Executive Brief

AI Is Already Inside Your Fund. Does Anyone Own the Risk?

By Robert Choynowski, Chair, HFA Cybersecurity Committee | SeaGlass Technology

Prepared for readers seeking practical AI governance guidance for hedge funds and alternative investment firms.

AI is already entering fund workflows through the tools teams use every day.

This brief outlines the cybersecurity, compliance, and operational risks that come with unmanaged adoption.

It also offers practical guardrails leadership can put in place now.

> *The real question is no longer whether firms will use AI. It is whether they will govern it before it governs them.*

The conversation around artificial intelligence in the alternative investment space has changed quickly.

Not long ago, fund leaders were asking whether AI was relevant to their firms. Now the more important question is this: How much AI is already being used inside the organization, and who actually owns the risk?

At many firms, the honest answer is uncomfortable. AI is already showing up in daily workflows, but the governance around it has not kept pace.

This is not simply a technology issue. It is a cybersecurity, compliance, and operational risk issue. For hedge funds and other alternative investment firms, that matters.

AI adoption is happening quietly

Most funds did not approve some formal, enterprise-wide AI initiative. AI arrived through the tools employees were already using.

Microsoft Copilot is being introduced into M365 environments. Team members are using ChatGPT to summarize notes, draft emails, review documents, or accelerate research. Operations teams are exploring automation platforms with AI built in. Vendors are also adding AI features into products firms already rely on.

In many cases, none of this went through a formal risk review. It did not get vetted by compliance. It was not reviewed against data classification policies. And it may not have been evaluated for investor confidentiality, regulatory implications, or access control concerns.

That is what makes this so important. The issue is not just intentional AI adoption. It is unmanaged AI adoption.

Why this is a cybersecurity issue

AI governance often gets framed as an innovation conversation or a productivity conversation. It is both of those things, but for fund managers it is also a cybersecurity issue.

There are several reasons why.

Sensitive data may leave the firm.

When employees paste internal material into public or consumer-grade AI tools, they may be exposing proprietary research, investor information, strategy documents, valuations, portfolio company details, or internal communications. Even when the risk is not obvious to the user, the exposure can be significant.

AI can create false confidence.

Generative AI tools often produce answers that sound polished and credible, even when they are wrong. In a fund environment, that creates real risk. A flawed regulatory summary, an inaccurate memo, or an AI-assisted response built on incorrect assumptions can introduce compliance and operational exposure very quickly.

Permissions problems become AI problems.

Enterprise AI tools do not magically create access issues. They expose the ones that already exist. If file permissions are overly broad in SharePoint, Teams, or other repositories, AI-powered search and summarization can make that problem far more visible. Employees may suddenly be able to surface information they technically had access to but never would have found otherwise.

Vendor risk is evolving.

Many third-party platforms now include AI features, whether firms realize it or not. That means traditional vendor due diligence may no longer be enough. Firms should understand where data is processed, what is retained, whether customer content is used to train models, how security incidents are handled, and what contractual protections are in place.

The bigger issue: nobody clearly owns it

One of the most common problems is not that firms lack smart people. It is that AI risk often falls into a gap between departments.

IT may assume compliance is handling it. Compliance may assume IT is reviewing the tools. Legal may only get involved once a vendor contract appears. Business leaders may allow teams to experiment because the tools seem harmless and productivity gains are appealing.

Meanwhile, AI use keeps spreading.

That lack of ownership is the real concern. Because once a productivity tool becomes embedded in daily processes, it gets much harder to unwind.

Funds do not need to panic, but they do need guardrails

The goal is not to block every AI tool. That is unrealistic, and in many cases unnecessary. The goal is to put guardrails in place before adoption outruns oversight.

A good starting point is often much simpler than people think.

1. Create an AI use policy

Every firm should have a basic, plain-English policy that answers a few practical questions:

- What kinds of AI tools are permitted?
- What kinds of firm data may never be entered into AI systems?
- Which uses require approval?
- Who is responsible for oversight?
- What review process applies before a new AI-enabled vendor or feature is adopted?

This does not need to be a 30-page manual. It just needs to be clear enough that employees understand the rules.

2. Inventory where AI is already in use

Before building a future-state AI strategy, firms should understand current-state exposure.

- Public AI tools employees are already using
- Enterprise tools with AI features enabled
- Third-party vendors that process or analyze firm data using AI
- Internal automation workflows with AI components

You cannot govern what you have not identified.

3. Revisit data classification and permissions

AI governance is tightly connected to basic cyber hygiene.

If firms do not know where sensitive data lives, who has access to it, and whether permissions are properly scoped, AI can amplify those weaknesses. In many environments, the quickest win is not an AI tool decision at all. It is cleaning up access controls, data sprawl, and file-sharing practices.

4. Update vendor due diligence

Vendor reviews should now include AI-specific questions. For example:

- Does the vendor use customer data to train models?
- Where is the data stored and processed?
- Is data retained after prompts or tasks are completed?
- Can AI functionality be disabled if needed?
- What transparency does the firm have into model behavior and outputs?
- What does the vendor's breach notification and incident response process look like when AI systems are involved?

These are no longer niche questions.

5. Train employees before a problem occurs

Policies alone are not enough. Employees need examples they can relate to.

Most people are not trying to create risk. They are just trying to save time. Training should focus on real-world fund scenarios such as:

- summarizing internal meeting notes
- drafting investor communications
- using AI to review agreements or policies
- analyzing spreadsheets or portfolio information
- researching compliance questions

The more practical the training, the more likely it is to work.

What leadership should be asking now

For fund executives, COOs, CFOs, CTOs, CISOs, and compliance leaders, this is the moment to ask a few direct questions:

- Do we know which AI tools are currently in use across the firm?
- Has anyone defined what data can and cannot be used with those tools?
- Are our M365 and collaboration permissions actually in good shape?
- Have our vendors disclosed where AI is embedded in their offerings?
- Who owns AI governance here, in practice, not just in theory?

The firms that handle this well will have an advantage

There is a tendency to frame AI risk and AI opportunity as opposing forces. They are not.

The firms that put sensible guardrails in place early will be in a far better position to benefit from AI responsibly. They will be able to move faster with more confidence, reduce unnecessary exposure, and demonstrate to investors, regulators, and internal stakeholders that adoption is being handled thoughtfully.

AI is already inside the fund environment, whether leadership intended it or not.

The real question now is not whether firms will use AI. It is whether they will govern it before it governs them.

About the Author

Robert Choynowski is Chair of the HFA Cybersecurity Committee and a member of the HFA Global Board of Directors. He is also a founder and Chief Visionary of SeaGlass Technology, which supports hedge funds and alternative investment firms with managed IT, cybersecurity, and strategic technology leadership.

Contact

If you'd like to compare notes on AI governance, cybersecurity, or technology risk within your firm, Robert can be reached at robert@seaglasstechnology.com or through the HFA member network.