FUNDGUARD

WHITEPAPER

Intelligence at the Core: The AI Divide in Investment Accounting

How Asset Managers, Asset Servicers and Asset Owners Should Implement AI

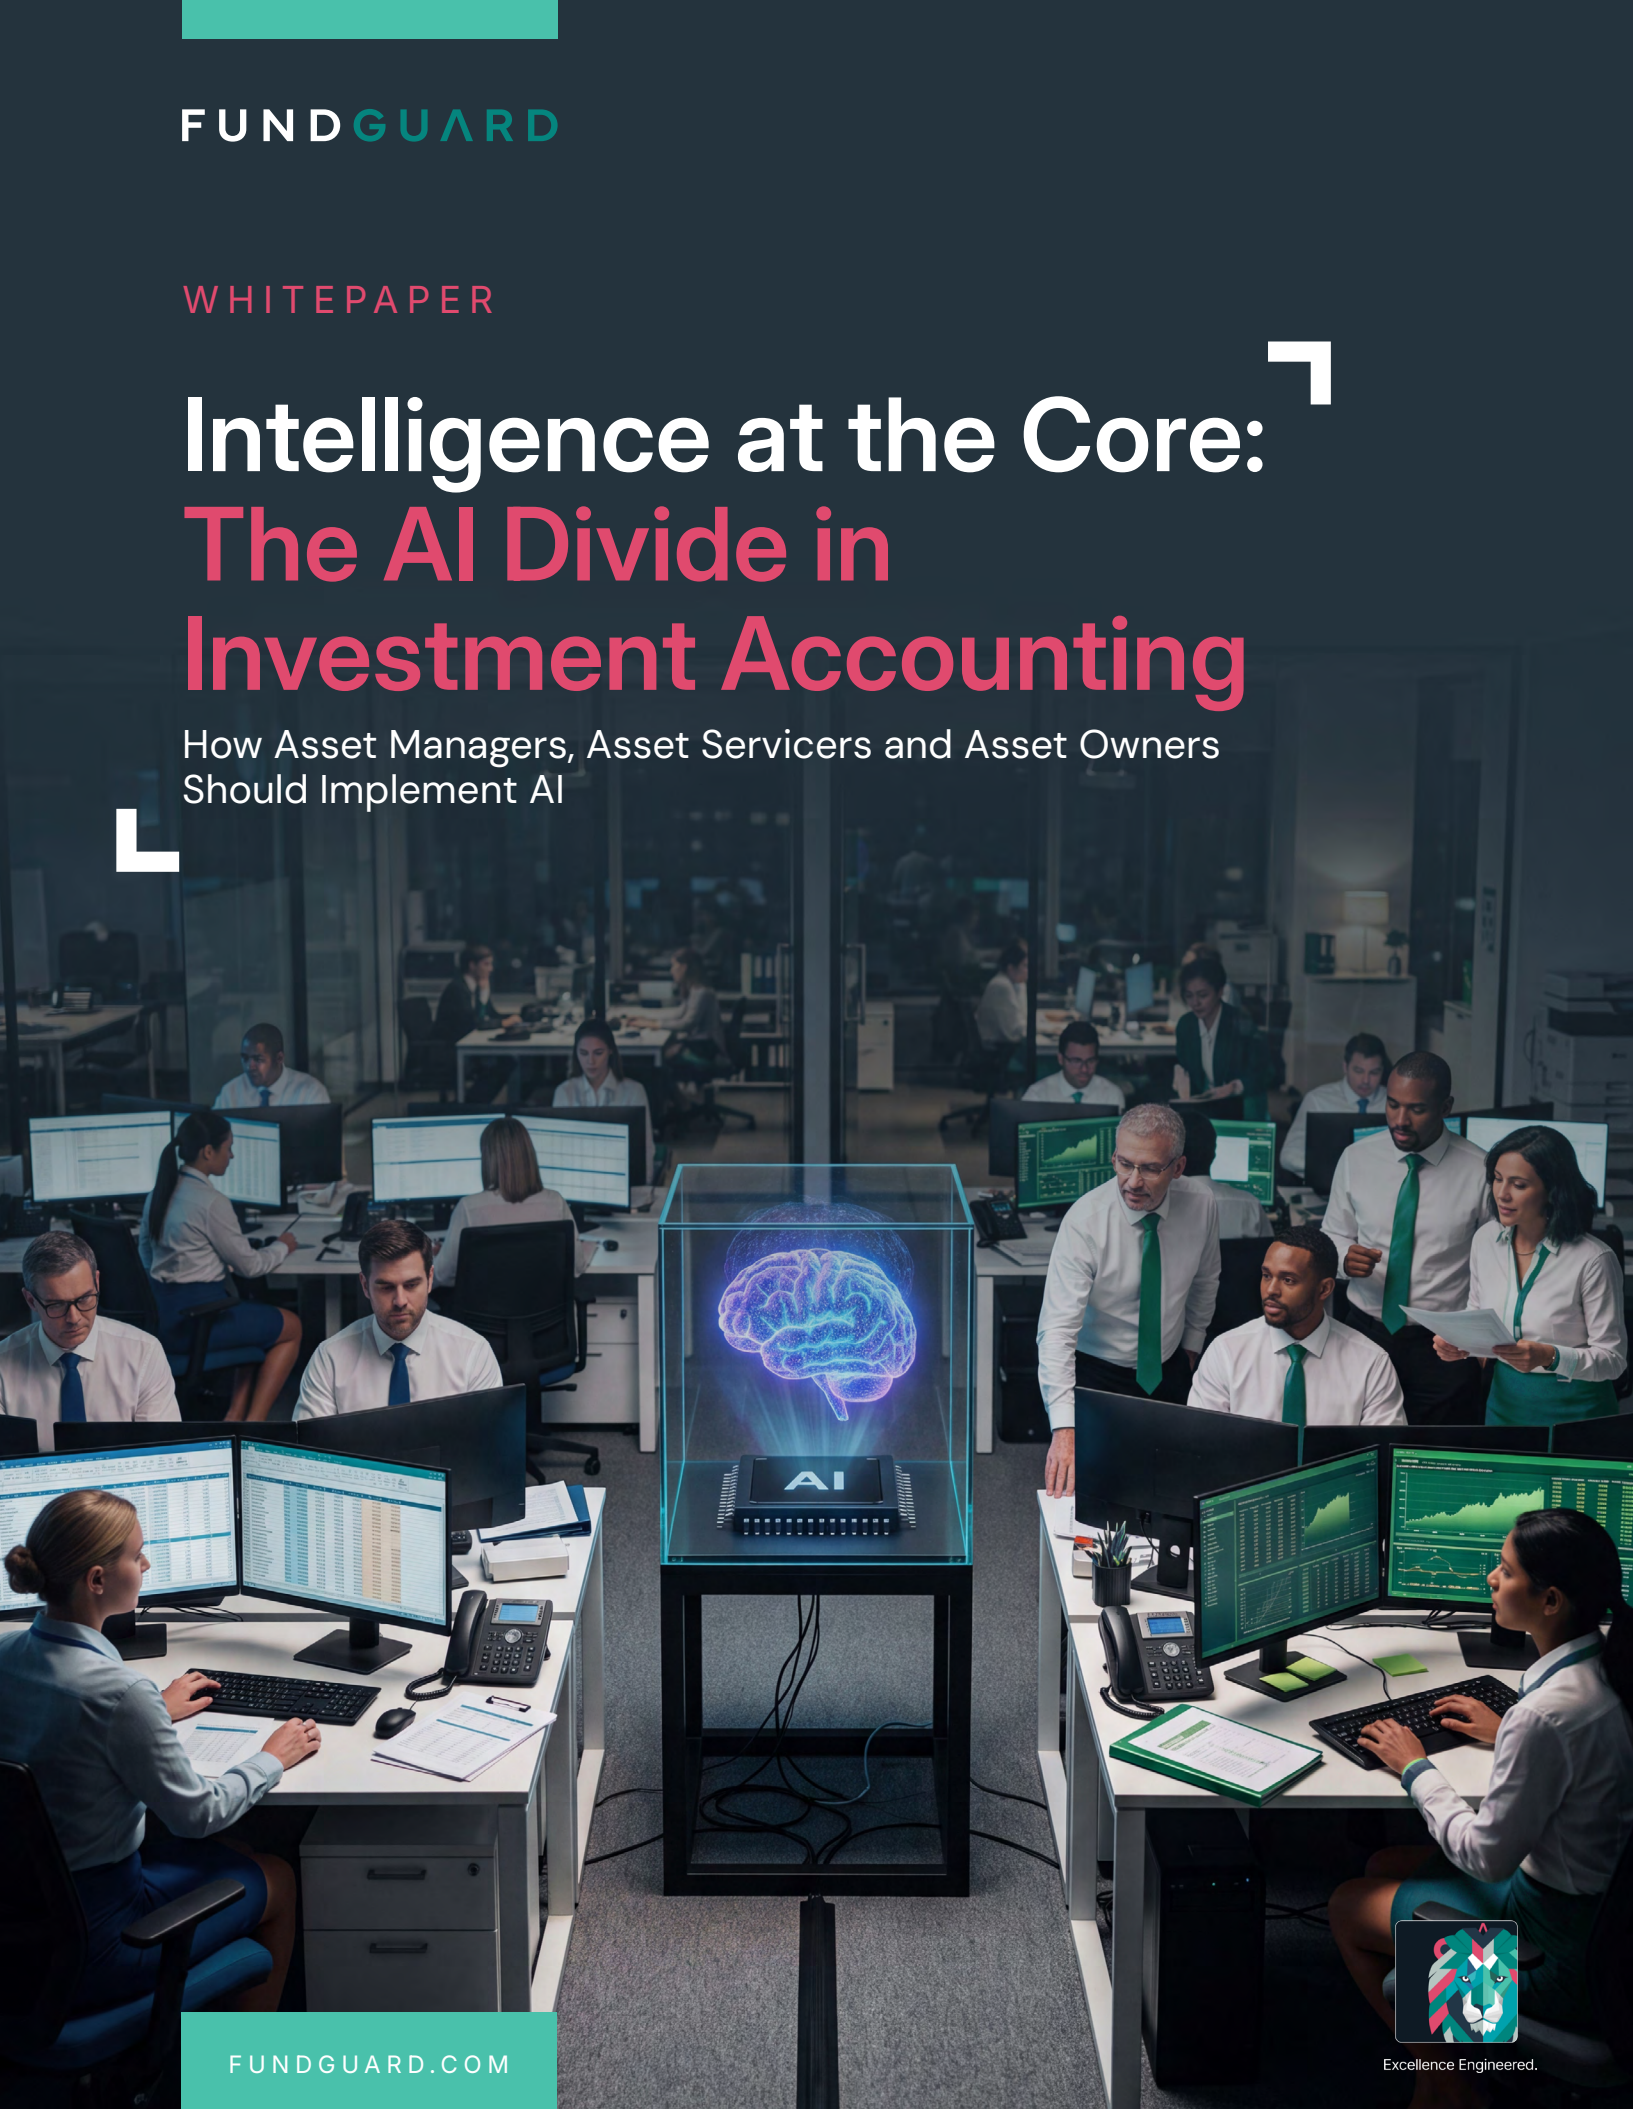

FUNDGUARD.COM

Excellence Engineered.

Contents

Introduction:
The AI Ambition-Integration Gap

Agentic AI is changing the economics of investment operations. Boston Consulting Group's Global Asset Management Report 2026 projects that agentic workflows could automate 70–80% of standard trading and execution activity, reduce investment operations costs by approximately 40% and materially expand operating capacity.

Capturing that value requires redesigning the operating model, not simply automating existing processes.

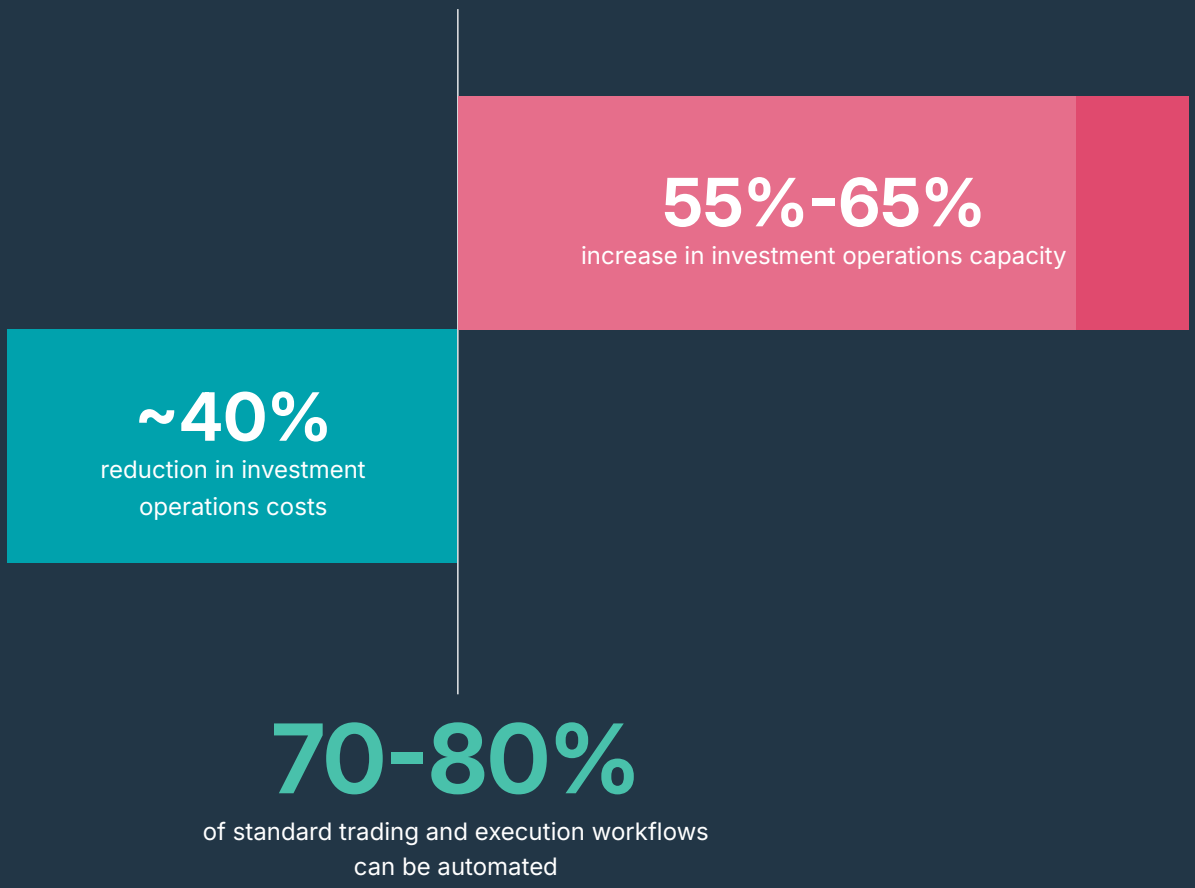

Here's how to bridge the gap between AI ambition and AI integration.

1

Legacy architecture caps AI

Batch processing, asset silos and overwrite-based data limit what any model can do.

2

Engine-level AI changes the operating model

Controls run before exceptions compound and exceptions get filtered, NAV validates continuously.

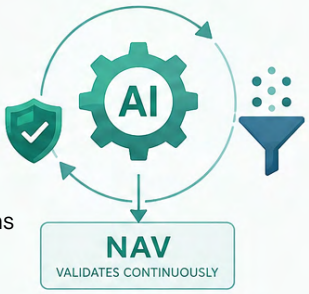

3

Architecture is the prerequisite for AI

AI only works in accounting when it is paired with cloud-native, bitemporal, multi-book architecture.

4

Your vendor must successfully demonstrate AI

Find seven questions in this whitepaper to truly evaluate vendor AI-readiness.

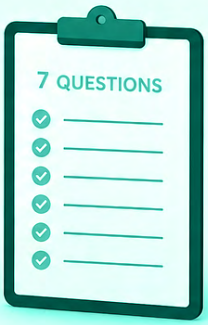

The gap between AI ambition and AI integration will shape investment operations over the next several years. Firms that get the architecture wrong will continue paying for exception queues, delayed NAVs and integration projects that never fully resolve the underlying fragmentation.

Boston Consulting Group reports that asset managers still trail banks and fintech firms in scaling AI across core processes. Most remain focused on pilots and incremental productivity gains, even though the larger opportunity lies in operating-model redesign.

Achieving that value requires unified data, embedded governance and an architecture capable of supporting agentic workflows across accounting, reconciliation, corporate actions and unstructured data. AI cannot deliver that transformation when it remains a feature layer above the same batch processing, asset silos and fragmented books of record.

FUNDGUARD

Break Free of Impressive Pilots That Lead to Disappointing Production Results

The gap between AI ambition and production results is fundamentally an architecture problem. Legacy investment accounting systems were designed for a world in which the answer arrived the following morning. AI operates differently. It requires continuous data, event-driven workflows, versioned state and a unified dataset across books and asset classes.

FundGuard provides the real-time accounting foundation for AI-enabled investment operations.

AI Built on a Real-Time Accounting System of Record

Embedded intelligence, trusted accounting data and governed access for client AI and agentic workflows

Modern Tech at Scale

Cloud-native, event-driven architecture designed to support real-time investment operations and accounting across institutional portfolios and operating models.

Unified Accounting Data

A single, bitemporal system of record spanning books, asset classes and accounting views, without relying on fragmented subledgers.

Proven in Production

Deployed with asset managers and asset servicers to support investment accounting, fund accounting, oversight and operational workflows at scale.

Deep Industry Expertise

Purpose-built for the controls, data, workflows and regulatory requirements of investment operations and accounting.

The AI-Readiness Framework

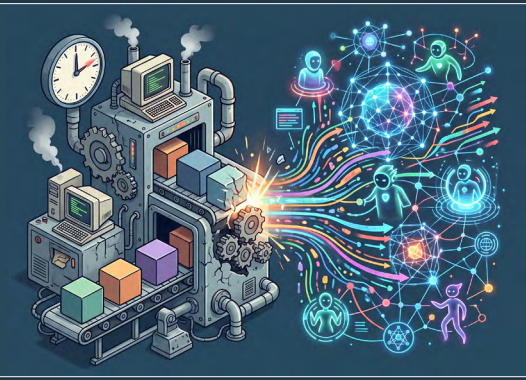

Legacy batch architecture

limits enterprise-scale AI by delaying data, fragmenting accounting state and restricting intelligence to after-the-fact analysis.

Integration

AI embedded within the accounting platform can support proactive controls, intelligent exception management, continuous validation and governed automation.

Prerequisite

Cloud-native, event-driven, bitemporal and multi-book architecture provides the trusted foundation AI requires.

Outcomes

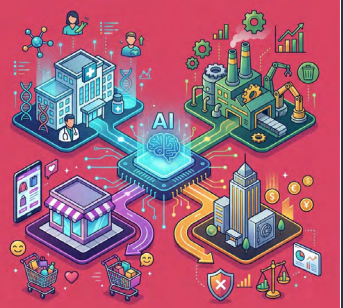

Asset managers, asset servicers and asset owners should evaluate AI against the specific operating outcomes they need to achieve.

Shared Vocabulary

Buyers need a consistent way to distinguish production-ready AI from capabilities layered above legacy systems.

Why Legacy Batch Architecture Cannot Be AI-Enabled

The investment accounting systems running most of the global industry today were designed in the 1990s and early 2000s, with some going as far as the 80s. Their defining characteristic is batch processing: data is collected throughout the day, processed overnight and reported the following morning. This was an entirely reasonable design for its era.

It is now the single biggest barrier to AI adoption in fund operations.

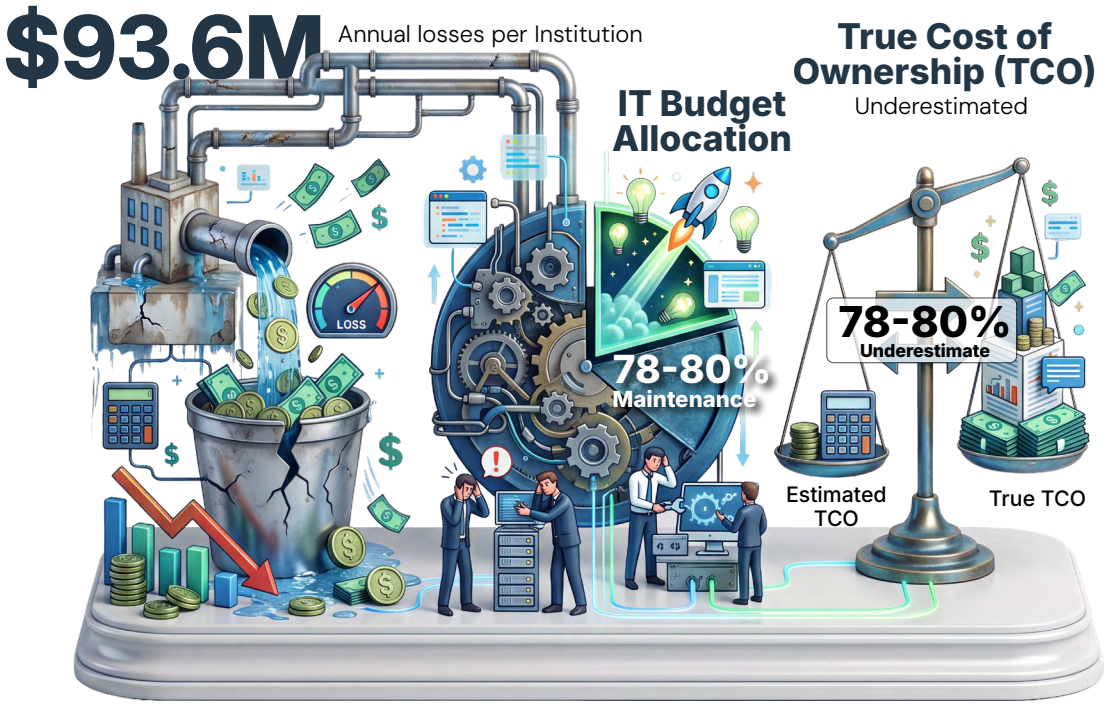

Legacy infrastructure also consumes the resources firms need for modernization. Research has estimated annual losses of $93.6 million per institution across operational inefficiency, processing friction, human error, rework, cybersecurity and compliance exposure. Separate analysis suggests that maintaining legacy systems can consume 70–78% of financial institutions' IT budgets.

That burden extends beyond cost. The more time and budget devoted to maintaining fragmented batch systems, the less capacity remains to modernize the accounting foundation on which AI depends.

Asset-class silos compound the problem. Equities may sit in one ledger, fixed income in another, derivatives in a third and private assets in spreadsheets or specialist subledgers. Corrections, cancellations and restatements fragment the historical record further. In these environments, AI can evaluate outcomes only after processes have completed. It cannot participate while accounting is taking place.

> Almost every interesting AI use case in investment accounting is inherently time-sensitive.

- Anomaly detection loses most of its value if it fires the morning after the NAV has already been struck
- Intelligent reconciliation cannot replace exception queues if the breaks are only visible after overnight processing

AI layered over batch processing remains retrospective. Fragmented books of record create fragmented data, and fragmented data turns every AI initiative into a custom integration project.

- A model trained on reconciliation breaks from one system cannot reason about exceptions in another
- Every additional asset class, every additional jurisdiction, every additional book of record multiplies the integration burden

The FundGuard MCP, Agents and A2A Ecosystem

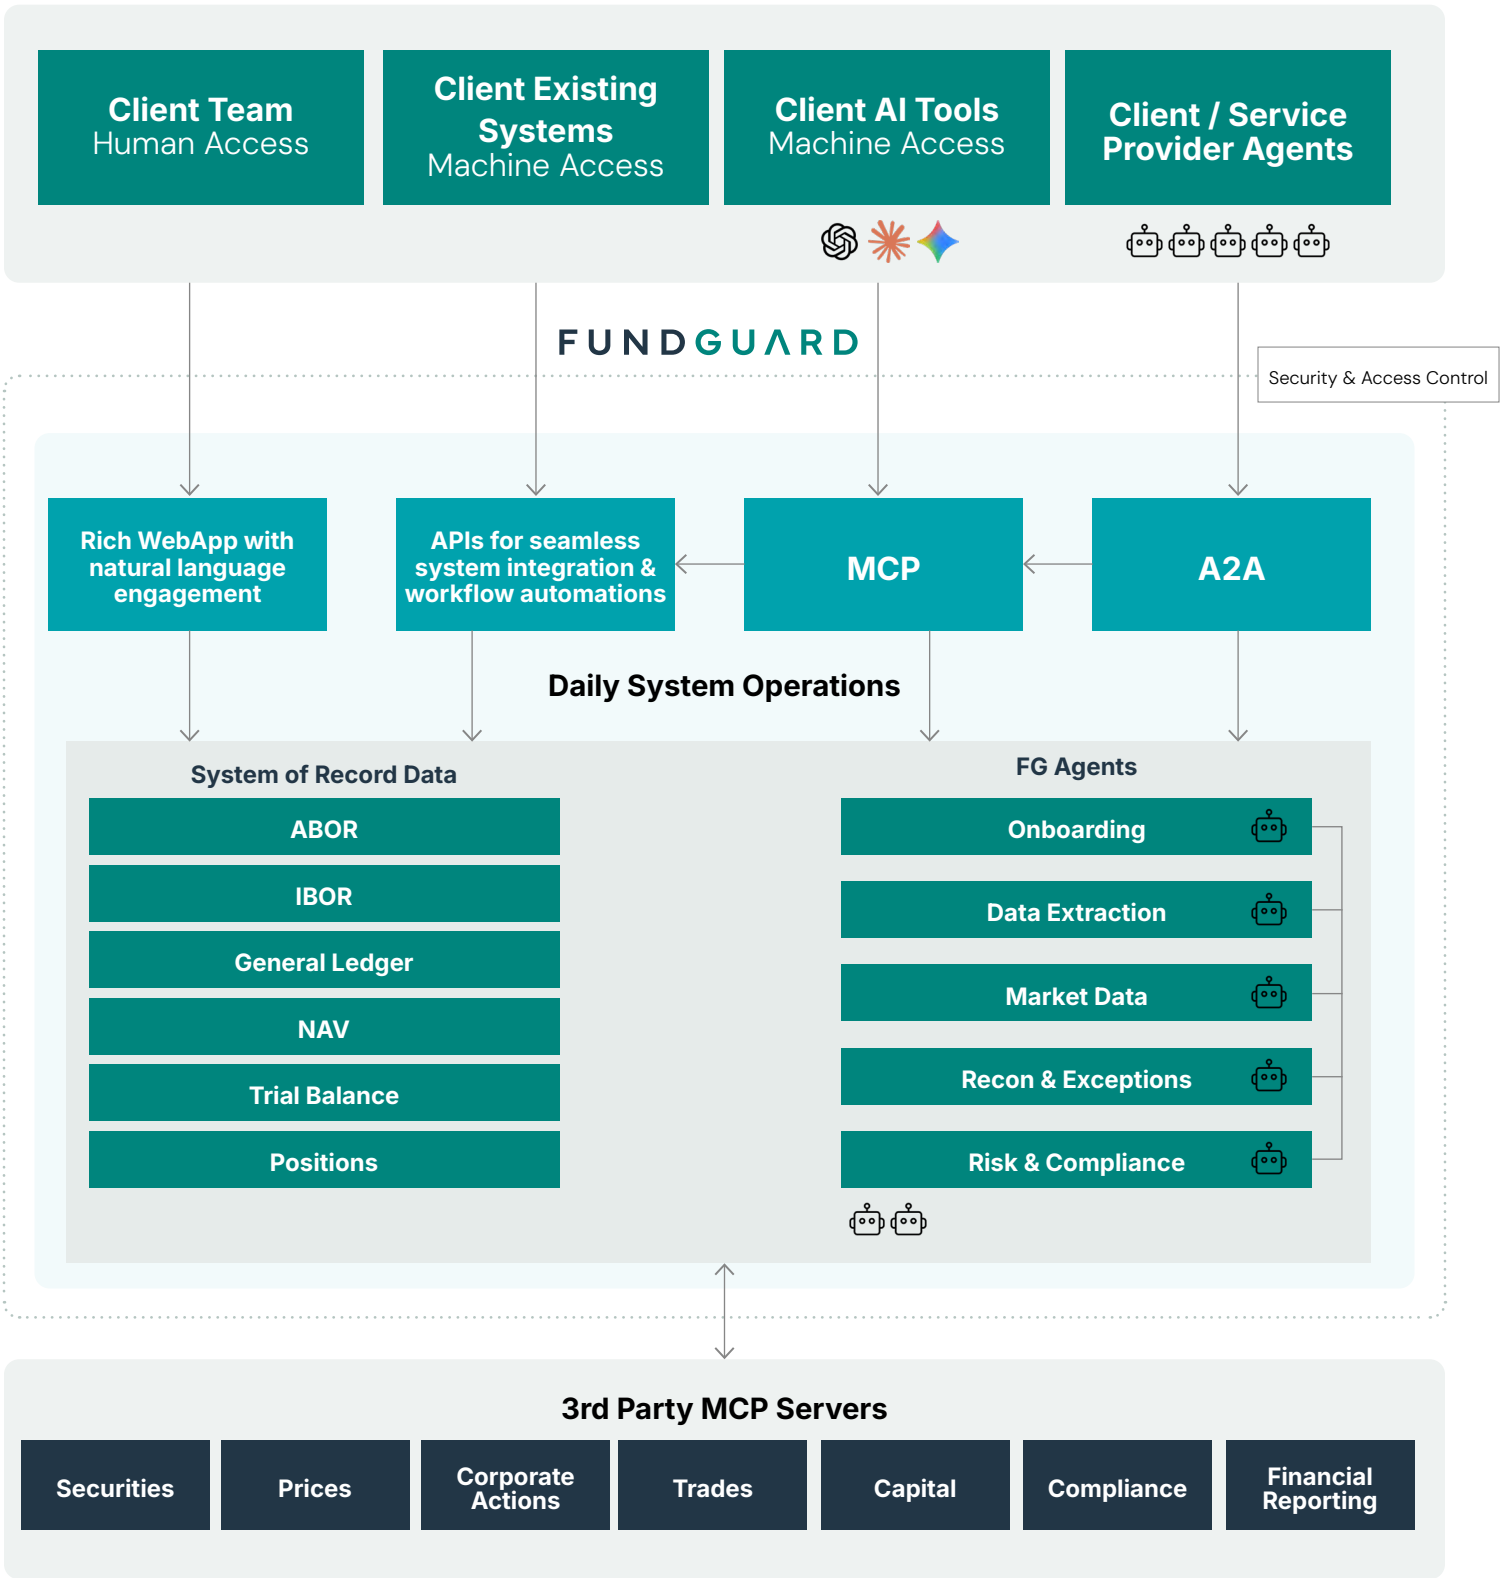

Legacy systems fragment data across monolithic applications and disconnected books of record. That fragmentation restricts the ability to develop AI models that can operate across complete, current accounting data. In one analysis of financial institutions, only 32% had successfully integrated AI into their core systems.

Legacy platforms also tend to overwrite state. Modern AI requires the opposite. To learn, explain and support audit, every correction and restatement must remain visible and traceable. When point-in-time balances are overwritten, the system loses both the historical context a model needs and the evidence regulators, auditors and clients may need to reconstruct a decision.

Legacy systems will not disappear overnight. Replacement programs are complex, and many firms will continue to patch, wrap and extend existing environments. The strategic issue is therefore not whether legacy systems can continue running. It is whether they can provide the authoritative, machine-readable accounting state required for AI-enabled operations.

Firms broadly face three paths:

1. Keep the incumbent system and layer AI or agents above it
2. Build a new accounting data foundation
3. Adopt an AI-native platform designed around real-time, governed and machine-readable data

Layering agents above the incumbent may accelerate access to information, but it does not resolve fragmented state, delayed processing or incomplete audit history. Building a new foundation internally requires firms to recreate accounting logic, permissions, controls and integrations. An AI-native system of record addresses the problem at the architectural level.

The defensible system of record must be more than a data layer with an API on top.

The defensible system of record has to be more than a data layer with an API on top.

AI's Role
Within the Accounting Engine

When AI is embedded into the accounting engine rather than layered on top of it, the operating model changes in four specific ways:

- Controls become proactive rather than reactive
- Exception management becomes intelligent rather than exhaustive
- NAV calculation becomes continuous rather than cyclical
- Operational alpha becomes a measurable outcome rather than a marketing phrase

1. Proactive Controls Instead of After-the-Fact Variance Reports

Traditional accounting controls are narrow and retrospective by design. They examine defined data points and depend heavily on manual investigation when an exception occurs.

AI broadens the control environment by comparing information across peer funds, securities, historical patterns, market conditions and related attributes. This allows the system to identify inconsistencies that deterministic checks alone may not detect.

Because these controls operate against streaming event data rather than batch extracts, they can surface issues before those issues propagate into downstream books and processes.

2. Intelligent Exception Management

Every legacy accounting environment generates a long tail of reconciliation breaks, price-tolerance failures, corporate-action mismatches and cash-timing differences. When traditional systems flag all of them and accounting teams must investigate each one, the cost of exception management grows with operating complexity.

AI changes the economics of this workflow in three ways.

Filters: The model learns which exceptions have historically been benign and which patterns are more likely to indicate a genuine issue.

Explains: When an exception requires review, the model surfaces the most likely cause, such as a late corporate action, missed FX rate or trade booked to the wrong strategy. The accountant begins with context rather than starting the investigation from scratch..

Resolves: For high-confidence, low-risk items, the system can resolve or approve the exception within defined thresholds while retaining a complete audit trail and escalating items that still require human judgment.

Stage 0: Intake

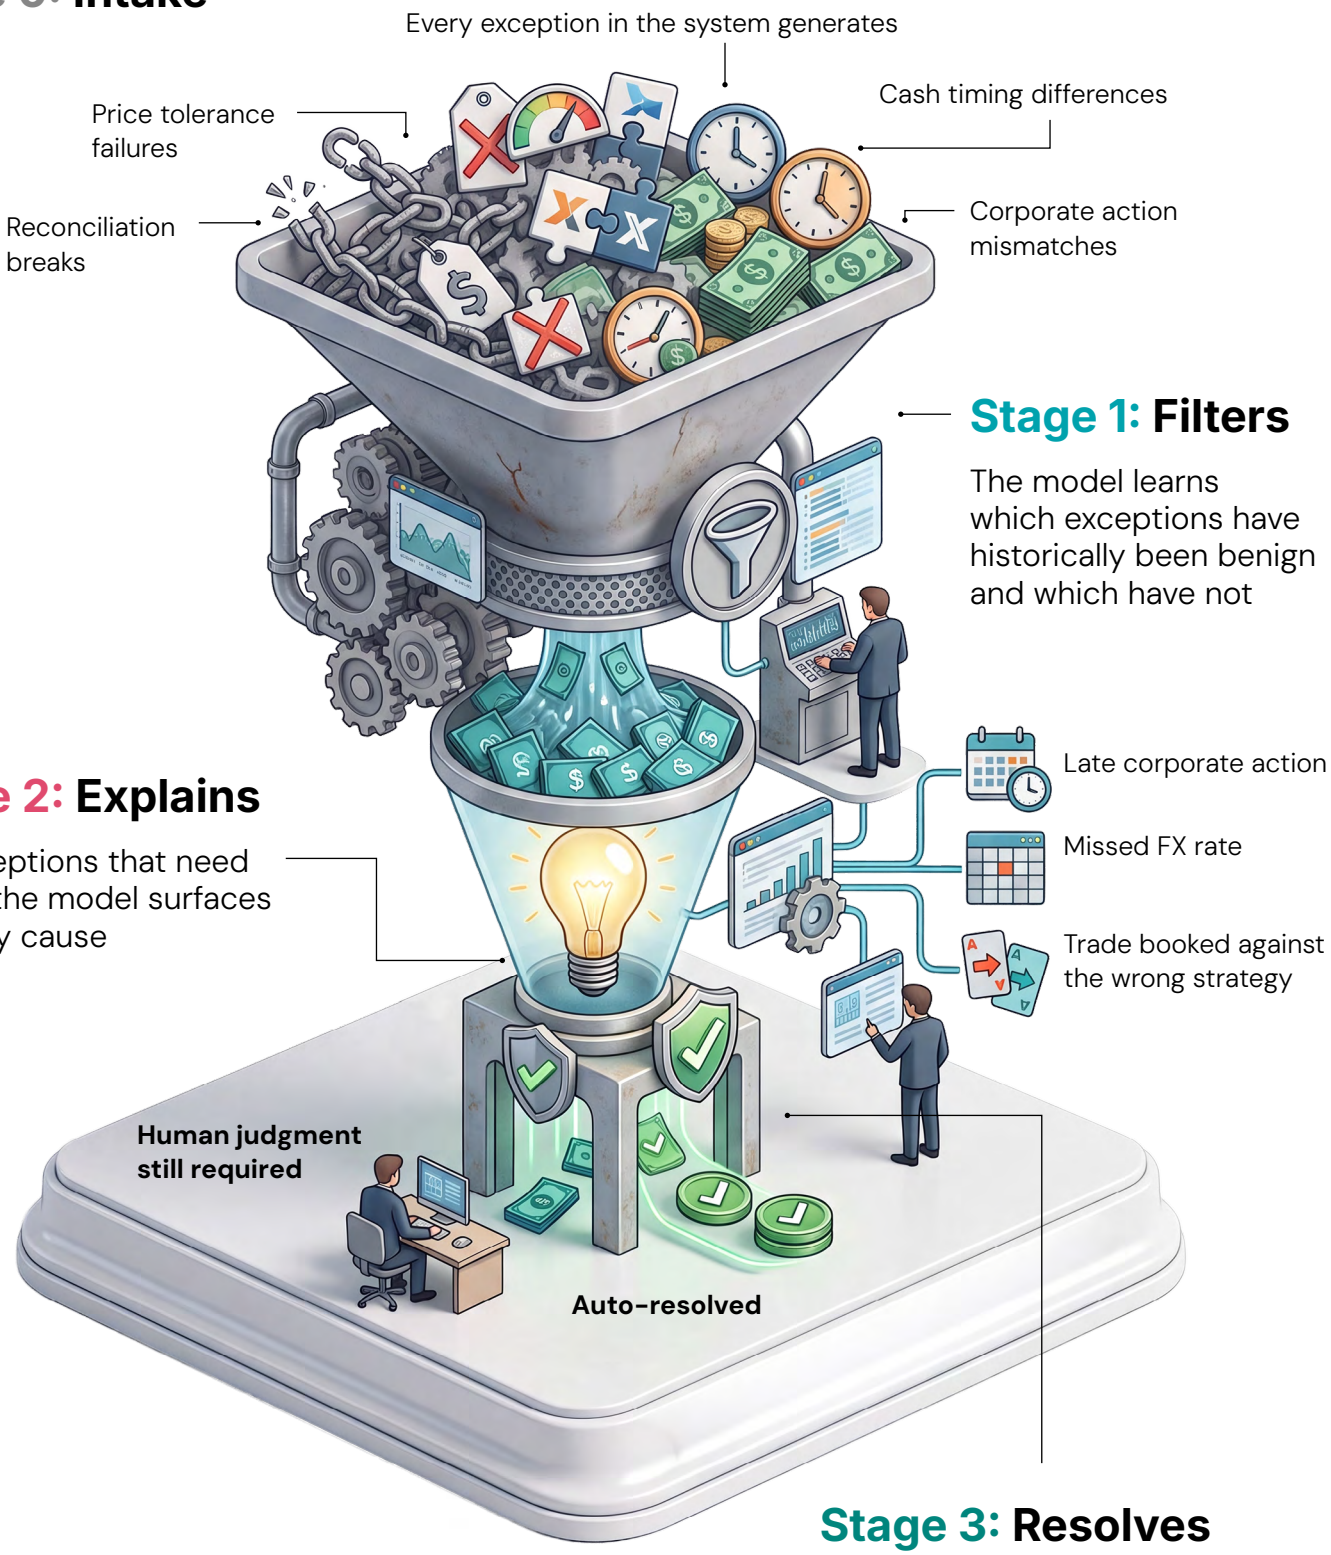

Stage 1: Filters

The model learns which exceptions have historically been benign and which have not

Stage 2: Explains

For exceptions that need review, the model surfaces the likely cause

Stage 3: Resolves

High-confidence, low-risk items auto-approve within defined thresholds with a full audit trail

Agentic workflows extend this operating model by coordinating actions across accounting, reconciliation, corporate actions, oversight and reporting.

In public markets, agents can help reconcile NAVs across custodians, process corporate actions and separate true exceptions from operational noise. In private markets, they can support capital-call processing, document interpretation and validation of complex calculations.

The role of the accounting professional also changes. Rather than manually triaging every item, teams define controls, review judgment-based exceptions and improve the logic that governs automated resolution.

3. Continuous NAV Validation and Readiness

Most funds still operate on a lagged NAV cycle because prices, trades, corporate actions, cash and other inputs arrive at different times and in different formats. The accounting book is finalized at defined points in the process, even though the underlying portfolio state continues to change.

AI-enabled continuous validation allows firms to assess data quality, positions and exceptions as information arrives. Incoming data can be evaluated against expected ranges, historical patterns and related portfolio activity before issues affect downstream valuation.

Formal NAV approval and publication remain governed processes. The difference is that validation and readiness no longer need to wait for the overnight cycle.

For asset servicers, continuous oversight can strengthen both the service proposition and the contingency model by maintaining an always-current view alongside the primary accounting process. For asset managers and asset owners, it provides earlier visibility into issues that could affect valuation, reporting or investment decisions.

4. Operational Alpha as a Measurable Outcome

Operational alpha is the capacity a firm frees when manual, repetitive and low-judgment work is absorbed by the system.

That capacity can then be redeployed into higher-value activities such as client service, product launches, investment support, oversight and process improvement.

In practice, operational alpha should be measured through changes in exception volumes, auto-resolution rates, investigation time, onboarding effort, NAV readiness and operating capacity per fund or mandate.

Operational alpha is the capacity a firm frees up when manual, repetitive, low-judgment work is absorbed by the system

Architecture as a Prerequisite

If AI-enabled operations are the outcome, cloud-native architecture is the prerequisite.

AI performance is constrained by the architecture beneath it. Production-scale AI in investment accounting requires five capabilities:

- Event-driven processing
- Bitemporal accounting data
- Governed and explainable AI controls
- Multi-book, cross-asset intelligence
- Secure interoperability for agentic workflows

Cloud-Native, Event-Driven Processing

Legacy systems were built on batch processing, while cloud-native systems are built on events.

In an event-driven architecture, each trade, price update, cash movement and corporate action is processed as a discrete event. That event can trigger validation, enrichment, accounting and downstream workflows as soon as it enters the system rather than waiting for an end-of-day cycle.

Cloud-native design also provides the elasticity required for AI workloads. Routine intraday monitoring and large-scale validation across millions of positions place very different demands on compute and storage. A cloud-native platform can scale those resources as needed.

As firms move toward agentic operating models, interoperability becomes as important as intelligence. Standards such as Model Context Protocol and Agent-to-Agent communication can give AI systems governed access to context and allow agents to coordinate actions across workflows.

These standards do not solve fragmented accounting data. Without a trusted system of record, agents inherit the same incomplete, delayed and inconsistent state that limits traditional automation.

As firms move toward agentic operating models, interoperability becomes as important as intelligence.

Enterprise-Grade Cloud Technology

Cloud-native microservices architecture designed for real-time investment operations and accounting

Cloud Native / SaaS

Real-Time/ Bitemporal

AI Native

- Real-Time and Bitemporal
- Multi-Book from a Single Dataset
- Cloud-Native SaaS
- API First
- End-to-End Processing
- Big Data Database
- AI (anomalies, rapid onboarding)
- Test Center/Library
- High Availability

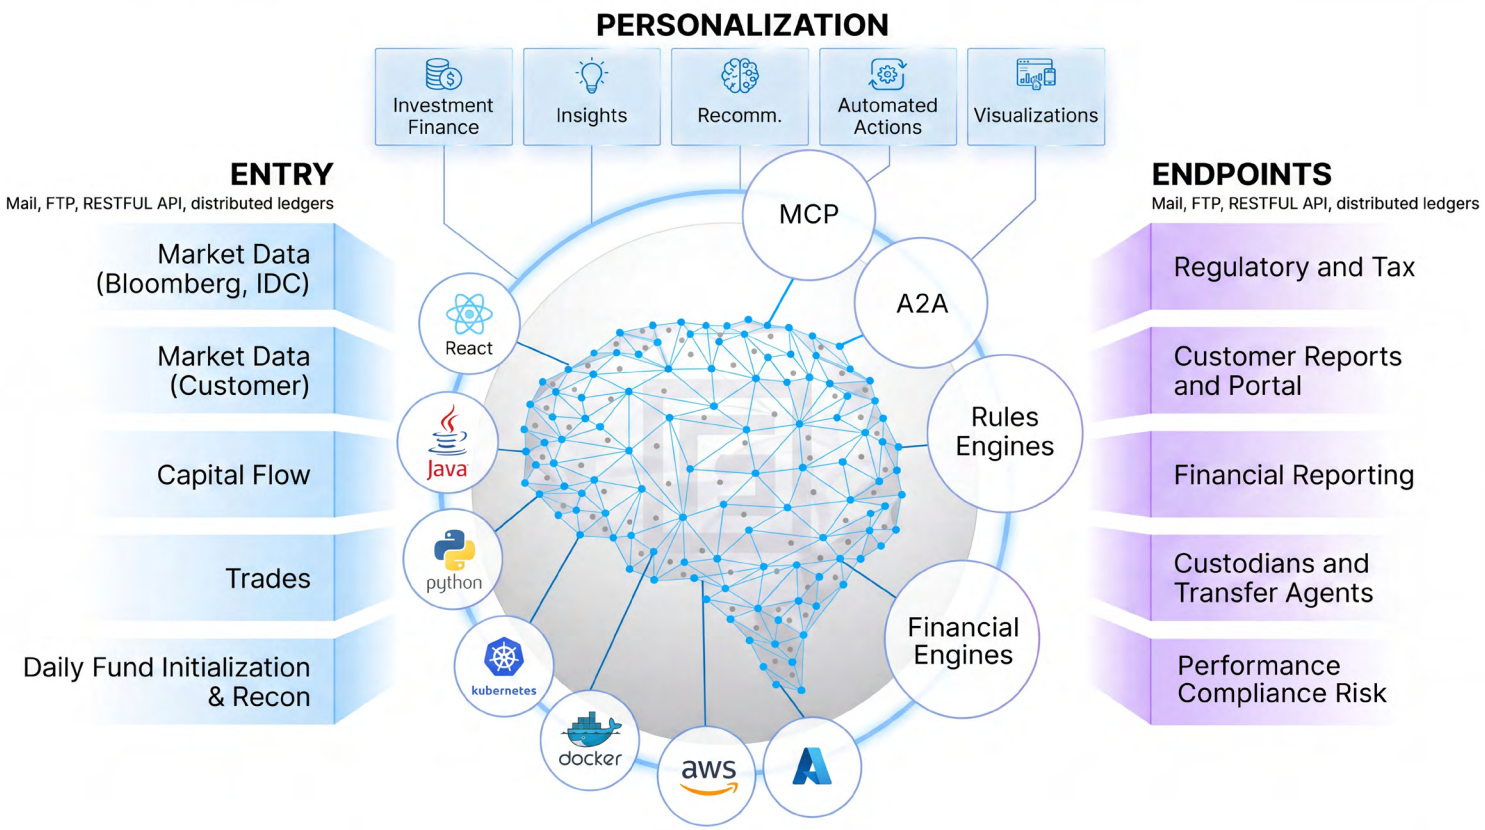

DATA SERVICES

kafka · Spark · Airflow · trino

DATA STORES

cassandra · elastic · redis · ICEBERG

ANALYTICS

databricks · snowflake · Power BI

Bitemporal Data

Bitemporal accounting is the foundation of explainable AI in accounting. Without bitemporality, the historical record is continuously overwritten.

With bitemporality, every correction and restatement remains visible and traceable.

Agents require explicit rules, permissions and process definitions to act safely. They also need to reconstruct what was known at any prior point because both their reasoning and their audit trail depend on it. A system that overwrites historical state cannot provide that context.

Bitemporality has two important consequences.

First, it allows regulators, auditors, clients and internal teams to reconstruct what was known, when it was known and how a decision was made.

Second, it gives AI a cleaner training and evaluation signal. The system can distinguish between a trade that was initially incorrect and later corrected and a trade that was correct throughout because both states are preserved. Overwrite-based systems discard that distinction.

Explainable AI requires more than a record of the final answer. It requires the full sequence of events, corrections and decisions that produced it.

A system that cannot reconstruct its own state through time cannot support that accountability.

AI You Can Trust

For investment accounting, AI cannot be evaluated solely on automation rates or model accuracy. It must operate within the control environment required for regulated financial services.

That means protecting confidentiality, preserving data integrity and ensuring availability. It also means maintaining explainability, governance and auditability while allowing AI to participate in operational workflows.

Accountability must remain with the institution, supported by defined permissions, decision thresholds, human oversight and a complete audit trail.

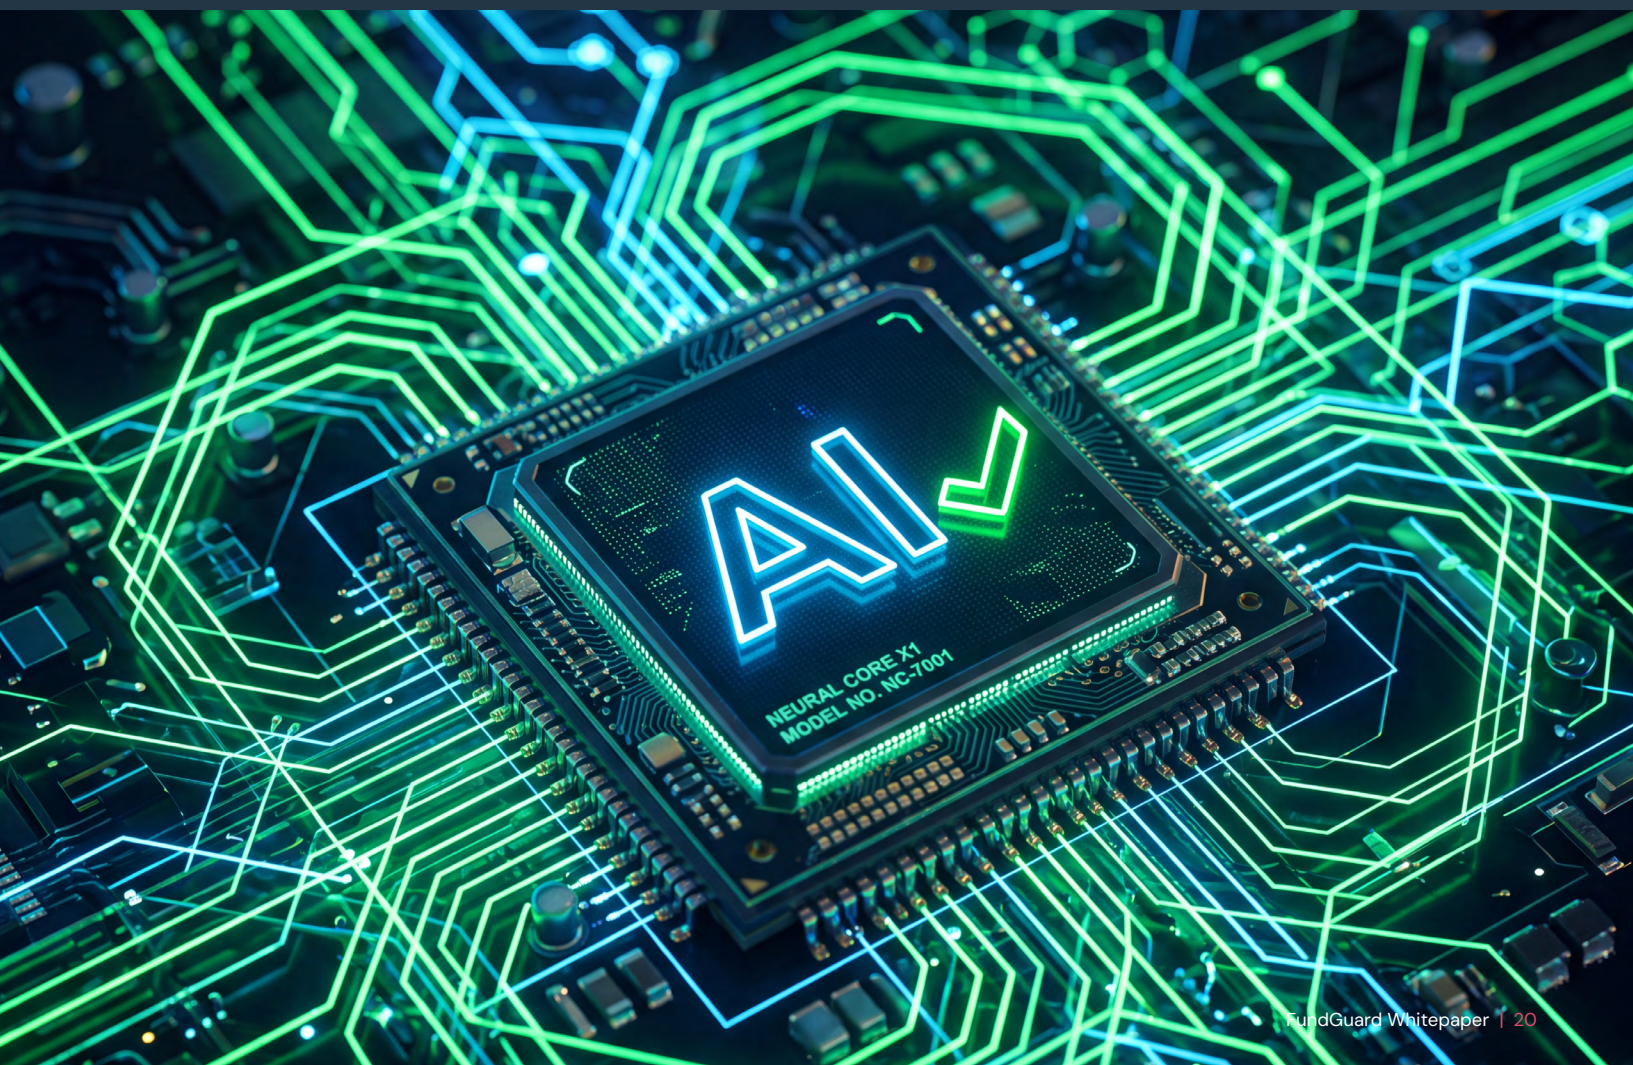

Multi-Book, Cross-Asset Intelligence

A unified multi-book architecture allows IBOR, ABOR, CBOR, PBOR and custom views to derive from the same underlying transactional dataset. Equities, fixed income, derivatives, private credit, private equity, real estate and digital assets share a common portfolio model rather than residing in disconnected asset-class systems.

This matters because the value of AI increases with the breadth and quality of the context available to it. A model evaluating listed equities alone can identify patterns within that dataset. A model that can also access related derivatives, corporate actions, issuer exposure and private holdings can identify relationships that would remain invisible within a single asset silo.

The platform should make cross-asset reasoning possible while allowing models to remain specific to the use case. Shared accounting data, controls and lineage provide the common foundation.

As private markets represent a growing share of industry revenue and portfolio complexity, firms will need an accounting platform that supports public and private assets within the same operating model.

From Automation to Agentic Operations

The next phase of AI adoption extends beyond task automation and decision support to systems that coordinate governed actions across workflows, data sources and operational processes.

For investment accounting, an agent can operate safely only when it has access to trusted data, clear permissions, explainable decision paths and a complete understanding of current operational state.

The accounting system must provide both the context an agent needs to reason and the authority model that determines what it is permitted to do.

As firms introduce agents across onboarding, reconciliation, oversight, reporting and exception management, competitive advantage will not come from access to models alone. It will come from providing those models with a real-time, governed and auditable accounting foundation.

Applications by Firm Type

Different firm types need different things from AI. It's a mistake to treat AI as a generic capability rather than as a specific answer to a specific operational constraint.

Asset Managers: Scaling Without Headcount Growth

Asset managers are being asked to support more products, asset classes and jurisdictions while protecting margins. The operating model must absorb that complexity without reproducing it in headcount, reconciliations and fragmented systems.

AI can support that objective when it operates within the accounting engine and against complete, current data. The value lies in reducing the manual effort required to onboard mandates, reconcile positions, investigate exceptions and support new investment products.

As assets under management continue to grow while profitability remains under pressure, operating leverage will depend increasingly on the architecture supporting investment operations.

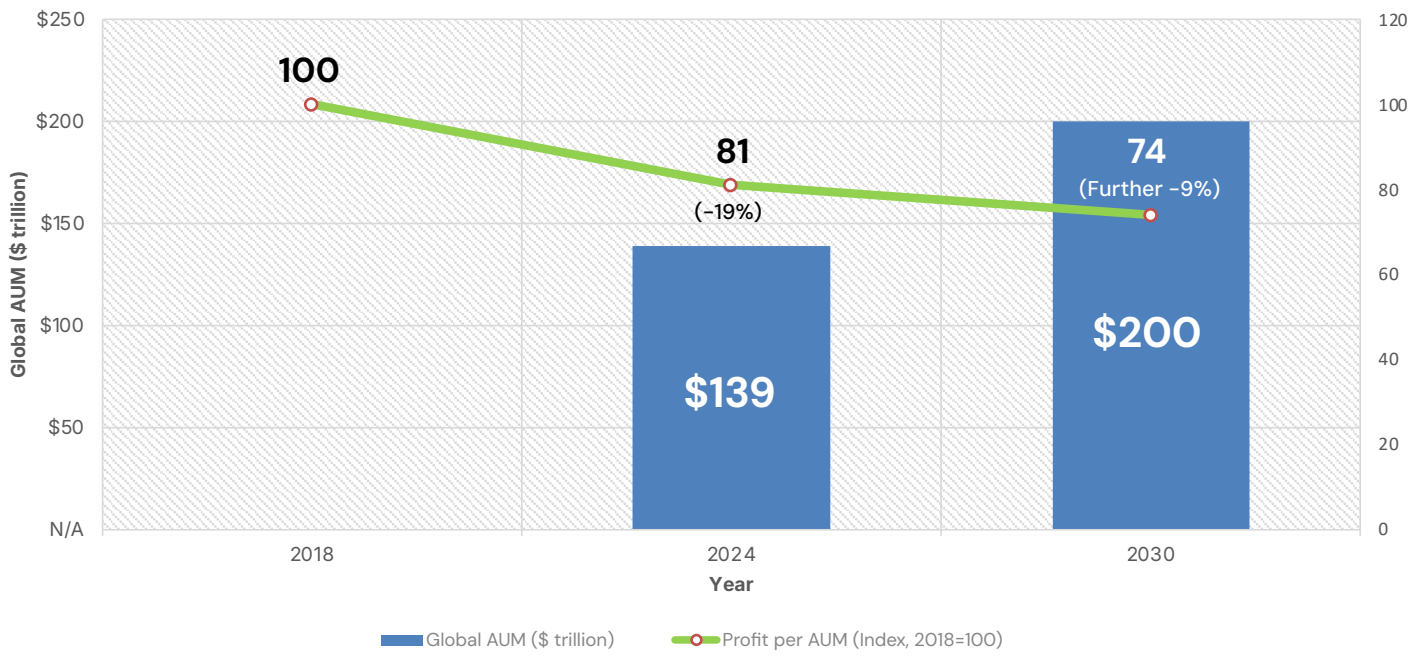

While assets are growing, profitability is not.

Asset Managers: Scaling Without Headcount Growth

Asset managers should evaluate AI by its ability to absorb product, asset-class and jurisdictional expansion without proportional growth in operational effort.

- AI-driven data mapping interprets diverse file formats, proposes classifications and reduces manual onboarding work for new mandates.
- AI-driven reconciliation resolves high-confidence breaks and filters the exception population so accounting teams focus on items requiring judgment.
- AI-driven anomaly detection identifies pricing errors, trade misbookings and corporate-action failures before they propagate through accounting and decision-making.

Asset Servicers: Quality Differentiation in a Fee-Compressed Market

Asset servicers face fee pressure while clients demand more tailored reporting, deeper oversight and faster turnaround. AI can increase operating capacity without proportional staffing by improving onboarding, exception management, continuous validation and client reporting.

The differentiator is not simply lower cost. It is the ability to support more complex client requirements, asset classes and operating models from the same platform.

Asset Owners: Consolidated Visibility and Independent Oversight

Asset owners have historically relied on managers and custodians to maintain books of record, then aggregated those views through internal workarounds. That model becomes harder to sustain as portfolios expand across public and private markets.

For asset owners, the value lies in consolidated oversight and independent validation. A unified accounting platform can ingest data from multiple managers, custodians and asset classes to provide a current view of the total portfolio and validate manager-reported information. AI helps interpret and normalize the heterogeneous, high-volume data required to support that oversight.

The Credibility Gap and How to Close It

Buyers face the same challenge with AI that they faced when distinguishing cloud-native systems from cloud-enabled ones: determining whether the capability is built into the operating foundation or layered above it.

The difference may be difficult to see in a demonstration and easy to miss in an RFP. It becomes clear in production, when the system must operate against full data volumes, live accounting workflows, control requirements and failure conditions.

AI layered above a batch engine remains constrained by batch processing. AI operating against asset-siloed ledgers remains constrained by incomplete context. AI connected to an overwrite-based record remains constrained by the absence of reconstructable history.

Vendor evaluation should therefore focus less on the quality of a curated demonstration and more on where the AI operates, what data it can access, what actions it can take, how its decisions are explained and whether it has been proven in production.

Vendor Questionnaire

Where does the AI operate: within the accounting engine or above it?

If the AI operates through a separate reporting layer, module or data warehouse, it can generally act only after accounting processes have completed. Ask for an architectural diagram showing where the model sits in the data flow and which workflows it can influence in real time.

Is the data bitemporal or overwrite-based?

Ask whether the system preserves every correction and restatement as a versioned event or overwrites prior state when new information arrives. This determines whether AI outputs can be reconstructed, explained and audited.

Can the AI use context across books and asset classes when the use case requires it?

Ask whether models operate against a unified dataset spanning books, public markets and private markets or against separate asset-class datasets that do not share context or signals.

How does the vendor handle explainability?

Ask the vendor to show the model's reasoning, supporting data and audit trail for a suggested match, flagged exception or automated resolution. Trusting the model without evidence is not acceptable for a regulated workflow.

What evidence shows that the AI operates at production volume?

Ask whether the AI has operated against millions of positions, complex funds and full transaction volumes in a live client environment. Request client references that can speak to the scale and operating outcome.

How does the AI behave under failure?

Ask how the system identifies model drift, failed data feeds and activity outside the model's confidence range. The system should degrade safely, preserve the workflow and escalate to human review.

What actions can the AI take, under what permissions and thresholds and where is human approval required?

Ask how the platform controls agent authority, segregation of duties, confidence thresholds, approvals and overrides. The vendor should be able to show exactly where automation ends and human accountability begins.

FUNDGUARD

Get True Intelligence Embedded Today with FundGuard

In 2024, the industry debated whether AI would matter for investment accounting. In 2025, firms began piloting individual use cases. In 2026, the focus has shifted to why many pilots fail to scale and what separates meaningful production outcomes from incremental automation.

The distinguishing factor is the architecture on which AI operates..

AI layered over legacy batch systems, asset-siloed ledgers and overwrite-based data remains constrained by that foundation. Cloud-native, event-driven, bitemporal and multi-book architecture allows AI to operate within the accounting flow rather than evaluate completed processes afterward.

That enables:

- Controls that identify issues before exceptions compound
- NAV data that is continuously validated and ready for formal approval
- Exception management that filters, explains and resolves within governed thresholds
- Operational capacity that can be redeployed to client service, oversight and growth

FundGuard provides a single, real-time accounting system of record across public and private asset classes, jurisdictions and investment products. Its AI-native platform supports investment operations and accounting from one unified transactional foundation.

FundGuard embeds intelligence within operational workflows while also providing governed access to trusted accounting data for clients' analytics, AI initiatives and external agents. As agentic operating models evolve, FundGuard provides the accounting foundation those systems require.

The firms that succeed with AI will not be distinguished by access to models. They will be distinguished by the quality of the operating foundation on which those models depend.

In investment operations, that foundation is accounting.

FundGuard provides that foundation by unifying books of record across public and private asset classes, jurisdictions and investment products in real time. The platform supports embedded AI capabilities and the next generation of governed, agentic operating models.

Request a demo today to see FundGuard in action

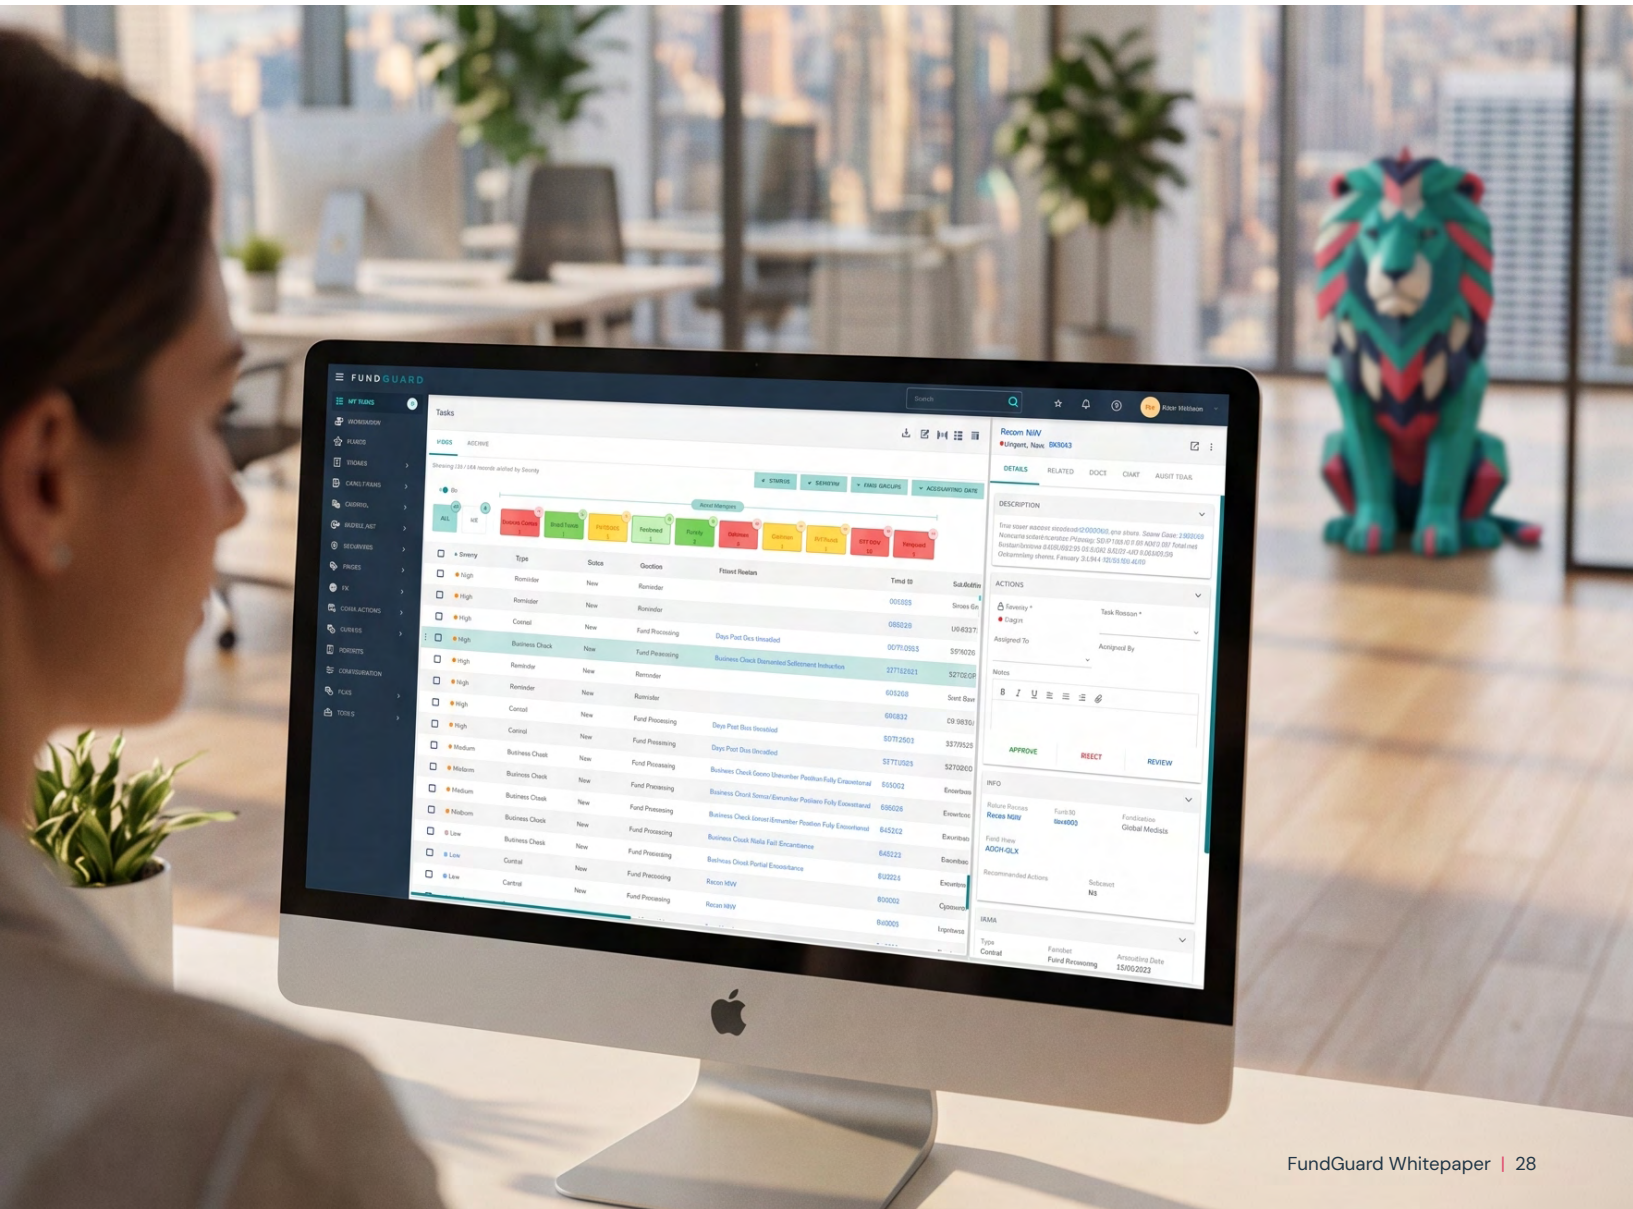

FUNDGUARD

Excellence Engineered.