRA have clear expectations about client data protection. When you send portfolio details, trading strategies, or client information to public AI servers, you may be creating compliance gaps. The data travels over public internet, sits on multi-tenant servers alongside millions of other users, and may be retained indefinitely.

Competitive Intelligence Leakage: Consider this scenario: you ask ChatGPT to help optimize your factor model. That conversation is now on OpenAI's servers. If they use customer interactions to improve their models (which most providers reserve the right to do), your proprietary approach could influence responses given to your competitors.

Data Sovereignty: You have no visibility into where your data goes, how long it is stored, or who can access it. For fiduciaries managing other people's money, this lack of control is increasingly untenable.

Why Most Firms Do Not Self-Host

If self-hosting solves the security problem, why is not everyone doing it? Because the infrastructure challenge is significant:

Complexity: Deploying AI models securely requires DevOps expertise, GPU infrastructure, security configuration, and ongoing maintenance. Most asset managers do not have dedicated AI engineers on staff.

Rapid Deprecation: AI models update every few weeks. The infrastructure you build today may be obsolete in months. Keeping pace requires continuous investment.

Cost and Time: Between hardware, cloud compute, engineering time, and security audits, building your own AI infrastructure can cost six figures and take months to deploy properly.

Evaluating Your Options

There is no single right answer for every firm. Some approaches worth considering:

Internal usage policies that restrict what types of data employees can input into public AI tools. Low cost to implement, but difficult to enforce and limits the utility of AI for your most valuable workflows.

Self-hosted or private deployment models that keep all data within your own infrastructure. Higher security ceiling, but historically complex and expensive to implement.

The right path depends on your firm's size, regulatory posture, and how central AI is to your investment process.

If you are thinking through this decision and want to talk through the tradeoffs, **including approaches that do not involve us**, we are happy to have that conversation.

SecureLLMs.org | michael@securellms.org