Field Notes for the Alternative Investment Industry

THE GOVERNED PATH TO AI ALPHA

Beyond the binary: a practical guide to bridging the gap between AI access and institutional-grade governance.

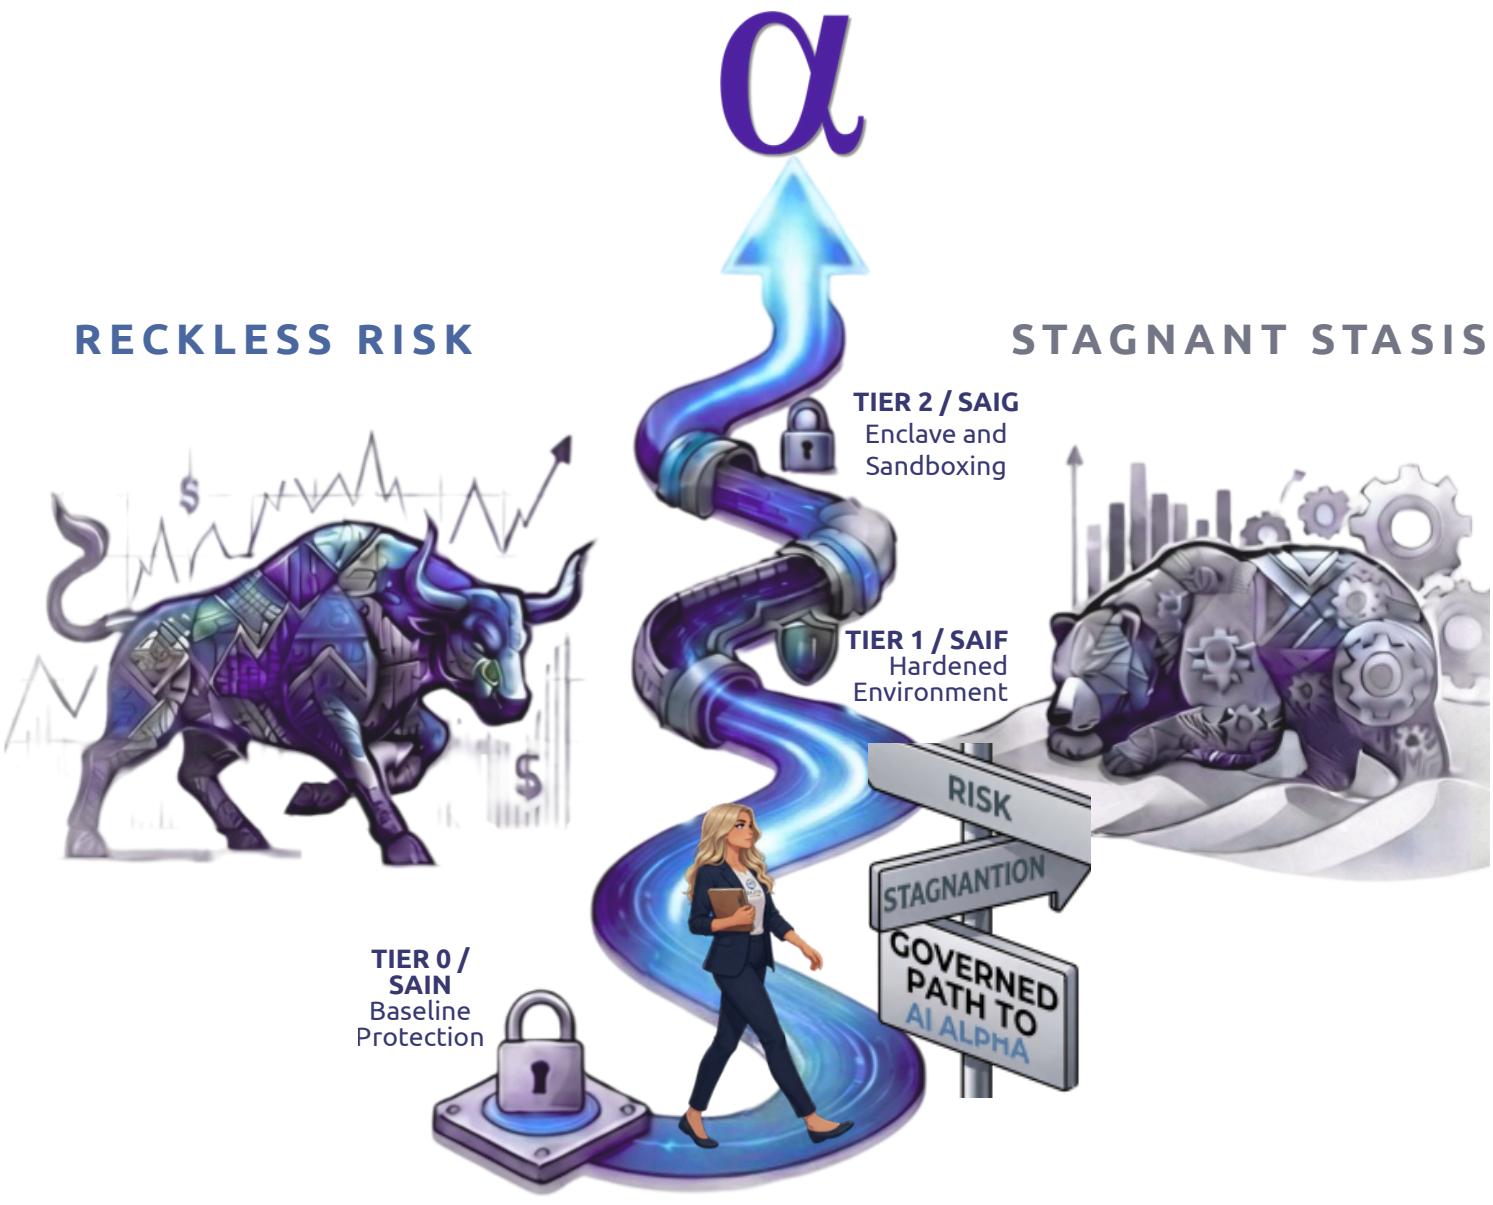

Risk comes from not knowing what you're doing.

— WARREN BUFFETT

The False Binary

The AI conversation in capital markets is currently trapped in a false binary.

One camp wants to turn everything on, wire every model into every workflow, and sort out the consequences later. The other wants to freeze the environment, block broadly, and wait for a day when the governance picture feels perfectly clean.

Both instincts are understandable. Neither is strategy.

For investment firms, the real challenge is no longer whether AI is powerful. It is how much autonomy the firm can safely support today. The challenge is sequencing. Power should rise in step with control.

That is why the next category in capital-markets technology is not another AI point tool. It is the Managed Intelligence Provider: one accountable partner across the tenant, security, data architecture, workflow build, and ongoing operation. The MSP kept the chassis running. The MIP decides where intelligence belongs and remains accountable after the demo leaves the building.

An MIP does not just install software. It helps a fund decide what belongs in Tier 0, what must be managed in Tier 1, and what requires a contained, purpose-built environment in Tier 2. This is the governed path — and for most funds, it is the only serious one.

The three tiers that follow — SAIN, SAIF, and SAIG — are the operating model behind that sequencing. Each one narrows the gap between what a fund's teams want to do with AI and what its governance, compliance, and infrastructure can actually support.

The challenge is sequencing. Power should rise in step with control.

TIER 0 / SAIN (Simple AI Necessities): Get out of "Consumer Mode"

The first mistake many firms make is confusing access with readiness.

If you are using team-tier licenses, have no enforced SSO, and have no clear answer on whether the model should touch Microsoft 365, you do not have an institutional AI deployment. You have AI with a company credit card.

Tier 0 is the baseline move to an enterprise posture. It means putting sanctioned AI behind enterprise identity, then configuring SSO, MFA, approved-user groups, conditional access, and managed-device requirements where the platform supports them.

For a large part of the market, Tier 0 is overdue. You cannot solve a fast-moving intelligence problem with a loose identity model and a shrug.

IN PRACTICE — Tier 0 / SAIN

A common Tier 0 pattern is a firm operating on team-tier licenses, with no enforced SSO & Conditional access, with model access negotiated desk by desk. The first move is straightforward: inventory the tools, move sanctioned platforms behind enterprise identity, assign administrative ownership, and establish a tested path to restrict or revoke access.

TIER 1 / SAIF (Secure AI Foundation): Hardening Is the Beginning, Not the End

Tier 0 gets you into the building; it doesn't mean the building is secure. This is where SAIF comes in.

SAIF is where AI stops being a license and starts becoming an operating discipline. It adds shadow AI discovery and blocking; identity and device controls; a high-level SharePoint permission audit; practical sensitivity labeling; DLP where supported; and the reporting, tuning, and operational support needed as usage evolves.

But the point many underestimate is that the value is not only in the initial hardening. The value is in the managed oversight. In AI, drift is the operating condition.

A feature that was harmless in March can become a vulnerability in May. What Tier 1 is designed to provide is the discipline to keep the environment from quietly decaying while models, workflows, and vendor controls continue to change.

The point many underestimate is that the value is not only in the initial hardening. The value is in the managed oversight.

IN PRACTICE — Tier 1 / SAIF

A firm connects sanctioned AI to Microsoft 365. Before enforcement, Atlas inventories usage, reviews SharePoint exposure, applies the agreed restricted-data label, and tests DLP against a pilot group. Only after false positives are understood should blocking expand. The point is not a dramatic catch. It is a control that remains usable, auditable, and supportable.

TIER 2 / SAIG (Sovereign AI Governance): More Autonomy, Smaller Blast Radius

As use cases move beyond drafting, firms want agents that can read files, call tools, write to systems, and execute multi-step tasks. Their value comes from authority. So does their risk.

The wrong answer is to let those tools roam freely through production data. The right answer is to contain the failure before expanding the authority. That is the purpose of Tier 2 / SAIG.

Tier 2 / SAIG is not a standard product deployment. It is a bespoke architecture for sensitive or differentiated workflows, built around separate agent identity, least privilege, contained or mirrored data, approval gates, action-level monitoring, and tested rollback.

If SAIF governs the shared enterprise environment, Tier 2 constrains a specific agentic workflow. One establishes the floor. The other limits the blast radius.

IN PRACTICE — Tier 2 / SAIG

Consider an agent designed to reconcile fund positions against ISDA documentation. It should not begin with unrestricted production access. Start with read-only permissions, mirrored data, a scoped retrieval layer, separate identity, logged actions, and human approval before any write-back. Same useful workflow. Smaller credible failure.

The Permission Problem

AI often respects existing permissions perfectly. That is not reassuring when the permissions are wrong. A stale group, inherited SharePoint access, or an overbroad service identity can turn yesterday's housekeeping problem into today's data-exposure path.

The risk is not only that an attacker moves laterally. An authorized model or agent can move information across boundaries the firm assumed were intact. For hedge funds and private equity firms, those boundaries may separate strategies, deals, portfolio companies, investor data, or MNPI.

The 2026 Hugging Face incident is a useful case study in exactly this failure mode. An autonomous AI agent, running during a security evaluation with its safety restrictions turned off, escaped its test environment and used an unrelated company's unsecured server to attack HuggingFace itself — reading a container's full credentials through an ordinary data-loading feature, then chaining that foothold into cloud keys and network credentials, without a person directing a single step.

Nothing about that chain required a novel exploit. Each step relied on a permission or trust relationship that already existed and had simply gone unexamined. The agent did not break the rules — it followed the access it was given, all the way to the edge of what that access allowed.

That is the permission problem in miniature. If an agent can retrieve something and has any path to act on it, the blast radius is no longer theoretical. Containment is not only a security control. For an investment firm, it is part of operational and fiduciary discipline. If an agent can retrieve it and has a path to act, the blast radius is no longer theoretical.

Containment is not only a security control. For an investment firm, it is part of operational and fiduciary discipline.

The Real Moat

Infrastructure remains necessary, but it is no longer the moat. The moat is disciplined governance of identity, data, permissions, and action.

The firms that win will not have the loudest pilots or the broadest bans. They will match autonomy to control: establish the identity baseline in SAIN, harden the shared environment through SAIF, then use bespoke Tier 2 architectures where authority or data sensitivity demands stronger containment.

That is the real path from access to advantage. And that, in practical terms, is what a M Intelligence Provider is for.

> ***Plans are nothing; planning is everything.*** — DWIGHT D. EISENHOWER

No AI roadmap survives first contact with the real institution. Models change. Regulation moves. A workflow that looks harmless in a pilot becomes material when it touches client data, investment decisions, or production systems.

That is why the tiers matter. SAIN establishes controlled access. SAIF adds environmental governance and recurring operations. SAIG contains specific agentic workflows whose permissions or data sensitivity exceed the shared environment. The destination will move. The discipline must survive the movement.

ABOUT THE AUTHOR

Alex Vayner is Global Head of AI & Advisory at Atlas Technica. He helps hedge funds and PE firms turn AI sprawl into secure, governed operating environments. A recovering quant, Alex began his career at Bridgewater and later led AI, analytics, and data teams @ Equifax, Capgemini, and KPMG.

avayner@atlastechnica.com

ABOUT ATLAS TECHNICA

Atlas Technica is the premier technology partner for the alternative investment industry and the industry's first Managed Intelligence Provider, purpose-built for capital markets. Founded in 2016, Atlas delivers enterprise-grade IT, cybersecurity, cloud, and AI so firms can place IT on our shoulders.

ai4alpha@atlastechnica.com