Alpha
Financial
Markets
Consulting

The Impact of AI on Asset Manager & Asset Owner IT Architecture

Is it time for asset managers and asset owners to evolve the way they design, build, and operate investment technology in the age of artificial intelligence?

The Impact of AI on Asset Manager & Asset Owner IT Architecture

Introduction

The technology architecture of asset managers and asset owners alike has followed a remarkably consistent playbook for the past decade. Consolidate on an enterprise investment management platform. Build a centralized Investment Book of Record. Outsource and/or augment technology capabilities. Keep internal teams small and focused on configuration rather than creation.

This playbook was rational. Enterprise platforms offered breadth of capability that most internal teams couldn't replicate. Centralized data stores promised a single source of truth. Vendor partnerships meant firms could focus on their core purpose of generating returns for clients and members. For many organizations, this approach has delivered exactly what it promised: operational resilience, regulatory confidence, and a stable platform for growth.

The AI inflection point of 2026 has forced each of these assumptions to evolve, opening new strategic options for firms to build on the foundations they have already established.

The opportunity is substantial. AI has the potential to compress the cost of building investment technology, dramatically accelerate research workflows, enable real-time portfolio simulation, and automate operational processes that currently consume significant human effort; however, several structural constraints must be navigated: fragmented data architectures, strict governance and explainability requirements, heavy existing investment in legacy platforms, and cultural resistance from investment teams accustomed to established ways of working. This paper examines the implications of these shifts for the IT architecture of asset managers and owners, aiming to offer a framework for understanding how AI changes the strategic options available while delivering a practical guide for navigating the transition.

The Scale of the Shift

In February 2025, OpenAI co-founder Andrej Karpathy coined the term "vibe coding", which is building software by describing what you want in plain English and letting AI write the code. Within months it went from novelty to norm: Y Combinator reported that a quarter of its Winter 2025 startups had codebases that were 95% AI-generated, and Google disclosed that over 30% of its new code was now written by AI.

By early 2026, Karpathy had moved on from the term entirely, arguing that the mature version of this practice, which he now calls "agentic engineering", is becoming the default professional workflow: "you are not writing the code directly 99% of the time, you are orchestrating agents who do."

This is the backdrop against which many asset owners/managers must now evaluate their technology strategy. A generation of tools is emerging that allows domain experts, not just developers, to build, test, and deploy analytical capability through natural language. For investment organizations, that unlocks possibilities that were previously gated behind months of vendor engagement or scarce internal development resourcing.

The Conventional Playbook

This urgency is underscored by industry data. Industry surveys indicate that 68% of asset management and private equity leaders are piloting AI agents, with close to a quarter already deploying them in their organizations. Global financial services spend on AI is expected to grow over 25% annually through the late 2020s. Meanwhile, industry operational research reinforces this point, indicating that data fragmentation is the leading challenge limiting front-office effectiveness, with firms' data teams spending the majority of their time on operational tasks rather than analysis, underscoring that data architecture is a foundational enabler of AI adoption.

Before examining what AI changes, it is worth articulating the three core assumptions that have guided asset manager and asset owner technology strategy.

Assumption 1: Consolidate on Enterprise Platforms

The prevailing wisdom has been that firms should select a comprehensive enterprise platform and consolidate as much functionality as possible onto it. The logic is compelling: reduced integration complexity, consistent data models, vendor-managed upgrades, and other operational efficiencies.

In practice, enterprise platforms have delivered genuine operational consolidation. However, where natural limits emerge is at the edges: highly differentiated investment workflows, asset class-specific analytics, and bespoke last-mile capabilities that reflect each firm's unique investment philosophy. These boundaries are not a failure of platform design but a feature of specialization, and they represent the precise space where AI-enabled tooling can extend platform value without displacing it.

Assumption 2: Build a Centralized IBOR

The Investment Book of Record has been one of the most consequential investments in modern investment operations, providing the control framework, auditability, and data consistency that regulator expect and operational teams depend on. It has been positioned as the foundation of modern asset manager/owner operations: 1) integrate internally managed assets' data in real-time with data sourced from custodians, third party managers, and market data providers; 2) reconcile and cleanse centrally; 3) load into a golden source; 4) distribute via standardized reporting views.

The challenge is that this supply-driven model produces data on the supplier's schedule, in the supplier's format, at the supplier's level of granularity. Investment teams needing data contextualized for their specific analytical question must manually blend, transform, and supplement. This is the "last mile" problem.

Assumption 3: Outsource Technology

Outsourcing has enabled firms to achieve scale, access best-in-class capabilities, and maintain operational resilience without the overhead of large internal technology functions. Asset owners have

historically maintained small technology teams, relying on vendors for platform configuration, integration, and support. Bespoke development has been reserved for only the most differentiated capabilities.

This assumption rests on a cost-benefit analysis that AI fundamentally alters. When the cost of creating bespoke analytical capability drops by an order of magnitude, the threshold for what is worth building in-house shifts accordingly.

Three Strategic Shifts

The emergence of production-grade AI tools creates the conditions for three strategic shifts.

Shift 1: From Enterprise Monolith to AI Agent Ecosystem

AI agents offer specialist capabilities that can be deployed, iterated, and retired independently. The enterprise platform retreats to its core competency in operational processing, trade lifecycle, and regulatory reporting.

Dimension	Conventional	AI-Enabled
Platform Scope	Spans all investment functions	Retreats to core operations and data
Analytical Capability	Configured within platform constraints	Delivered by specialist AI agents
Iteration Speed	Vendor release cycles (quarterly/annual)	Continuous agent refinement (days/weeks)
Customization	Limited by platform architecture	Purpose-built for each use case

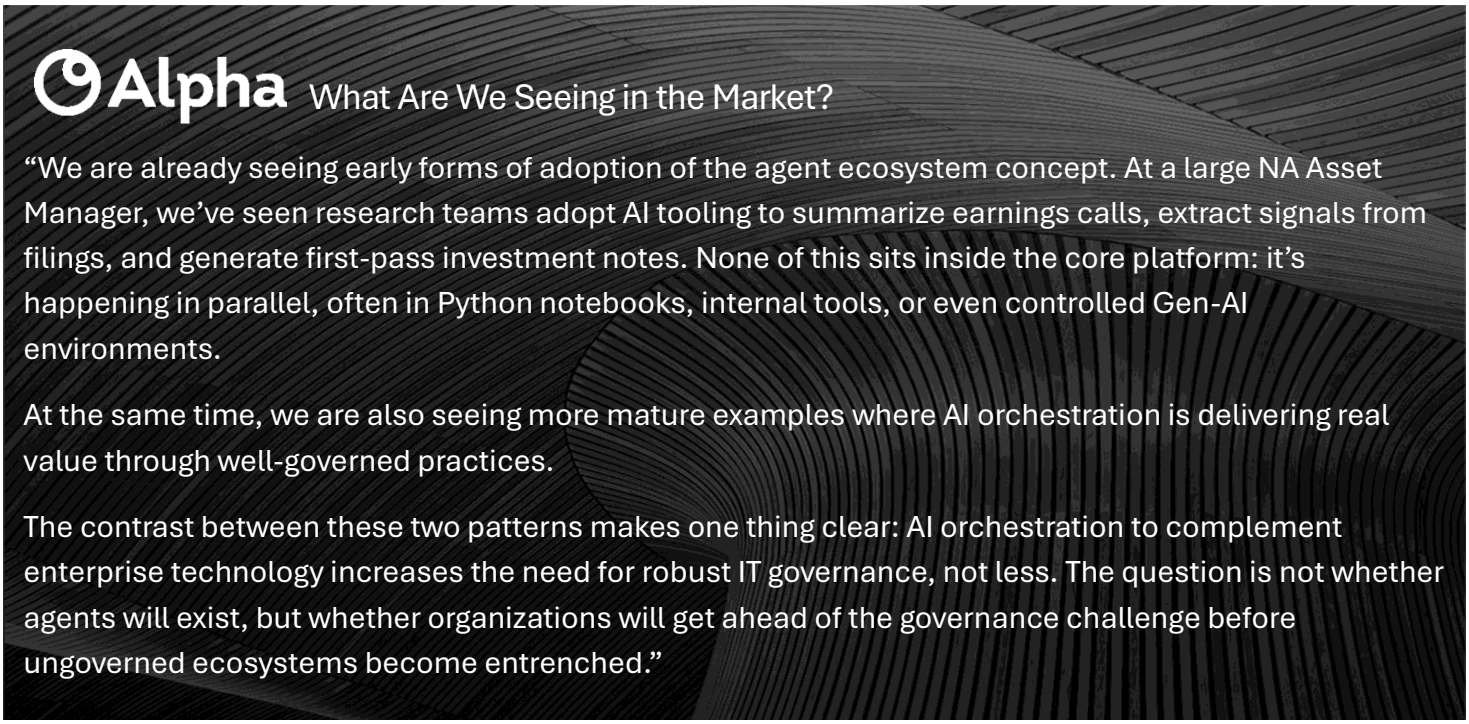

Why it Matters

Investment professionals interact with specialist agents that understand their domain context, rather than navigating a general-purpose platform designed primarily for operations.

Reality AI Must Overcome

The skeptic's case:

- Agent proliferation risks recreating the fragmentation enterprise consolidation was designed to eliminate
- Enterprise platforms provide a single point of accountability. In an agent ecosystem, "whose agent broke?" is a non-trivial operational risk at 3:55pm

The path through:

- Orchestration, observability, and governance must be treated as core infrastructure: agent registry with defined owners, validation gates, and lifecycle management
- Centralized audit trail of every agent invocation, with programmatic SLAs and automated alerting that surfaces issues before they reach downstream processes

Agents prove value in low-stakes analytical use cases first; enterprise platform retains operational processing until trust is demonstrably earned.

Shift 2: From Supply-Driven IBOR to Demand-Driven Data

AI has the potential to enable a demand-driven alternative where an agent identifies required data sources, applies mastering rules on demand, reconciles in real-time, and assembles a purpose-built dataset with full audit trail.

Supply-Driven IBOR	Demand-Driven IBOR
Source, reconcile, load on schedule	Assemble and deliver on demand when question is asked
Distribute pre-packaged views	Build contextualized datasets per use case
Data is always slightly stale	Point-in-time accuracy for the specific need
Manual blending for last-mile gaps	AI-managed assembly with audit trail

A critical requirement for any demand-driven model is point-in-time data integrity. Investment decisions, regulatory submissions, and audit reviews all require the ability to reconstruct what was known at a specific moment, not just what the data says today. AI-assembled data sets must therefore carry adequate data lineage.

Organizations adopting demand-driven approaches should treat lineage and auditability as first-class architectural requirements, not retrospective additions.

Alpha What Are We Seeing in the Market?

"In practice, what we're seeing at several large asset managers/owners is that the IBOR is still very much the control point. However, investment teams are increasingly bypassing it for analytical workflows because it doesn't provide the right granularity, it may have limited adaptability, or it may not handle private market complexity well. Analysts end up exporting IBOR data and blending it with external data and layering their own logic in Excel or Python. This is the impasse and it's absolutely real.

Leading investment organizations are addressing this directly through a 'data package' approach: blending IBOR data with other datasets to deliver a complete, curated data solution to the end consumer, rather than expecting investment teams to fend for themselves. Where AI comes into play is not to replace IBOR, but to formalize what is currently happening in an uncontrolled way. The headline: AI doesn't eliminate IBOR but can step in to replace the unmanaged layer sitting on top of it."

Why it Matters

Data mastering still matters. The difference is when, how, and for whom it is applied. The demand-driven model starts from the investment question and works backward.

Reality AI Must Overcome

The skeptic's case:

- The IBOR is a control framework, not a data warehouse. If teams assemble competing views with different on-the-fly mastering rules, you have competing versions of reality
- "Real-time reconciliation" understates the complexity: breaks are often genuine counterparty disagreements requiring human judgment, not faster pattern-matching
- Regulators expect a standing book of record with clear lineage, not an on-demand assembly
- The last-mile gap effects ~10% of use cases. Redesigning the entire architecture to optimize for that 10% risks the 90% that works reliably

The path through:

- The IBOR remains the golden source. AI agents operate as a consumption layer. They enrich, classify, and contextualize but cannot override reconciled positions or apply alternative logic
- AI accelerates the reconciliation workflow by automating the 85-90% of mechanical matching, freeing operations teams to focus on breaks requiring genuine expertise
- Every AI-assembled analytical view carries full machine-readable lineage, more transparent than the analyst spreadsheet workarounds it replaces
- AI solves the last mile on top of the IBOR, not instead of it, eliminating ungoverned manual workarounds without touching the operational foundation that surfaces issues before they reach downstream processes

Shift 3: From Outsource-First to AI-Augmented In-House

When AI tools can generate, test, and deploy capabilities with minimal human coding, the cost structure changes fundamentally, with differentiated capabilities now candidates for AI-augmented in-house development. Commoditized, transaction-heavy capabilities remain best served by established platforms.

Alpha What Are We Seeing in the Market?

"Right now, many asset managers/owners are experimenting with AI-augmentation but in a fragmented way. A portfolio manager may build an effective AI-powered tool, but the capability disappears once the individual leaves and isn't utilized for other desks / teams. We've seen this pattern repeatedly, and AI is only accelerating it. At the same time, vendors are rapidly embedding AI into their platforms, offering integrated workflows, built-in governance, and lower operational risk. As a result, firms are facing a renewed build-versus-buy decision: build for flexibility, differentiation, and speed; or buy for stability, support, and control. Leading firms are converging on a clearer model to navigate this trade-off:

1. Core operations (IBOR, accounting, trading) should remain platform-led with in-house AI supporting exception processes and oversight
2. Move differentiated analytics (alpha generation, portfolio insights) in-house with AI
3. Enable ad hoc / exploratory workflows through AI self-service

Why it Matters

Investment professionals with deep knowledge can build bespoke analytical tools that would previously have required months of vendor engagement and/or engineering effort.

Reality AI Must Overcome

The sceptic's case:

- AI lowers the cost of creating code, not operating it. "Analyst builds tool, analyst leaves, tool becomes unmaintainable black box" is a pattern AI accelerates
- Domain expertise alone is insufficient. Production-grade tools require data quality handling, edge case management, and operational resilience
- Vendors are building AI into their platforms, which may deliver 80% of the value at 20% of the risk

The path through:

- A tiered model redraws the build-vs-buy line deliberately: operational infrastructure stays on platforms; differentiated analytics are built in-house; ad-hoc tools become self-service where the alternative today is an ungoverned spreadsheet
- AI-generated code can be held to higher standards by default: standardized structure, automated testing, version control. Lower rebuild costs reduce black-box risk as replacement becomes viable
- The "investment technologist" role bridges domain expertise and technical fluency, ensuring in-house tools meet production-grade resilience and auditability standards

Today vs Tomorrow

Three scenarios illustrate the practical impact of these architectural changes:

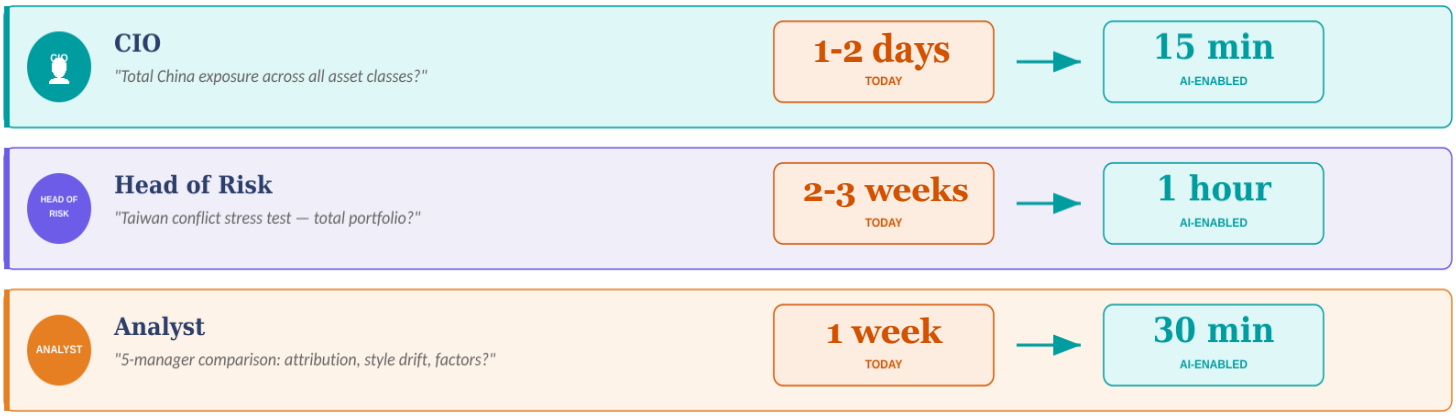

The CIO's AI-Enabled Morning

An AI-enabled morning might begin with the CIO reviewing an AI-generated portfolio briefing summarizing overnight market activity, current portfolio exposures, and emerging macro signals. An internal research assistant aggregates insights from earnings transcripts, internal research notes, and macroeconomic datasets. Portfolio managers review AI-generated scenario analyses showing potential impacts from geopolitical developments. Portfolio construction tools simulate thousands of portfolio combinations in minutes. Compliance and operations teams review AI-driven exception alerts that are flexible and adaptive, rather than relying on pre-coded frameworks dependent on human setup.

A minute-by-minute walkthrough of how a CIO's question flows through an AI-enabled technology stack:

Time	Step	What Happens
8:15 AM	The Question	CIO asks: "Show me total portfolio semiconductor supply chain risk across all asset classes and external mandates."
8:16 AM	Agent Activation	Orchestration decomposes into sub-tasks across four asset classes, maps classifications, resolves look-through for external mandates.
8:18 AM	Data Assembly	IBOR agent sources from core platform, PE administration, infrastructure databases, and external manager portals. Real-time reconciliation with quality flags.
8:22 AM	Analysis	Specialist agents run in parallel: supply chain mapping, factor exposure decomposition, geographic concentration. AI contextualizes against geopolitical developments.
8:25 AM	Insight Delivered	Interactive dashboard with semiconductor exposure by asset class, geography, and risk factor. CIO asks "Reduce by 20%?" Agent models rebalancing options immediately.

Ten minutes. Not ten days. Same question, same rigor, same audit trail but fundamentally different speed and depth.

The Architectural Shift

The architectural implications of the three shifts can be summarized in a comparison between today's enterprise-centric model and the AI-enabled alternative:

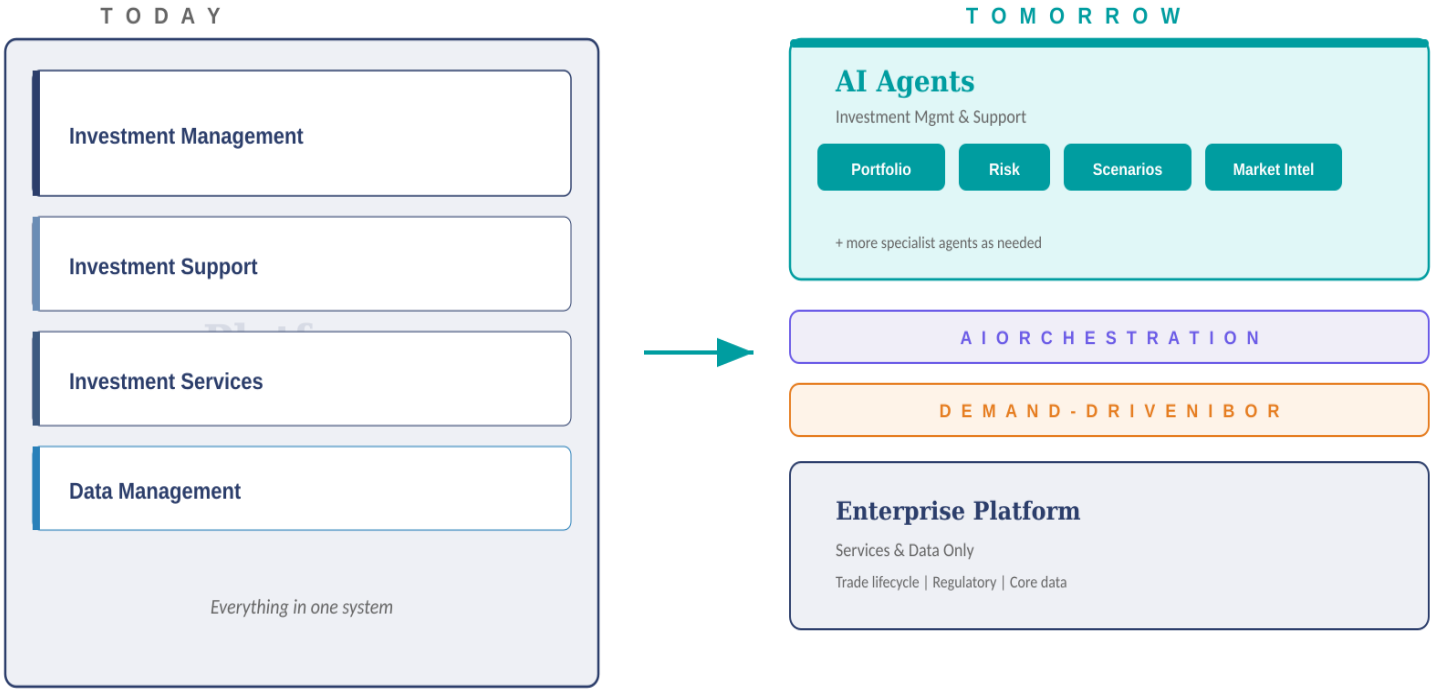

The degree of near-term AI impact varies significantly by business function. The following assessment reflects where AI-enabled architecture is most likely to transform existing workflows and team structures:

Business Function	AI Impact	Key Impact Areas
Research	**Critical**	AI fundamentally changes sourcing, synthesis, and presentation of investment research
Operations	**Critical**	Automation of reconciliation, exception management, NAV production, and reporting workflows
Portfolio Construction	**High**	AI augments scenario analysis and optimization; human judgment remains central
Risk Management	**High**	Real-time stress testing and exposure analysis; governance oversight still required
Trading	**Medium**	Near-term impact lower; increasing as execution analytics and liquidity tools mature
Compliance	**Medium**	Rule-based monitoring and exception alerting; regulatory interpretation remains human-led

Challenges and Risks

The transition to an AI-enabled architecture introduces significant risks that must be navigated with rigor and governance. Model risk is a first-order concern. AI agents are decision-making capabilities, and without rigorous documentation, performance monitoring, and drift detection, their outputs can degrade in ways that are invisible to the teams relying on them.

Governance gaps compound this. Agent proliferation without formal ownership or lifecycle management creates vacuums in accountability that are difficult to reverse once embedded in operational workflows. Regulators are actively developing expectations around AI explainability and auditability, and organizations should assume that scrutiny is coming rather than wait for frameworks to be finalized.

Data lineage sits at the center of all of this. Regulators expect firms to demonstrate not just what an AI output was, but why this was the case, tracing back through every data source, transformation, and decision rule that contributed to it. Building governance ahead of regulatory pressure, rather than in response to it, is the posture that will distinguish firms from transitioning well versus firms falling behind.

Risk Area	Severity	Key Considerations
Data Governance & Quality	**Critical**	AI-assembled data requires even higher data quality standards. Without robust data lineage, quality scoring, and governance, AI agents amplify problems rather than solve them.
AI Agent Governance	**Critical**	Each AI agent represents a decision-making capability that must be governed, validated, and auditable. Model risk management must extend to cover agent behavior.
Architectural Integrity	**Critical**	Architectural integrity is critical for AI enablement because it ensures scalable, secure, and interoperable foundations that allow AI capabilities to be deployed reliably and evolve without destabilizing core systems or creating 'shadow AI solutions'.
Regulatory Acceptance	**High**	Prudential regulators including APRA are still developing frameworks for AI. Organizations must demonstrate AI outputs meet the same standards of accuracy and explainability.
Integration Complexity	**High**	An AI agent ecosystem requires sophisticated orchestration. The challenge shifts from connecting to one platform to orchestrating a network of specialist capabilities.
Talent & Skills Gap	**High**	New hybrid roles required: AI-fluent investment professionals, orchestration architects, and investment technologists bridging domain expertise and technical capability.
Vendor Lock-in Shift	**Medium**	Lock-in risk shifts rather than disappearing. Organizations must assess AI agent dependencies and ensure portability from the outset.

The challenges above are amplified in private markets. Data is less standardized, valuations are infrequent, and coverage from third-party data providers is materially thinner than in public markets. At the same time, the analytical upside from AI is potentially the highest of any asset class. This combination makes governance doubly important: **the environments where AI can add the most value is also where data quality and lineage controls are hardest to establish.**

Operating Model Implications

An AI-enabled architecture requires changes across three dimensions of the operating model:

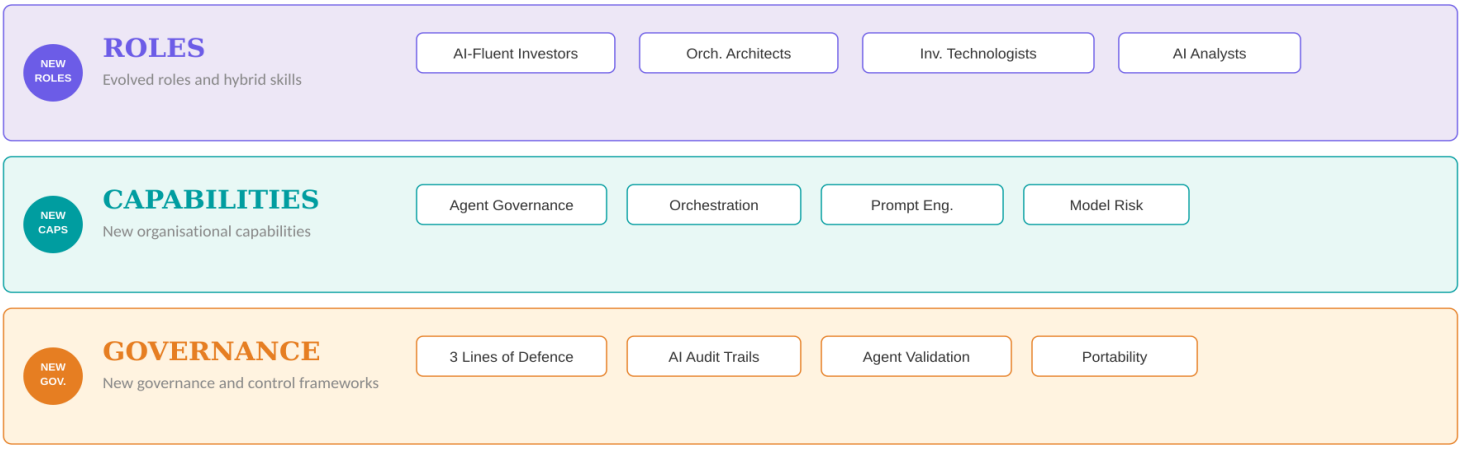

Evolved Roles and Skills

Investment professionals must become AI-fluent. Data engineers evolve into orchestration architects. A new "investment technologist" role emerges combining deep investment knowledge with technical capability in AI tools.

Alpha What Are We Seeing in the Market?

"We are already seeing early versions of these evolved roles in the market. At a large pension fund, a small team sits between investment and technology, helping translate investment questions into data and analytics workflows. At a hedge fund, "quant developers" are increasingly acting as orchestrators of tools rather than builders of models from scratch. What's changing with AI is that these roles become increasingly more focused around designing workflows, governing outputs, and ensuring reliability. In practice, the role of an investment professional fundamentally shifts to designing prompts and workflows, validating AI outputs, ensuring data lineage, and maintaining reusable analytical components."

New Organizational Capabilities

Organizations need capabilities in AI agent governance, orchestration management, prompt engineering, and AI model risk management, requiring a blend of investment domain expertise and AI fluency.

Governance and Control Framework

The three lines of defense models must extend to AI-specific risks: model validation for agents, vendor management for AI dependencies, and audit trail standards meeting regulatory requirements for explainability.

Where to Start?

Augment. Don't replace. The transition should follow a phased approach that proves value in bounded use cases and scales only when governance and readiness have been demonstrated.

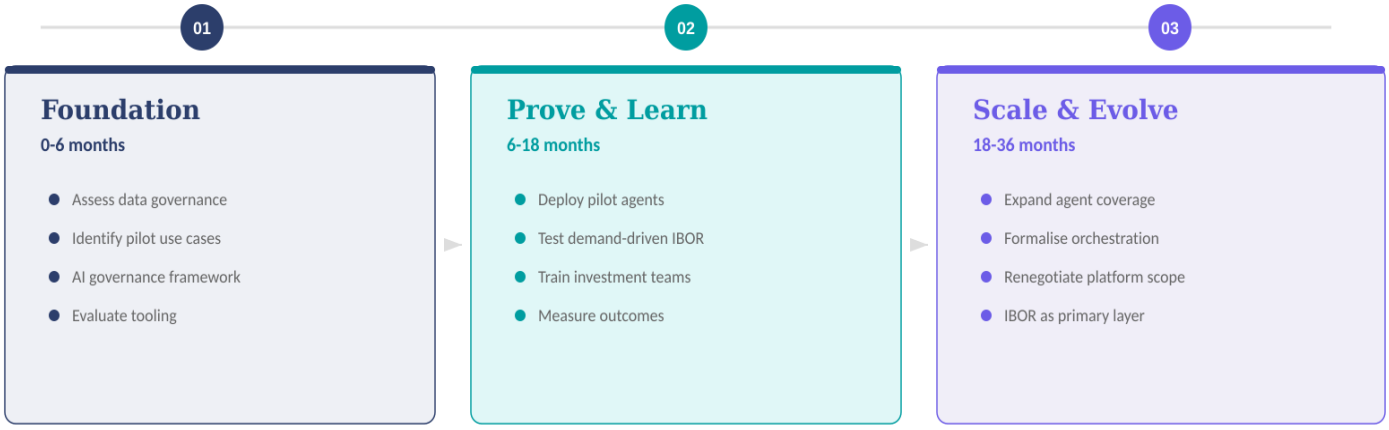

Meeting the Moment

What Does this Mean for Vendors?

Clearly the OMS and platform vendors have not been resting on their laurels and have made significant AI advances of their own, both in their SDLCs to accelerate product development as well as adding AI-driven tooling and functionality to enhance their software. Initially this had been with the slightly prosaic addition of the likes of copilot prompts, but increasingly with the release of complex AI model-driven analytics, document intelligence, workflow logic and AI/MCP enabled integration points.

In Alpha's recently released ***Outlook26***, our experts comment on where industry leaders are focusing.

> *"Adopting modular architectures – in many cases those based on a core 'front to back' hub – are differentiating due to their strong enterprise data platforms allowing for faster more agile API integrations."*

For further insights, read the full piece online here: Alpha's Asset & Wealth Management Outlook 2026

Conclusion

The AI inflection point does not render enterprise platforms obsolete. It does, however, fundamentally change the scope, role, and strategic importance of these platforms within asset managers' and owners' technology architectures. Organizations that approach this transition with discipline and pragmatism will unlock significant advantages: faster analytical capability, more responsive data delivery, and the ability to build differentiated tools reflecting their unique investment philosophy.

The path to sustainable AI adoption demands disciplined governance, uncompromising data integrity, and a deliberate progression, from proving value in contained deployments to building scalable foundations that enterprise-wide AI transformation requires.

How Alpha Can Help

Industry research suggests that 70-80% of enterprise AI initiatives fail to deliver meaningful business value. In financial services, these rates are often worse, with data fragmentation, regulatory requirements, and legacy architecture compounding the challenge. The gap between a compelling AI proof-of-concept and an operationalized capability embedded in investment workflows is where most programs stall.

Alpha has worked with clients navigating this exact transition. Our experts understand what separates initiatives that scale from those that stall. If you are evaluating where AI fits in your operating model or looking to course-correct a program that hasn't delivered, we would welcome the conversation.

For boards and executive teams, seven questions should guide the strategic conversation:

People

1. Are we investing in AI fluency for investment professionals, or only for technologists?
2. Do we have the hybrid roles - investment technologists, orchestration architects - needed to bridge domain expertise and AI capability?

Process

3. Could AI agents deliver our analytical capabilities faster and more flexibly?
4. Have we redesigned our investment and operational workflows to harness AI, or are we layering new tools onto legacy processes?

Technology & Data

5. Is our enterprise platform delivering genuine value, or just expensive data infrastructure?
6. Is our data governance mature enough for AI-assembled, demand-driven data?
7. What would our technology architecture look like if we started from scratch today?

Organizations that ask these questions now, and act on the answers with urgency and discipline, will be best positioned to navigate the architectural transformation that AI demands.

Contributors

Bruce Russell is a Senior Partner and Alpha's Global Head of Asset Owner Client Segment. He possesses an extensive background in management consulting, having served in a variety of leadership roles in the space. Bruce has driven several successful client transformations during that time, acting as a trusted advisor to executive leaders across the industry.

Laura Kayrouz is a Senior Partner and Alpha's Global Co-Head of Investments. She brings over 20 years of experience in front office and technology transformation, with expertise in agile delivery and business process redesign. Her background includes significant program delivery roles supporting investment platform implementations.

Kevin Faulkner is a Partner based in Sydney, and Alpha's Head of Investments for APAC. With a technology background, Kevin has 25+ years' experience across the asset management industry, Big 4 consulting and platform transformation. As APAC Investments Practice Lead, Kevin is tasked with leading the transformation of our clients' platforms and operating models.

Chris Rushworth is a Director at Alpha FMC based in London, where he advises asset managers, asset owners, and asset servicers on data strategy, target operating model design, and technology-led transformation. His work focuses on helping investment operations leaders cut through the complexity of legacy environments, simplifying data architectures, reducing total cost of ownership, and building operating models that are fit for a cloud-native future.

Luke Ferreira is a Consultant at Alpha FMC based in Boston, Massachusetts, and is aligned to Alpha's Investments Practice. He has supported a wide range of engagements, with experience across enterprise strategy, operating model design, vendor evaluation and selection, and large-scale implementation delivery.