

publicis
sapient

canoe

The Intelligence Era in Asset Management

Industry playbook for how asset managers can move from fragmented AI initiatives to intelligent operating systems

Building an enterprise designed for intelligence

What will it take to build an enterprise that can compete in the age of AI?

Over the past year, we brought together senior leaders from across the asset management industry to explore this simple but important question. While firms were at different stages of AI maturity, executives consistently described many of the same challenges. The discussion shifted from individual AI use cases to a broader question: How do organizations need to evolve if intelligence is to become an enterprise capability rather than another point solution?

Across those conversations, the same barriers kept coming up. Despite differences in size, operating model and technology landscape, leaders pointed to the same structural barriers preventing AI from delivering enterprise-wide value. More importantly, they increasingly saw AI not simply as another technology, but as a catalyst for rethinking how knowledge, decisions and expertise move across the organization.

This paper explores that shift, and what it will take for asset managers to move from fragmented AI initiatives to an enterprise designed for intelligence.

Table of contents:

- 04 Why is AI hard to scale in asset management?
- 08 From fragmentation to intelligence
- 09 What are the four layers of enterprise intelligence?
- 10 Where should asset managers start with enterprise AI?
- 14 Six moves leaders must make to scale
- 15 Build what AI can't buy

Why is AI hard to scale in asset management?

Access to AI may be getting easier, but scaling it is not. Firms still need to connect fragmented data, knowledge, workflows and governance throughout the enterprise.

Across asset management firms with different operating models, technology landscapes and levels of AI maturity, executives identified eight themes shaping intelligence:

01 AI is advancing faster than organizations are changing

Technology capabilities are evolving rapidly, while operating models, governance and ways of working struggle to keep up.

02 Intelligence remains fragmented across the enterprise

Critical knowledge sits across people, functions, platforms, documents and partners, making it difficult to access, reuse and apply consistently.

03 Data alone is no longer enough

Years of data modernization have created stronger foundations, but data by itself doesn't create better decisions.

04 AI initiatives are repeatedly rebuilding the same foundations

Teams recreate prompts, governance models, integrations and business logic for individual use cases, slowing delivery and making AI harder to scale.

05 Operational complexity has become the biggest barrier to transformation

Adding AI to fragmented processes can create more complexity rather than less, making workflow redesign as important as the technology itself.

06 Competitive advantage is shifting from models to enterprise intelligence

As foundation models become widely accessible, differentiation increasingly comes from combining proprietary data, business knowledge and operational context.

07 Governance is more important than ever

Risk, compliance and oversight need to sit inside enterprise workflows so intelligence can scale safely and consistently.

08 Building intelligence is fundamentally an operating model transformation

AI changes how work, decisions and expertise move through the organization.

Taken together, these themes suggest the industry isn't struggling because it lacks AI investment or capable models. Intelligence remains fragmented as it tries to move through organizations that were never designed to connect and apply it at enterprise scale.

Three structural barriers explain why:

01 Intelligence remains trapped inside organizational silos

Asset managers possess enormous amounts of knowledge. Investment research, portfolio insights, regulatory interpretation, operational expertise, client interactions and historical decisions all represent valuable intellectual capital. Yet much of this intelligence remains isolated within business functions, individual teams or technology platforms.

AI can generate new knowledge, but it cannot easily access intelligence that remains fragmented across disconnected systems, inconsistent processes and institutional memory. As a result, organizations repeatedly solve the same problems, recreate the same analyses and make similar decisions without benefiting from collective enterprise learning.

Until organizations make enterprise intelligence discoverable, connected and reusable, AI will continue to amplify fragmentation rather than eliminate it.



The data is not centralized and just to gather the data you need to go through many different tools and places. It's just very hard.

– Lucas Cortini, Investment Product Analyst, Schroders

02 Organizations optimize individual tasks instead of redesigning decision-making

Many AI initiatives begin by automating existing activities. While these improvements create measurable productivity gains, they rarely change how the organization actually operates.

The executive discussions highlighted a common challenge: many firms have become better at accelerating individual tasks without fundamentally redesigning how decisions are made across the enterprise.

True transformation requires organizations to rethink entire workflows, not simply replace isolated activities with AI.

The greatest opportunity lies in redesigning how intelligence flows between people, systems and decisions, enabling humans and AI to work together as part of a connected operating model.



Teams are spending hundreds of hours every month collecting documents, storing them, but then having to type data points from a PDF into a downstream system. It's a workflow that shouldn't exist today.

– Oliver Wedlake, Director of Sales, EMEA, Canoe

03 Every AI initiative starts from scratch

Perhaps the most consistent frustration expressed by executives was the repeated effort required to deliver every new AI capability, which wastes valuable time and resources.

The result is an organization that scales AI projects rather than enterprise intelligence—and these are often doomed to fail. As AI adoption accelerates, this fragmented approach becomes increasingly expensive, difficult to govern and almost impossible to industrialize.

Leading firms are beginning to recognize that intelligence itself must become a shared enterprise capability, supported by common governance, reusable context and operational standards that allow every new initiative to build upon the last rather than start again.



The repetition of things that go on throughout the organization is such a massive waste. But at the same time we can't seem to get organized to show these are the central things that we're going to do to enable people to experiment in a safe way. These are the standard kits, that means we only have to do that once.

– Yugo Ashida, Head of ISS Enterprise Architecture, Fidelity International



”

I’ve personally seen tens of AI pilots fail. The models and tools are amazing. They go through a two-, three- or four-week proof of concept and get there, and then someone asks, ‘How are you going to get that data into one place?’ There’s no answer. So we see a lot of projects just fall down there.

– Cloud Architecture Leader

Taken individually, each of these challenges appears manageable. Together, however, they expose a far more fundamental issue.

The traditional enterprise was designed to move information between systems and people. The enterprise must be designed to move intelligence, continuously, securely and at scale.

This requires more than new technology. It requires a different architectural foundation—one built not around applications or processes, but around the creation, orchestration and application of enterprise intelligence.

From fragmentation to intelligence

The debate over whether organizations should adopt AI has largely been settled. The question now is how enterprises need to change so intelligence can move continuously across people, systems and decisions.

For decades, technology transformation has focused on digitizing processes, integrating applications and modernizing data. The intelligent enterprise operates differently. Rather than treating intelligence as something created by individuals or isolated applications, it treats intelligence as a strategic enterprise asset, one that is continuously captured, connected, governed and reused across every business function.

This changes where competitive advantage comes from. As foundation models become more powerful, accessible and interchangeable, the model itself becomes less differentiating. Sustainable advantage comes from what surrounds it: proprietary knowledge, business context, operational expertise, governance and the ability to orchestrate across the enterprise.



The organizations that succeed won't be those with the most AI. They'll be the ones that create an enterprise where every decision, every interaction and every outcome makes the organization smarter.

– Richard Doherty, Head of Asset and Wealth Management, Publicis Sapient

What are the four layers of enterprise intelligence?

While every organization will evolve differently, the architecture of intelligence rests on four interconnected layers: knowledge, intelligence, orchestration and learning.

01 Enterprise knowledge

The foundation is the organization's collective knowledge: structured and unstructured data, business policies, investment research, regulatory guidance, operational documentation, client interactions and institutional expertise. Rather than remaining isolated within individual functions, this knowledge becomes shared [enterprise context](#) that the business can access and reuse consistently.

02 Enterprise intelligence

Knowledge alone doesn't create value. It becomes intelligence when information is interpreted, connected and enriched with business meaning. AI agents, reasoning models, workflows and domain expertise turn raw information into trusted insights that can support decision-making across functions.

03 Enterprise orchestration

Intelligence only creates value when it influences decisions and actions. Orchestration connects intelligence with business processes, human expertise and enterprise systems so the right insight reaches the right person or agent at the right point in the workflow. Governance, security and compliance sit inside this layer, allowing intelligence to move safely across the organization.

04 Enterprise learning

Intelligence doesn't stop once a decision is made. Every decision, exception, interaction and outcome creates feedback that strengthens the enterprise over time. Learning becomes a continuous organizational capability rather than a periodic exercise.

Together, these layers create a closed loop: knowledge becomes intelligence, intelligence drives decisions and actions, and the results create new knowledge that improves what happens next.



The model is only one component. What creates lasting value is the system around it—the knowledge, context, governance and workflows that allow intelligence to flow across the enterprise.

– Head of AI & Architecture, Leading Asset Manager

Where should asset managers start with enterprise AI?

Once the architecture is mapped, the next challenge is deciding where to begin. This was one of the most debated topics in our executive discussions. While every organization faced different commercial priorities, regulatory obligations and technology landscapes, there was striking agreement on the approach.

The firms making the greatest progress aren't attempting to deploy AI across every function at once or pursuing isolated proofs of concept with no path to enterprise scale. Instead, they're deliberately selecting a small number of high-value use cases that could solve immediate business challenges while simultaneously creating reusable enterprise intelligence.

The goal is to establish the foundational capabilities that can be reused across future AI initiatives, helping teams adopt AI faster while reducing implementation effort, governance overhead and operational complexity.

Viewed in this way, every successful use case becomes another building block in the enterprise intelligence architecture.

What makes a strong starting point?

Across our discussions, executives consistently prioritized use cases with four characteristics:

01 They operate across multiple business functions

The greatest opportunities often sit where intelligence must move between investment teams, operations, compliance, risk and technology. These cross-functional workflows expose the fragmentation that prevents organizations from operating as a connected enterprise.

02 They rely on complex knowledge rather than simple automation

Processes involving regulations, investment mandates, client documentation or operational policies benefit disproportionately from AI because they depend on reasoning, context and domain expertise rather than repetitive task execution.

03 They combine measurable business value with reusable capabilities

The strongest initiatives deliver immediate operational improvements while creating assets that other teams can use, from governed knowledge repositories to decision logic and shared enterprise context.

04 They strengthen enterprise governance from the outset

Successful programs embed governance, human oversight and auditability into their operating model from day one. That foundation makes it possible to reuse intelligence safely across the organization.

Focus on priority capability areas

While every organization will define its own roadmap, our discussions highlighted several domains where firms are already generating measurable value while establishing reusable enterprise capabilities.

- **Regulatory interpretation and policy intelligence**
- **Investment guideline management and monitoring**
- **Investment research and knowledge management**
- **Client servicing and relationship intelligence**
- **Operational exception management**
- **Enterprise search and knowledge discovery**
- **Workflow orchestration across investment operations**
- **AI-enabled decision support for portfolio, risk and compliance teams**

These aren't simply AI use cases. Each implementation strengthens the shared knowledge, governance, workflows and decision capabilities upon which future initiatives build. Over time, those individual initiatives form a connected intelligence ecosystem: knowledge becomes easier to discover, decisions become more consistent and the organization becomes progressively better equipped to deploy AI at scale.

USE CASE:

Guideline intelligence agent turns guidelines into action with up to 70 percent fewer interpretation errors

Problem: For most firms, creating and interpreting investment guidelines is still a manual process, slow and prone to errors. This process varies across teams and individuals, has no audit mechanism and cannot scale across the enterprise. This creates a bottleneck at compliance, one of the most crucial functions in the business.

Opportunity: Agentic AI provides a unified system capable of multi-step reasoning, validation and contextual decision-making with speed, consistency and auditability.

Manual guideline monitoring



Agentic guideline monitoring

with human in-the-loop governance



What is the guideline intelligence agent? The agent solves the compliance challenge by accurately interpreting and operationalizing investment guidelines at scale. Ingesting huge volumes of unstructured data, it rapidly creates investment guidelines with real-time validation and monitoring. And it does all this with human oversight, automating guideline interpretation while also incorporating confidence scores to indicate where human interpretation is needed. Its built-in auditable reasoning trails align with regulators' expectations, transforming compliance from a function that exists to find errors to one that intelligently prevents them.

Impact: Guideline intelligence agent delivers a robust ROI by lowering risk and cost while improving time to market. By moving away from individual judgment, the agent improves regulatory confidence and reduces interpretation-related issues by up to 70 percent. This also cuts down on the manual interpretation workload to keep teams lean and their work scalable. And with onboarding times slashed to hours instead of weeks, guidelines can be updated faster, products can be launched quicker and investigations can happen sooner.



”

By building the intelligent infrastructure to reduce manual guideline monitoring activity, we are releasing up to 20 hours of operational capacity each day.

– Richard Doherty, Head of Asset and Wealth Management, Publicis Sapient

Six moves leaders must make to scale

A strong use case can prove the value of enterprise intelligence. But scaling that value requires something more: a common operating foundation that prevents each new initiative from becoming another point solution.

Here are six moves leaders should make to scale:

- 01 Set one enterprise ambition.** Define the value target in productivity, control quality, service responsiveness and time-to-market rather than counting pilots.
- 02 Fund the platform first.** Approve budget for data products, context services, model gateway, evaluation harness and governance tooling before approving a long tail of isolated use cases.
- 03 Name one accountable executive.** Split accountability slows scale. One executive should own platform outcomes end to end, with domain business owners accountable for adoption and value.
- 04 Mandate governed context.** Require that priority AI workflows use approved data products, source retrieval, permissions and evidence capture.
- 05 Adopt supervised autonomy as policy.** Draw clear thresholds for what can draft, what can recommend, what can act under approval and what must remain fully human-led.
- 06 Insist on a 90-day result.** The first production workflow should be live in one quarter with a clear value baseline, quality target and control dashboard.

Together, these actions provide a practical foundation for building enterprise AI capability in a sector where regulators already expect strong governance, traceability, testing and management-body oversight.

Sapient Bodhi is Publicis Sapient's agentic platform built for AI orchestration at enterprise scale. It brings agents, models and workflows into a single system designed to execute business processes according to embedded logic and rules without cloud or model lock-in. Bodhi's library of pre-built, customizable agentic solutions runs with the right enterprise context, policies and data from the start. Teams design, run and monitor their agents from a centralized dashboard, making it easier to track value and manage risk across the business.

Build what AI can't buy

Asset managers need an enterprise AI platform that turns data, enterprise context and organizational knowledge into better decisions at scale.

That foundation starts with trusted data, shared enterprise context and organizational memory. It combines agentic architecture that coordinates work across the enterprise with governance built into every interaction and workflow. Together, these capabilities make AI more accurate, more transparent and easier to scale in a highly regulated industry.

As foundation models become widely available, competitive advantage will come less from the models themselves and more from the enterprise operating system built around them. Firms that establish that foundation today can be better positioned to deploy new AI capabilities faster, strengthen governance and improve decision-making without rebuilding the platform for every new use case.



Start building your intelligence advantage

Publicis Sapient helps asset managers identify where intelligence can create the greatest business value and deliver the platform, governance and operating model required to compete in the age of AI.

Speak with our financial services and AI transformation leaders:



David Murphy
Head of Financial Services,
EMEA & APAC



Dan Pitchenik
Head of Financial Services,
North America



Richard Doherty
Head of Asset and
Wealth Management



Pinak Kiran Vedalankar
Group Vice President, Technology,
Financial Services

[Book a strategy session with our experts](#)



Publicis Sapient is a technology company that provides enterprise AI platforms and services. With over 30 years of digital business transformation experience, we enable enterprise clients to transform how they operate and serve their customers, unlocking new value and enabling them to thrive in an AI-driven world. Our platforms Sapient Slingshot, Sapient Bodhi and Sapient Sustain use AI built off this deep enterprise context to help organizations modernize their legacy tech systems, build agentic solutions, and automate their IT operations. The combination of our AI platforms and the expertise of our people enables us to deliver faster and more effective outcomes through solutions that are specific to the unique needs of our clients' businesses, their industries and their customers. Publicis Sapient is the technology hub of Publicis Groupe, uniting 20,000 people worldwide across 28 countries.

For more information, visit publicissapient.com.