KPMG

Unifying the asset management value chain with AI-ready data

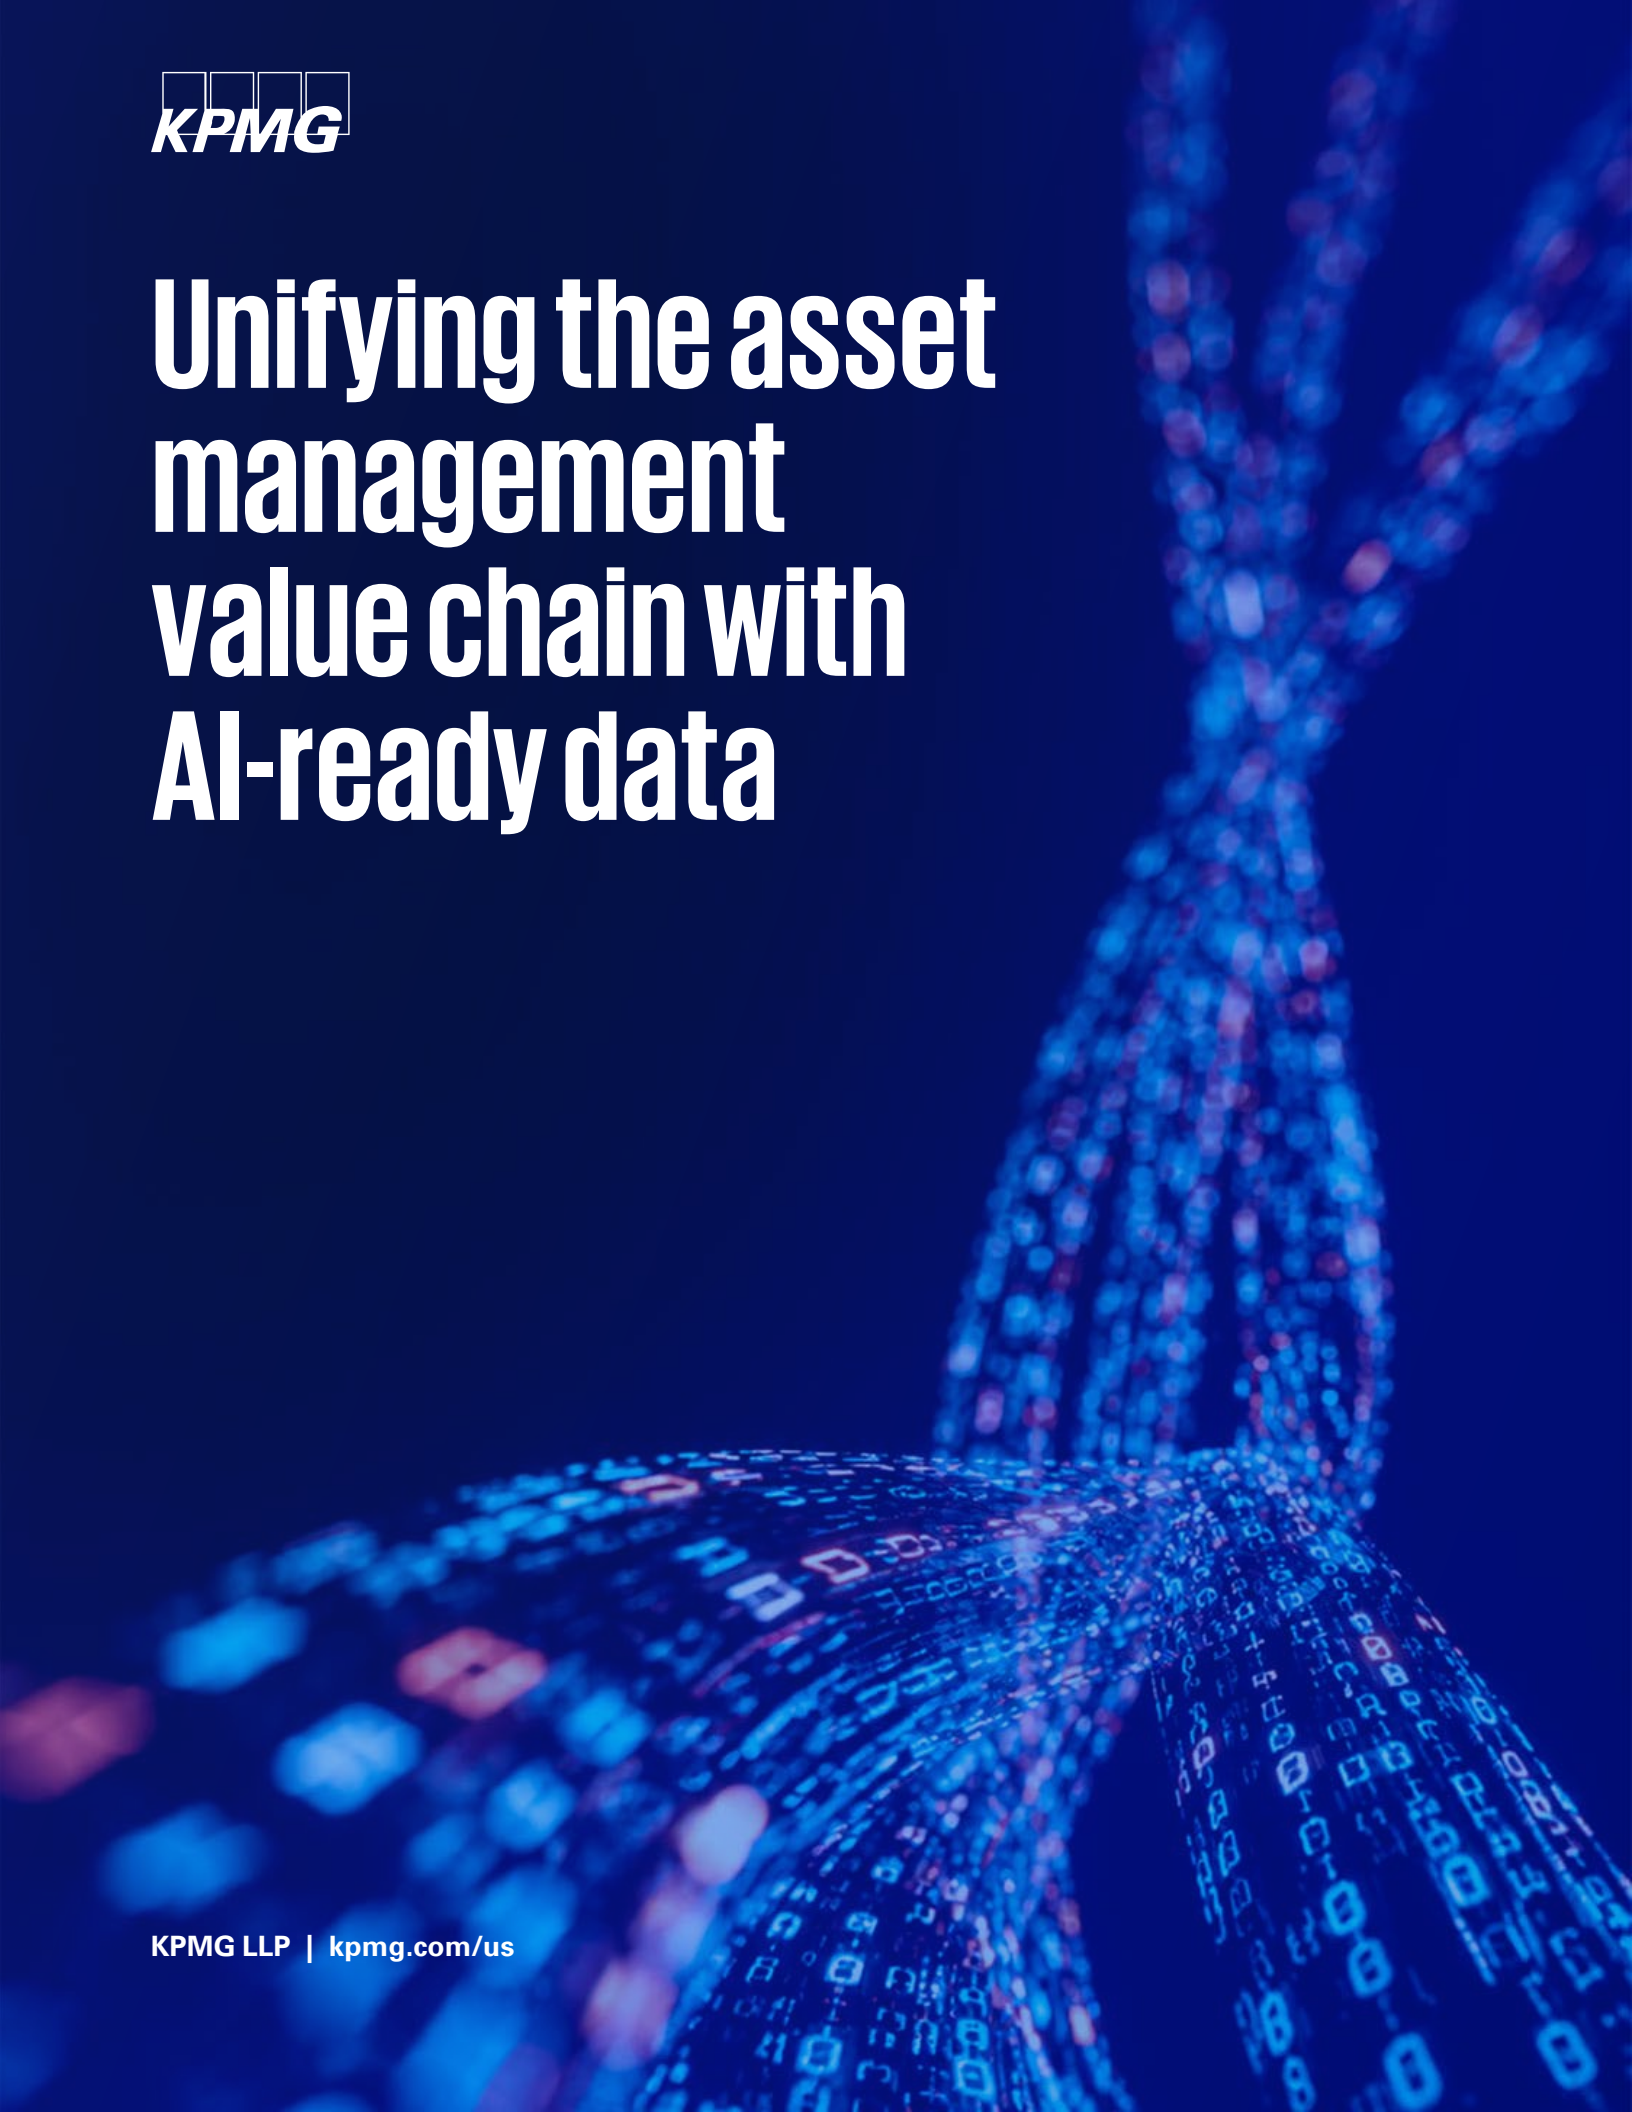

KPMG LLP | kpmg.com/us

Despite heavy investment, many asset management firms are struggling to achieve significant, sustainable ROI from their artificial intelligence (AI) initiatives. AI initiatives often stall at the pilot stage, and scalability remains elusive. The root cause is a foundational weakness: Data is fragmented and siloed because each department processes data independently.

The front office's pricing data might be updated at a different frequency or cleansed using different rules than the back office's data procedures. No single department can claim to have the golden source for truth, leading to constant data reconciliation challenges and a lack of confidence in analytics and reporting across the firm.

Unifying fragmented data is only the first step. To achieve meaningful ROI from AI, asset managers must ensure their data is AI-ready, which differs from the data traditionally used in spreadsheets and dashboards. It is machine-readable and prepared for retrieval, reasoning, composition, and governed execution.

By harnessing proprietary data and strengthening the underlying data foundation, asset managers can move beyond incremental gains and unlock the transformative potential of AI.

The goal is to reframe the conversation today:

From: **"What can AI do?"**

To: **"What must our data be to power AI effectively?"**

Building a unified, AI-ready data foundation is the most critical factor for success.

In this paper, we will explore how cultivating a foundation of semantically consistent, observable, and machine-consumable data is the key to unlocking compounding value, transforming the entire asset management value chain, and establishing a sustainable competitive advantage.

What is AI-ready data?

AI-ready data is contextualized with a semantic layer, a knowledge graph, and an ontology that explicitly defines business rules and relationships. This foundation is embedded with trust and provides a clear, auditable trail to verify AI reasoning. Ultimately, it allows probabilistic AI agents to operate safely and reliably within deterministic business logic.

Build a data foundation that can deliver AI-ready data

Building an AI-ready data foundation requires shifting from legacy, siloed infrastructure to an integrated, governed, and contextual architecture. This transition ensures enterprise data is searchable, contextualized, and trusted for AI agents, retrieval-augmented generation (RAG), and autonomous workflows.

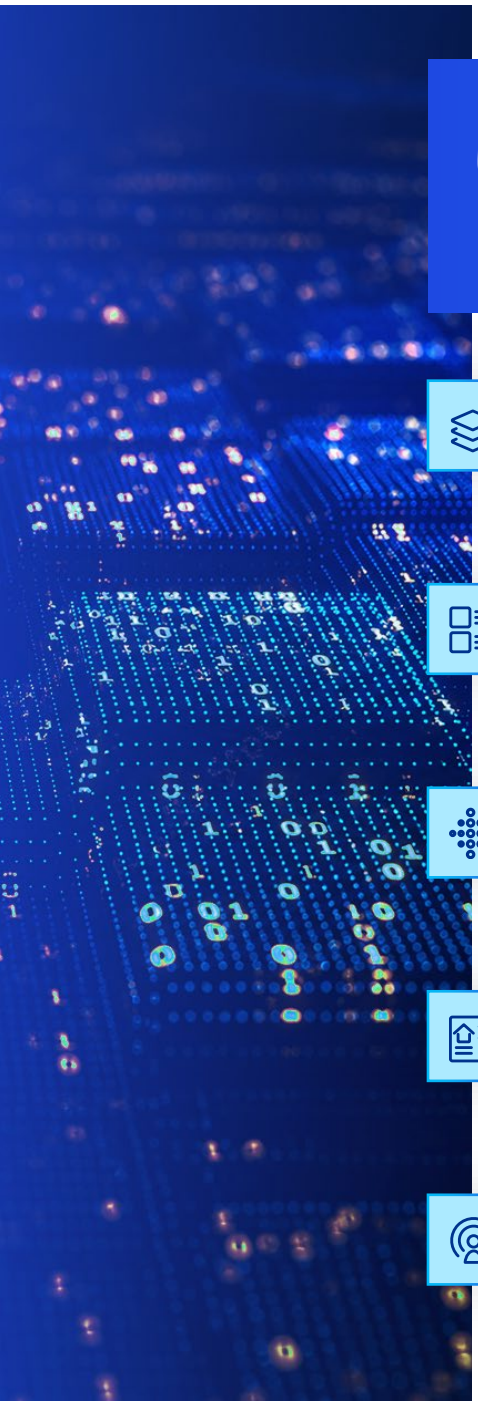

Five steps to building an AI-ready data foundation

- **Determine your golden source**
- **Standardize data definitions and semantic layers**
- **Build AI-ready data products and pipelines that flow throughout the value chain**
- **Implement robust governance, lineage, and security**
- **Design an agentic workflow factory**

Determine your golden source

AI-ready data begins with establishing authoritative sources for critical business information and creating the infrastructure needed to connect data across the enterprise. Rather than maintaining redundant copies of data across disparate systems, organizations should create a trusted foundation that enables consistent access, context, and interoperability.

Connect, don't copy: A data fabric allows AI models to query data in real time from its original location (whether a legacy database or a cloud application) without having redundant copies and variations in different systems.

Single source of truth: When an AI agent pulls a client profile or a trade execution record, it is accessing the single, undisputed golden source of that data, eliminating reconciliation errors.

Trust and governance by design: Every data asset should include clear ownership, lineage, quality standards, and access controls. AI agents must be able to trace where data originated, understand who can use it, and validate that it meets governance requirements. This transparency strengthens trust, reduces risk, and supports regulatory and fiduciary obligations.

Standardize data definitions and semantic layers

AI can only act on the information it receives. To enable consistent and reliable outcomes, data leaders must establish a semantic layer that standardizes business definitions across the firm and creates a shared understanding of data.

Context is key: AI can only generate accurate, reliable outputs when it understands the business meaning behind the data it consumes. If the front office defines a term differently than the back office, then AI's interpretation will be inconsistent. A semantic layer provides a consistent business vocabulary that maps technical fields to business terms, ensuring data is interpreted consistently across the enterprise.

With contextualized data, AI is equipped to reason, and most importantly, to act reliably, a prerequisite in asset management. For example, a trade, portfolio position, or client relationship may have different meanings and requirements depending on whether it is being used by portfolio managers, operations teams, risk professionals, or compliance functions.

Ontology: Building on the semantic layer's shared business vocabulary, an ontology defines business concepts, relationships, and rules in a machine-readable format. It establishes how entities interact and the logic that governs those interactions. For example, it can define relationships among clients, accounts, funds, trades, and portfolio positions. This additional layer of context enables AI to understand not just what data represents, but how and why it is connected, supporting more sophisticated reasoning and decision-making.

Agentic harness: The harness controls the planning and execution of agentic behaviors, using the context provided by the ontology to guide actions safely and reliably. It acts as a structured framework for AI agents, helping ensure that even as they operate autonomously, they remain aligned with business goals and constraints.

Build AI-ready data products and pipelines that flow throughout the value chain

Data powers every function across the investment cycle. To support AI, organizations must access and share data in a federated capacity and treat data as a product.

Federated data access: Enable data to remain in its source systems while making it accessible across the enterprise. Rather than creating redundant copies, AI systems should access trusted data where it resides, using common definitions and business context to ensure consistent interpretation across front-, middle-, and back-office functions. This approach helps maintain data quality, reduces reconciliation challenges, and ensures AI operates from a single version of the truth.

Data as a product: A data product is a collection of data assets that have been taken from raw to curated status, for a specific business purpose. It's built for reuse and actively managed from definition through retirement. The more that a data product can be reused or monetized, the more it's worth. Data products that can create value simply, reliably, and at scale are the most prized.

Implement robust governance, lineage, and security

The back office demands trust, and regulators demand auditability. Organizations must be able to demonstrate exactly where AI-generated insights originated and how they were produced.

Data governance: Establish clear ownership and accountability for all high-value data assets. This involves assigning data stewards who are responsible for creating and enforcing policies on data quality, usage rights, and lifecycle management.

Data lineage: Use automated tracking so that every output generated by an AI agent can be traced back to its raw data source. This provides a transparent audit trail for regulators and allows internal teams to quickly debug, audit, and validate AI-generated outputs.

Context-aware access controls: AI agents should operate under context-aware access policies that govern both human and nonhuman identities. Access decisions should consider data sensitivity, business purpose, and regulatory requirements to ensure agents only access information needed to perform their task. For example, an agent supporting portfolio analysis may require access to holdings and market data, while another supporting compliance functions may require access to supervisory and regulatory information. Continuous monitoring and telemetry provide visibility into agent activity, strengthening trust and supporting auditability.

Design an agentic workflow factory

AI is changing how people interact with technology. Rather than relying solely on applications, workflows, and manual processes, users increasingly engage with AI through conversational and agent-driven experiences. As firms adopt AI agents at scale, they should establish repeatable frameworks for designing, governing, and deploying agentic workflows. Clear governance models should define when agents can act independently, when human oversight is required, and when approval is needed before execution.

An AI agent can analyze market and portfolio data in the front office, identify exceptions in the middle office, and prepare reconciliations or reports in the back office. Depending on the task and risk level, AM firms can choose the level of human oversight of agents, including full human oversight (Human-in-the-Loop [HITL]), shared oversight (HITL Lite), and exception-only intervention (Human-out-of-the-Loop).

From cost reduction to compounding value

Early AI initiatives have largely focused on productivity gains, cost reduction, and operational efficiency. As AI adoption matures, asset managers are increasingly expected to demonstrate measurable business value and ROI. Yet the greatest opportunity extends beyond cost savings. When applied across the investment lifecycle, AI can accelerate revenue growth, reduce risk, and enable sustainable scale.

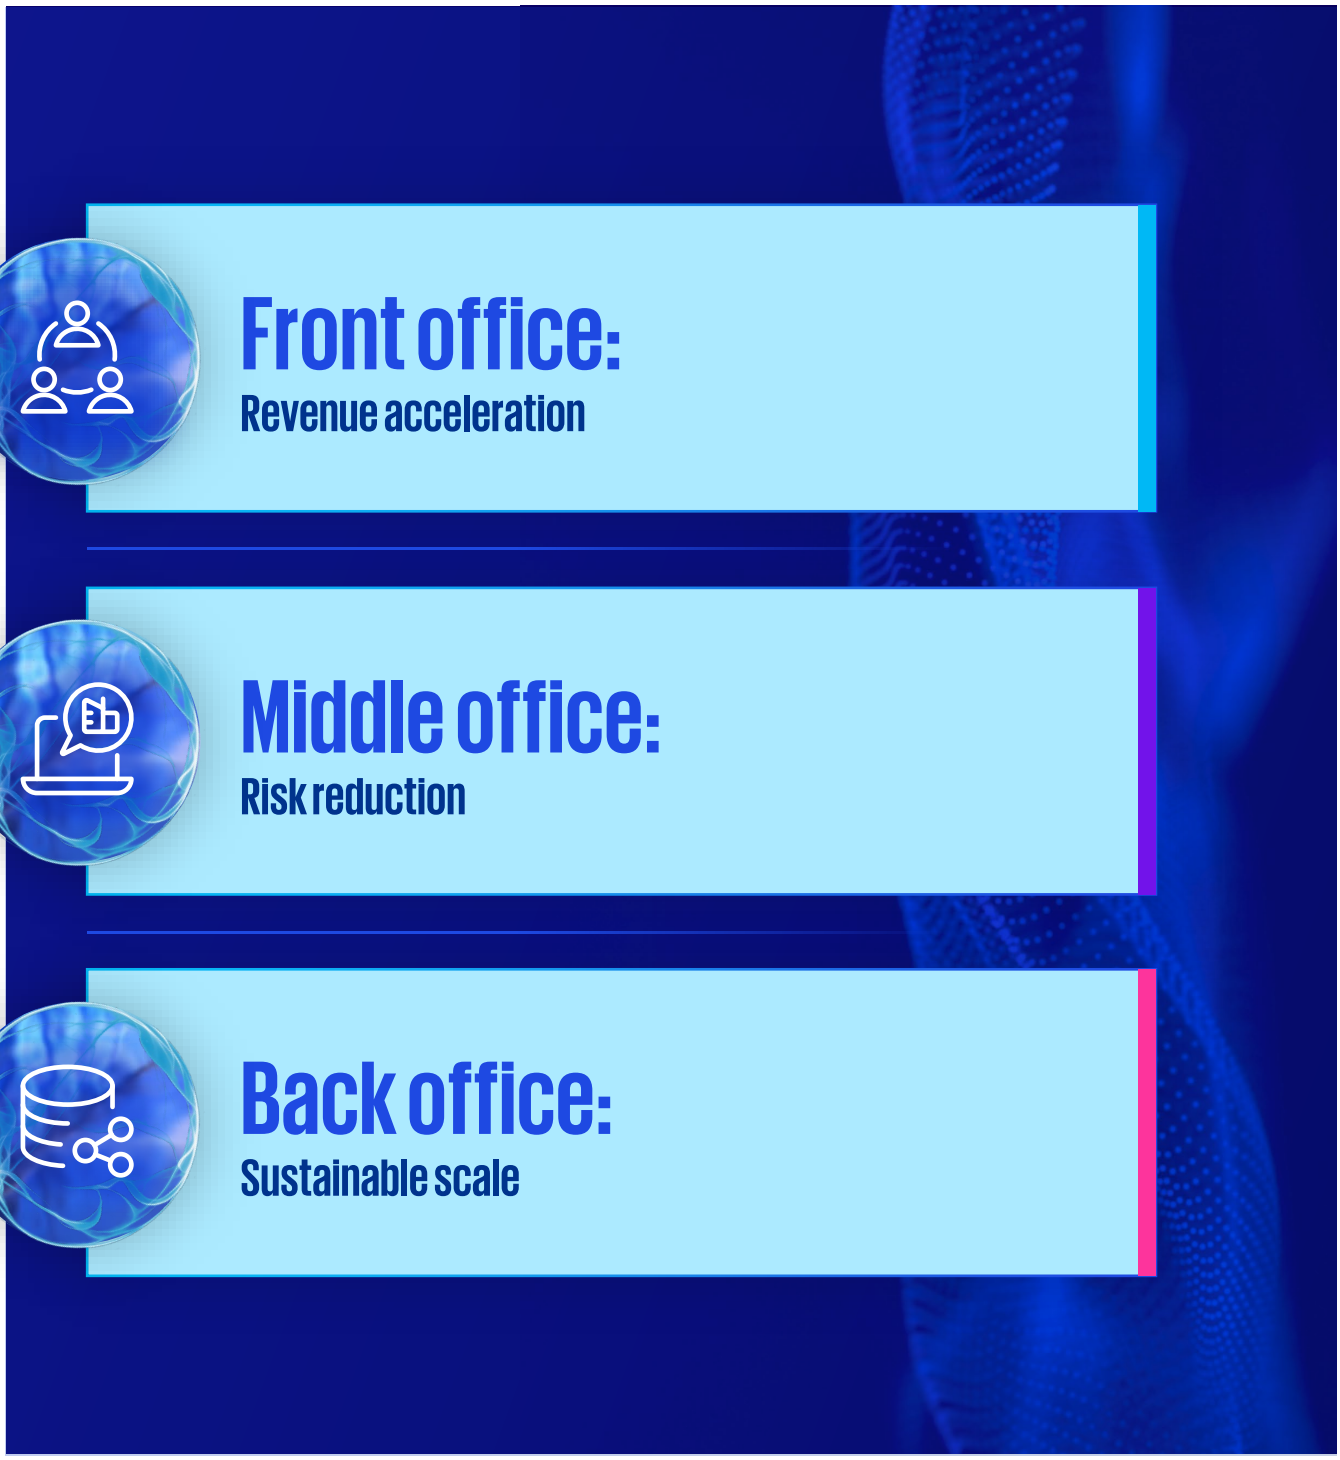

Front office:
Revenue acceleration

Front-office teams must continually connect market intelligence, portfolio information, research, and client insights to identify opportunities and make informed decisions. AI can help synthesize these inputs into timely recommendations—but only when the underlying data is trusted, contextualized, and AI-ready.

AI creates value by connecting information across the investment lifecycle and turning fragmented data into actionable insight. By delivering data as a product, the front office can deliver value through an agentic experience.

For example, a large investment manager transformed how nonstandard client inquiries and reporting needs are answered. What once required weeks of manual research across disparate data sources can now be delivered through a self-service interface that generates answers to client inquiries in minutes.

In one example, AI reduced salespeople's prep time for each meeting from 45 minutes to 7–8 minutes, which is even more impressive when it scales. For a 100-person sales team, that equates to nearly 2.5 hours saved per person, per day.[1,2]

It turns regular days into more productive days; **a more productive sales team can conduct an additional 200 client meetings daily,** allowing them to engage a wider network of financial advisors, deepen existing relationships, and significantly shorten the sales cycle across the board.[2]

[1] Based on the average salesperson having four meetings per day
[2] KPMG LLP, proprietary data, July 30, 2026

Middle office:
Reduction of risk premium

Middle-office breaks cause trade and process failures, directly impacting costs, risks, and net asset value.

AI helps by cutting through the operational noise of alerts and false positives that teams must sift through. It also reduces latency and friction across the trade lifecycle by automating trade confirmation and validation and accelerating presettlement processes.

A wealth management firm identified significant potential savings through an enterprise AI and automation strategy designed to streamline operations, reduce inefficiencies, and improve execution across the value chain.

AI can help operations teams identify anomalies, prioritize exceptions, and accelerate investigations, but only when it understands the business context behind the data. By embedding semantics and ontological relationships into the data foundation, firms can improve the accuracy of AI-driven decisions and reduce operational risk.

Back office:
Sustainable scale

Fragmented data across systems force endless back-office manual reconciliations.

Due to poor validation and governance upstream, the back office operates as a reactive hub, constantly resolving errors rather than executing workflows. With an AI-ready data foundation, the back office can scale operations more effectively across the enterprise. The goal is standardization, enabling firms to handle growing operational complexity and transaction volumes with greater consistency and efficiency.

AI's role in the back office is to significantly reduce manual reconciliation and exception handling. By automating routine operational activities, AI enables operations teams to shift their focus from transaction processing to oversight, analysis, and continuous process improvement, without compromising quality.

Escalation-based models, where AI handles over 80 percent of tasks autonomously and humans only review exceptions, delivered a **median productivity gain of 71 percent.** This is more than double the 30 percent gain seen in models requiring human approval for every output.[3]

71 percent gain in the back office means moving from manually checking thousands of items to investigating a few dozen incidents. This frees up the back office to focus on high-value initiatives such as identifying the root causes of errors, proactively resolving operational issues, and designing more resilient workflows. The data just needs to be trusted, normalized, and machine-readable, which AI-ready data is.

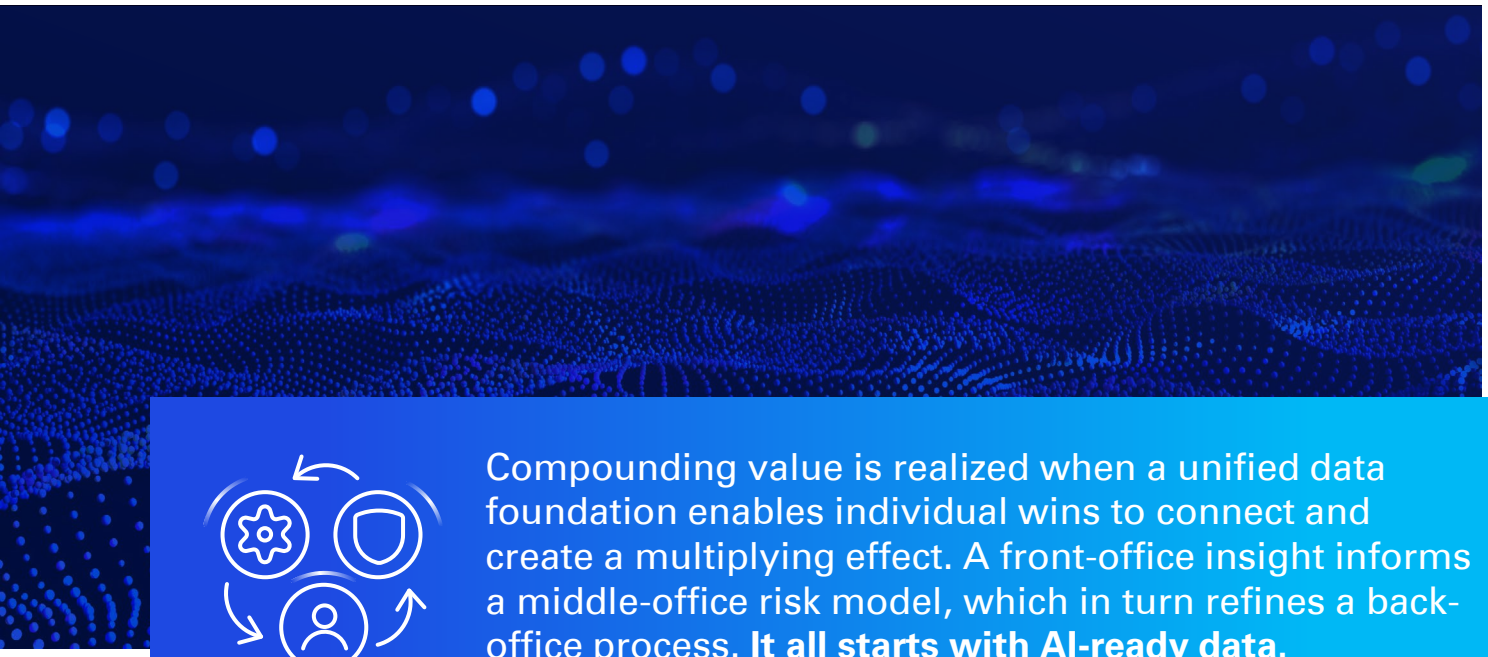

Compounding value is realized when a unified data foundation enables individual wins to connect and create a multiplying effect. A front-office insight informs a middle-office risk model, which in turn refines a back-office process. **It all starts with AI-ready data.**

[3] Elisa Pereira, Alvin Wang Graylin, and Erik Brynjolfsson, (2026). *The Enterprise AI Playbook: Lessons from 51 Successful Deployments.* Stanford Digital Economy Lab, Stanford University, p. 11, 30, April 2026.

Build the operating model for an AI-ready enterprise

Successfully scaling AI requires governance frameworks that establish trust through clear standards, policies, accountabilities, and controls. Governance defines the destination and the rules of the road, ensuring data remains trusted, secure, and fit for purpose across the enterprise.

However, governance alone is no longer enough. As organizations move toward agentic AI, the distinction between data governance and AI governance begins to disappear. Data, AI models, and AI agents increasingly operate as part of a unified ecosystem, requiring firms to govern not only data quality and lineage, but also how AI systems retrieve, reason over, and act upon enterprise information.

This shift requires a federated operating model where business, technology, and data product teams collaborate around reusable data products.

Federated governance and data product operating model

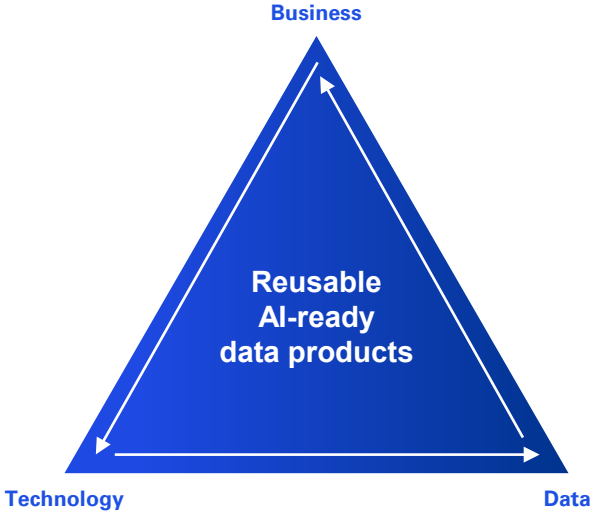

Business, technology, and data product teams collaborate to create trusted, reusable, AI-ready data products.

From traditional governance to AI-ready operations

Most asset managers have established foundational governance capabilities, but relatively few have evolved their operating models to fully support AI-ready data and agentic workflows.

1

Traditional governance

Rules, ownership and controls are defined, but execution is often manual or function-specific.

2

Data product governance

Governance begins to travel with reusable data products and shared business context.

3

AI-ready operations

Quality, lineage, access and observability are embedded into AI-enabled workflows.

4

Autonomous operations

Agents operate with context, controls, auditability and exception-based oversight.

Key takeaway

Most asset managers today operate within Stages 1 and 2. Ultimately, the goal is not simply to deploy AI; it is to evolve to Stage 3 and progress to Stage 4, creating an operating model that enables probabilistic AI systems to operate safely, reliably, and repeatedly within deterministic business processes.

Your data, your differentiator

The journey from fragmented, human-centric data to a unified, AI-ready data foundation requires more than technology modernization. It demands a fundamental shift in strategy, architecture, governance, and operating models.

Five steps to building an AI-ready data foundation

1. Determine your golden source
2. Standardize data definitions and semantic layers
3. Build AI-ready data products and pipelines that flow throughout the value chain
4. Implement robust governance, lineage, and security
5. Design an agentic workflow factory

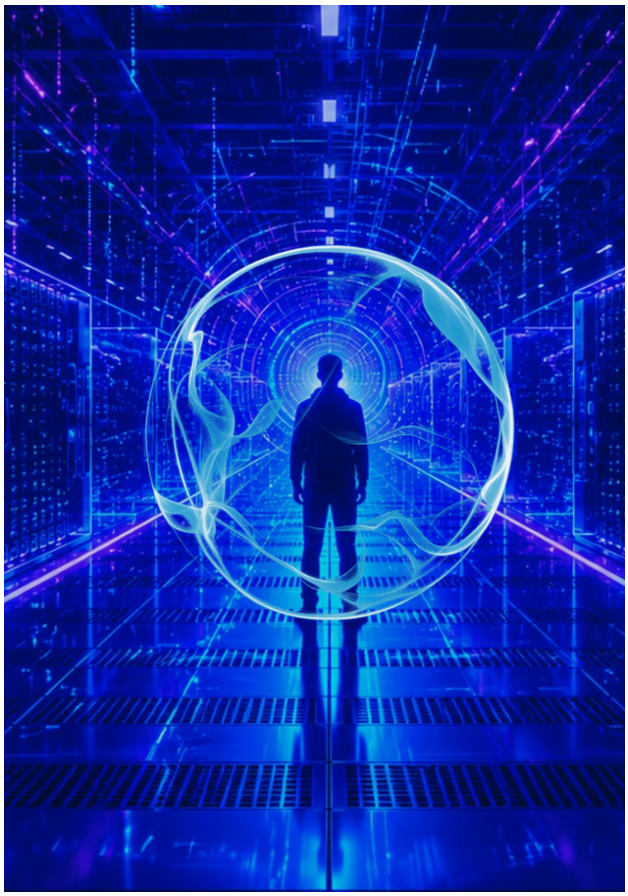

Together, these steps enable portfolio managers, operations teams, risk professionals, compliance functions, and client-facing teams to work from the same trusted foundation. The result is not simply greater efficiency—it is an enterprise that can accelerate revenue opportunities, reduce operational risk, achieve sustainable scale, and create compounding value across the investment lifecycle.

In the end, your data is your ultimate differentiator.

The firms that succeed will not be those with the most AI tools, but those with the strongest foundation to support them. By investing in AI-ready data today, asset managers can create a durable competitive advantage and position themselves to capitalize on the next generation of AI-driven innovation.

How KPMG can help

Navigating the digital transformation of asset management requires an advisor with deep industry knowledge, technical specialization, and a clear understanding of the regulatory landscape. KPMG LLP is well positioned to collaborate with your firm, guiding you through every stage of your AI-ready data journey, from strategy to execution.

Our multidisciplinary teams have extensive experience in AI, data science, and the asset management industry, and have helped our clients...

Develop a strategic AI roadmap

We work with you to define a clear vision for AI, identifying high-impact use cases across the front, middle, and back office that align with your business objectives. We help you build a compelling business case internally and a practical journey map for AI-ready data implementation.

Establish an AI-ready data foundation

Recognizing that AI-ready data is the bedrock of AI's value, we help organizations establish the foundations required to support AI at scale. Our services include designing semantic layers and ontologies, enabling federated access to trusted data, establishing governance frameworks, and improving data quality, context, and accessibility for AI-driven workflows.

Implement and scale AI use cases

We help design, build, and deploy AI solutions tailored to asset management. From advisor and client engagement in the front office to exception management in the middle office and operational efficiency in the back office, we help organizations turn AI ambitions into measurable business outcomes.

Navigate risk and regulation

We help you build trust in your AI systems. Our subject matter professionals provide guidance on establishing strong AI governance, risk management, and control frameworks to help ensure your use of AI is responsible, ethical, and compliant with evolving regulations. We can help you manage model risks, address potential biases, and ensure the security of your AI applications.

Deploy agentic workflows

We help you transition from legacy systems to deploying AI agents that act as digital coworkers. We efficiently integrate autonomous AI workflows into your operations to compound value and drive unprecedented efficiency.

Enable an AI-efficient workforce

A successful AI transformation is as much about people as it is about technology. We provide support for change management and help upskill your workforce, helping ensure your teams have the necessary capabilities to thrive in an AI-augmented environment.

The journey to AI-ready data is essential for leveraging AI and automation throughout the asset management value chain.

KPMG is there every step of the way, providing strategic guidance, technical capabilities, and industry experience to help your firm unlock the full potential of AI, driving sustainable growth and creating an enduring competitive advantage.

Authors

Agnel Kagoo
US Consulting Sector Leader, Asset Management and Capital Markets
KPMG LLP
akagoo@kpmg.com

K Agrawal
US Offering Leader, Asset Management Digital Distribution,
KPMG LLP
kagrawal@kpmg.com

Murali Sreedhar
Principal, Capital Markets, Financial Services Operations
KPMG LLP
msreedhar@kpmg.com

Matteo Colombo
Global Consulting Cloud, Data and AI Leader
KPMG LLP
matteocolombo@kpmg.com

Daisy Tung
Partner, Customer and Operations Financial Services,
KPMG LLP
daisytung@kpmg.com

James Gardner
Principal, Capital Markets and Asset Management – Technology, Data and AI
KPMG LLP
jgardner@kpmg.com

Jeremy Edwards
Director, Finance Data Strategy and Operations,
KPMG LLP
jeremyedwards@kpmg.com

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Learn about us: **kpmg.com**