AI in Asset Management

Smart Infrastructure for Self-Learning Portfolios

How AI is reshaping research, portfolio construction, risk management, and the operating model in BFSI.

Whitepaper

Written by Arjuna Gridharan
Financial Analyst

Executive Summary

AI as the core intelligence layer in asset management

The asset management industry is dealing with more data, tighter margins, and growing client expectations around speed and transparency. Traditional operating models and legacy systems are struggling to keep up, which makes small improvements harder to justify for firms looking to stay competitive long term.

Simultaneously, advances in artificial intelligence, cloud infrastructure, and modern data platforms have enabled scalable processing of diversified datasets. Capabilities once considered experimental are now dependable enough to support daily investment and operational workflows. As a result, asset managers now have an opportunity to turn data from an operational burden into a **measurable competitive advantage.**

> **Future differentiation in asset management will come from positioning AI as a core intelligence layer across the investment lifecycle.**

This white paper explains how firms can use AI in investment and operational workflows to create adaptive portfolio processes with a scalable operating model.

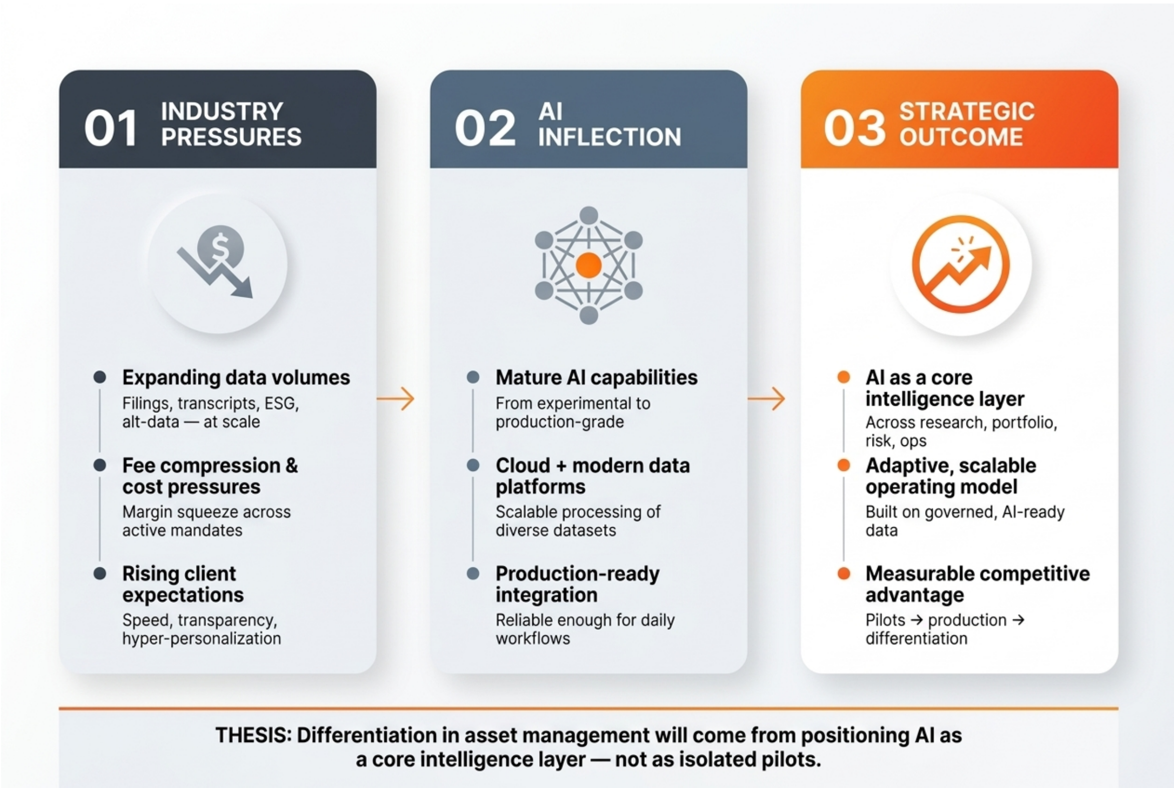

Implementation Hurdles

Unstructured data and fragmentation infrastructure - the real blockers

For asset and wealth managers we work with, the main blockers to meaningful AI adoption are unstructured data and fragmented systems, which continue to slow investment operations.

a. Unstructured Data

Industry surveys[1] estimate that 70–80% of **enterprise data is unstructured** and inaccessible to current AI workflows without preprocessing.

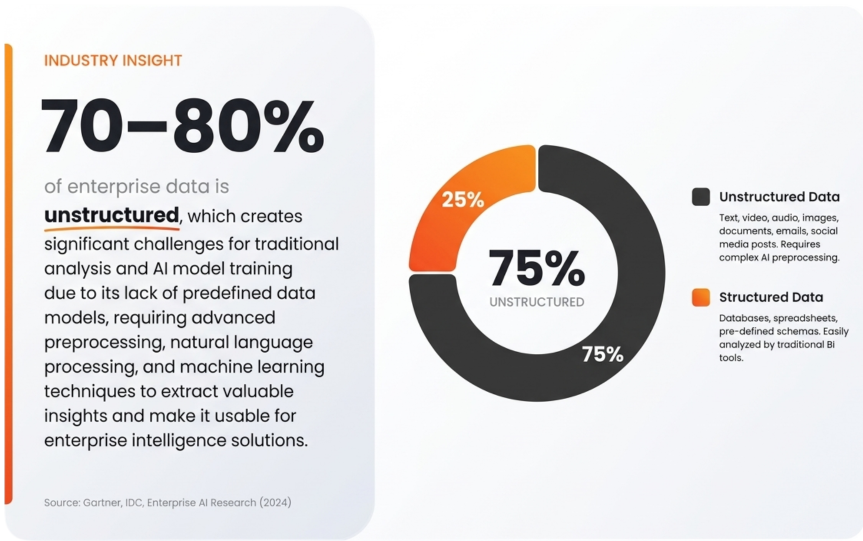

The consequence is headcount spent on low-value reconciliation rather than differentiated analysis, long cycle times from research insight to trade, and a backlog of AI use cases that stall when teams find documents that are not machine-readable or consistently tagged.

From a risk and compliance perspective, the stakes are even higher. Poorly parsed documents and incomplete client data can lead to inaccurate ESG reporting and gaps in trade surveillance.

Many firms launch underwhelming AI initiatives to unlock unstructured data, but without an **industrial-grade ingestion and enrichment layer** (OCR, entity resolution, topic and sentiment tagging, and document lineage) each new pilot creates its own mini-pipeline. The result is fragile, duplicated infrastructure and model outputs that cannot be reliably traced back to source documents when decisions are challenged.

b. Fragmented Data Infrastructure

McKinsey[2] notes that many managers operate with **siloed, asset-class-specific stacks** and lack a fit-for-purpose, front-to-back platform. This raises operational complexity and makes it hard to integrate multiple data sources for AI.

For AI and systematic strategies, fragmented systems make it harder to work with consistent and reliable data. Key datasets often update at different times, creating delays and reconciliation issues between OMS, PMS, and accounting systems.

Supervisors[3] have shown willingness to levy large fines when data quality and integration failures cause erroneous orders or misstated positions – treating poor data architecture as a risk-management failure rather than an IT issue.

This architectural debt turns prototypes into **"pilot factories"**, with models unable to scale without governed, AI-ready data. Each new product or data source compounds point-to-point complexity, raising cost and operational risk.

How Smart Infrastructure Helps

Taken together, unstructured data bottlenecks and fragmented infrastructure are **not 'IT problems'** – they are the main constraints on scalable AI in asset management. These constraints determine whether a **CIO** can safely deploy models into investment and client workflows, a **COO** can demonstrate control to regulators, and a **CFO** can capture efficiency gains rather than just fund experiments.

The fastest-moving firms are investing in **smart, domain-specific data infrastructure.** This includes governed data platforms and unstructured data hubs designed for financial content, with clearer ownership of data quality and lineage. This is the foundation on which AI moves from being a software promise to becoming an integral part of the operating model.

The Rise of AI

From innovation labs to daily investment workflows.

AI has transitioned from innovation labs to integration within daily workflows, where it consistently increases **speed, coverage, and control** across research, portfolio construction, and risk management. Leading systematic firms incorporate AI as a core component of their operating systems (supported by data engineering) rather than treating it as a peripheral initiative. Mainstream asset managers are steadily adopting these practices wherever the business case is compelling and operationally feasible.

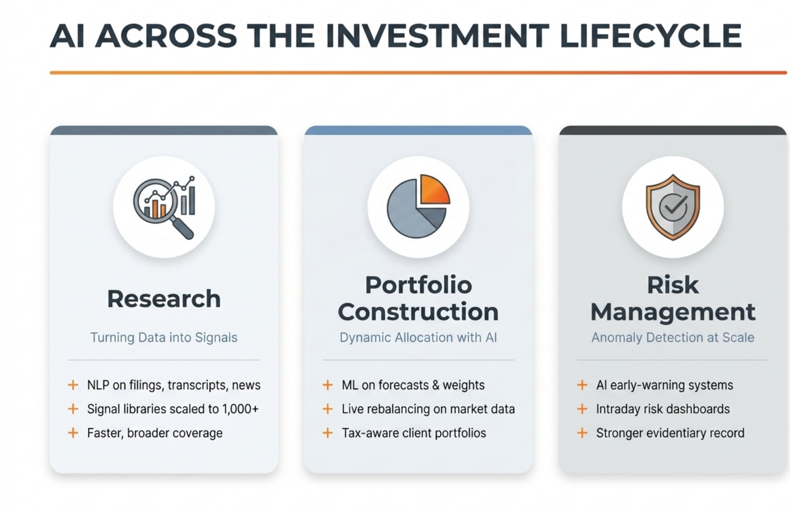

a. Research: Turning Data into Signals

AI acts as a **force multiplier in research**, allowing systematic teams to use machine learning and natural language processing to convert filings, transcripts, news, and alternative datasets into structured features for signal libraries. This extends analytical coverage beyond human limits.

For example, **BlackRock's systematic group** uses AI on both traditional and alternative data to generate security-level forecasts. The number of active signals has grown from a handful in the 1980s to **over 1,000 today**, with models retrained and evaluated regularly. Analysts can screen hundreds of earnings transcripts overnight using AI assistants. These systems flag shifts in guidance tone, accounting language, and market stance before insights reach stock-ranking dashboards for morning meetings.

Impact

Leading firms benefit from faster research, broader coverage of factors and themes, and a measurable reduction in manual effort spent on document parsing and data integration.

Strengthening Risk Mitigation & Resilience

Partnering with governments to enhance climate resilience through funding wildfire-resistant construction, improved flood infrastructure, and stricter building codes.

Promoting proactive risk reduction by incentivizing mitigation measures for property owners and expanding parametric insurance for faster, event-based payouts.

Example: State Farm offers discounts to homeowners who harden their properties against wildfires

b. Portfolio Construction: Dynamic Allocation with AI

AI now supports portfolio construction and rebalancing well beyond alpha ranking. Teams apply machine learning to improve forecasts and optimize portfolio weights and constraints. It also helps inform security selection for multi-strategy and liquid alternative products.

In wealth and asset management, AI-driven model portfolios use live market data, correlations, tax lots, and client preferences to generate suitable rebalancing actions. These systems **act faster than calendar-based models,** adjust risk budgets as volatility shifts, identify weak portfolio segments, conduct scenario analysis, and produce validated reallocations.

Impact

Integration with feature stores, model registries, and OMS/PMS platforms helps firms execute trades consistently and manage portfolios.

c. Risk Management

AI provides clear, measurable benefits in risk management. Studies[4] show that **AI-powered risk assessment is materially more accurate** at predicting volatility and downside risk than standard models. Firms use AI-driven early-warning systems to monitor large streams of market, credit, liquidity, and behavioural indicators – spotting anomalies and tension patterns before they surface in profit and loss.

Intraday risk dashboards use anomaly-detection models to analyse trade flows, factor exposures, and liquidity metrics across portfolios. These dashboards **issue alerts when portfolio behaviour resembles previous stress** events – even if headline Value at Risk remains within limits.

Impact

Fewer unexpected losses, faster hedging or de-risking decisions, and a stronger evidentiary record for regulatory or board inquiries about risk monitoring and management.

Smart Infrastructure with Indium

Across these domains, AI reaches operational maturity only when supported by strong data and workflow foundations. The key takeaway for firms is to invest in AI-ready infrastructure that supports reliable and scalable AI workflows. **Indium** is one such provider, specializing in AI-first, cloud-native data and analytics platforms for financial services.

Smart Infrastructure

The governed foundation for scalable AI

Smart infrastructure brings together the data and technology needed to support investment and client workflows with reliable data instead of disconnected spreadsheets and point-to-point feeds.

In practice, it replaces fragile, batch-oriented systems with a **unified foundation** where research, portfolio management, risk, and reporting all draw from the same trusted data sources. It also enables AI services to be deployed seamlessly, without repeated reconfiguration across the organisation.

Indium's AI-ready platforms help banks and wealth managers move away from fragmented data systems and make better decisions with connected data.

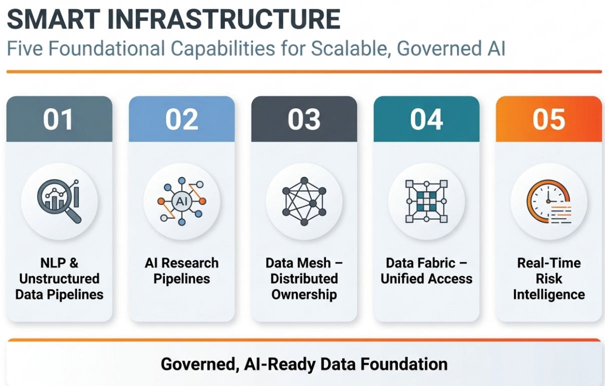

a. NLP & Unstructured Data Pipelines

Smart infrastructure treats unstructured content (filings, transcripts, ESG reports, credit agreements, research PDFs, and internal notes) as **first-class input.** Firms that cannot systematically process this data, risk missing alpha and valuable risk insights.

Indium's The Lifter illustrates this in practice. By automating extraction and analysis of complex financial documents, we have measured:

INDIUM IN ACTION | THE LIFTER

~40%	~90%	~45%
Faster Processing Time	**Extraction Accuracy**	**Reduction in Manual Effort**
Across complex financial documents	Across structured & unstructured inputs	On document-heavy BFSI workflows

These metrics indicate not only speed but also consistency and auditability. When ESG scores, credit flags, or controversy indicators are derived from documents through a governed pipeline, a firm can demonstrate exactly which sources and models supported a decision, reducing greenwashing and suitability risk.

b. AI Research Pipelines

Once the data foundation is in place, smart infrastructure unlocks **AI research pipelines** that move quants and data scientists from idea to validated signal, and to production through clear stages: feature access, model experimentation, back-testing, review, and controlled promotion.

In practice, this involves a **shared feature store** integrated with the native data platform – where standardised factors, NLP features, and alternative data signals are versioned and documented. Researchers build models (machine learning, hybrid quant, reinforcement learning) on containerised platforms with automated back-test harnesses.

Indium's Managed Analytics & MLOps

Helps improve operational efficiency by ~20% while reducing delivery costs by another ~20%. Models also remain continuously monitored and retrained.

c. Data Mesh – Distributed Ownership with Clear SLAs

Data Mesh decentralises data ownership to business domains, giving each domain responsibility for its own data products through clear quality and access agreements, reducing bottlenecks and accountability gaps common in centralised models.

It also establishes **explicit SLAs** for critical data such as NAV, exposures, and performance attributions. Investment and risk teams no longer have to wait for basic changes and can work with more consistent data definitions.

Indium supports this through **domain-driven data modelling, productization, and data-quality observability** frameworks.

d. Data Fabric – Fragmented Data, Unified Access

Data Fabric connects fragmented sources across multi-cloud, legacy, and on-premises systems through a unified logical layer that provides central governance, metadata management, policy enforcement, and consistent access. This makes **regulatory reporting, audits, and access control** more consistent and efficient.

Positions from custodians, trades from OMS/PMS, prices, curves, risk measures, and reference data can all be exposed through the same fabric, even if they sit on different platforms or in different regions.

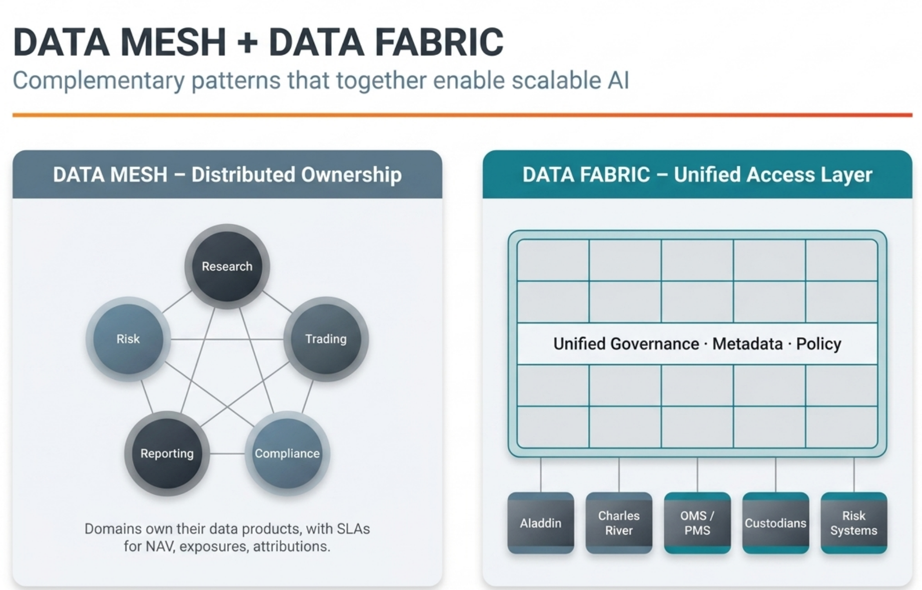

For instance, Indium's data fabric implementations connect platforms like core banking, BlackRock Aladdin, and Charles River into a unified layer. This gives financial firms better visibility into their data while reducing integration effort for new AI initiatives.

> *Together, Data Mesh and Data Fabric reduce reconciliation effort and shorten the path from concept to production – by making integrations reusable and dependable.*

e. Real-Time Risk Intelligence

The final element of smart infrastructure is real-time risk intelligence, the ability to monitor exposures, stress behaviour, and anomalies across portfolios as markets move, rather than waiting for end-of-day reports. As markets and regulators expect near-real-time oversight, firms need continuous access to risk and compliance data to maintain effective controls.

When an earnings shock hits a concentrated sector, Indium's platform can recalculate factor exposures, scenario losses, and liquidity metrics within seconds, alert the relevant desk and risk team. If limits are breached, it can automatically throttle trades or propose hedges, all with a full audit trail. Tools such as Striim and Azure Synapse enable real-time replication from core systems to the cloud.

Conclusion

Smart infrastructure is the precondition for more adaptive and self-learning portfolios. With native data infrastructure, unstructured pipelines, AI research tooling, and real-time risk intelligence in place, firms can safely let models update signals, propose allocation changes, and respond to new information within clear risk and governance boundaries. The next pillar builds on this foundation.

Self Learning Portfolios

Adaptive, RL-driven strategies under real-world constraints.

The evolution of asset management is increasingly defined by the transition from **static portfolio construction** to adaptive, self-learning systems. Unlike traditional approaches that rely on periodic rebalancing and fixed assumptions, self-learning portfolios continuously evolve, incorporating new data, market feedback, and investment outcomes.

Indium's work on AI-led investment platforms (including advisory experiences for HNIs and UHNIs) shows how data and models can be embedded directly into live portfolio and client workflows.

At the core of this transformation lies the application of **reinforcement learning (RL) and hybrid AI models,** which enable portfolio strategies to dynamically adjust to changing market conditions while adhering to investment constraints.

a. From Static Allocation to Adaptive Learning Systems

Traditional portfolio construction techniques (such as mean-variance optimisation) operate under static assumptions about risk, return, and correlation. Effective in stable environments, these models struggle to adapt to:

- regime shifts
- market volatility
- non-linear relationships across assets

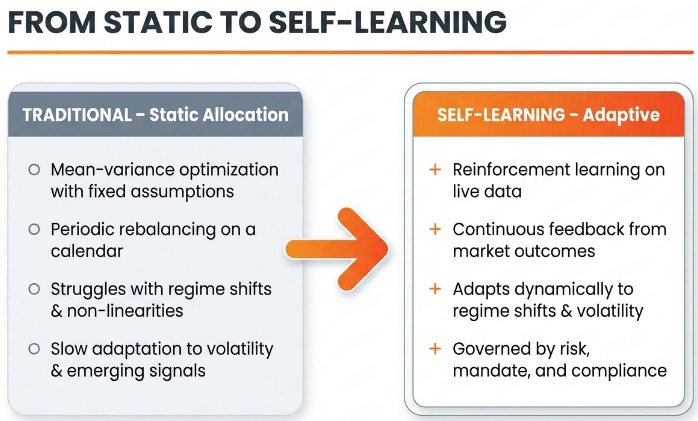

Self-learning portfolios address these limitations by introducing continuous feedback loops, where investment decisions are evaluated against outcomes and iteratively improved over time.

b. Constrained Reinforcement Learning for Portfolio Optimisation

A key part of self-learning portfolios is constrained reinforcement learning, which helps AI models make investment decisions within defined risk and regulatory limits.

In a reinforcement learning setup:

- The agent = portfolio strategy
- The environment = financial market
- The state = current market conditions + portfolio composition
- The action = asset allocation decisions
- The reward = risk-adjusted return (e.g., Sharpe ratio, alpha)

The objective is to learn a policy, or a mapping from states to actions, that maximizes long-term reward.

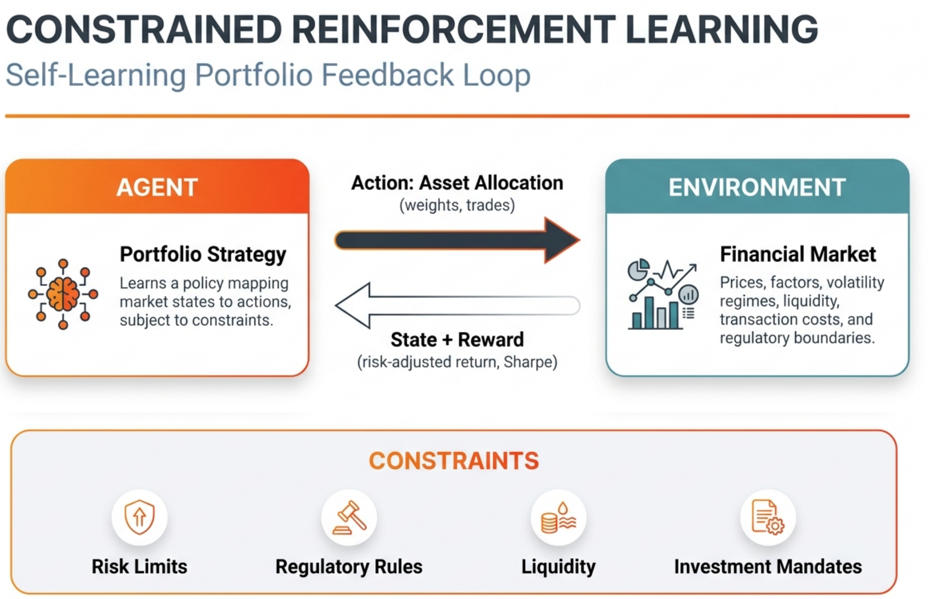

Unlike theoretical models, portfolio management must satisfy constraints such as:

- Risk limits
- Regulatory requirements
- Liquidity
- Investment mandates.

This transforms the problem into a constrained optimization framework, where the model must learn optimal strategies without violating predefined boundaries. For enterprises exploring these approaches, Indium helps integrate RL models into existing systems and controls.

Why Reinforcement Learning Matters

Constrained RL enables:

Dynamic allocation under uncertainty – Studies on S&P 500[5] sector indices showed annual returns and Sharpe ratios that were roughly 1.85× higher than mean-variance portfolios over a 10-year backtest.

Continuous learning from market feedback – Research[6] showed portfolio Sharpe ratios improving from 0.42 to 0.483 in highly volatile market conditions through continuous position adjustments.

Integration of multiple objectives (return, risk, cost) – Recent studies[7] show that when transaction costs, leverage, and regulatory constraints are embedded directly into the reward, RL policies can jointly optimise return, risk, and cost. This approach outperforms Markowitz strategies that ignore cost terms in realistic trading environments.

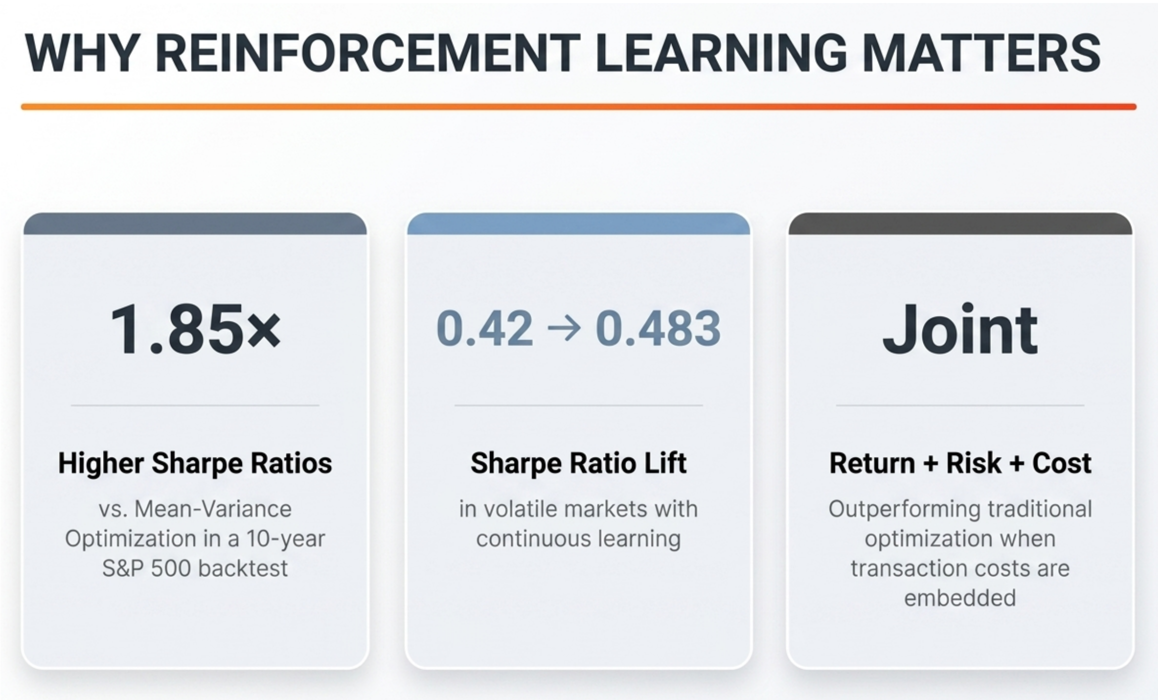

> This is a significant shift from traditional optimization toward policy-driven investment systems that evolve over time.

The Road Ahead

An AI-first operating model - scalable, explainable, sound.

Smart infrastructure and self-learning portfolios are becoming strategic imperatives. Smart infrastructure creates the governed foundation for scalable AI adoption. Self-learning portfolios convert that foundation into adaptive, policy-driven investment outcomes.

Indium is positioned as an **AI-first services partner** with enterprise-grade platforms and execution capability. **The Lifter**, an Agentic AI platform, builds end-to-end intelligence for legacy systems, applications, and data flows to reduce manual legacy archaeology. It maps dependencies, quantifies technical debt, and exposes modernization risk before capital is committed.

A[i]LPHA, Indium's investment analyst solution, automates the investment-research workflow from ingestion to insight synthesis. It processes structured and unstructured financial data through NLP and AI pipelines – turning filings, transcripts, ESG reports, and market data into **consistent, auditable outputs.**

Indium's combination of **agentic AI platforms, research and portfolio tooling, applied AI expertise, and data engineering depth** is designed for that mandate: to help asset managers move from **pilots to production,** and from fragmented initiatives to an AI-first operating model that is scalable, explainable, and financially sound.

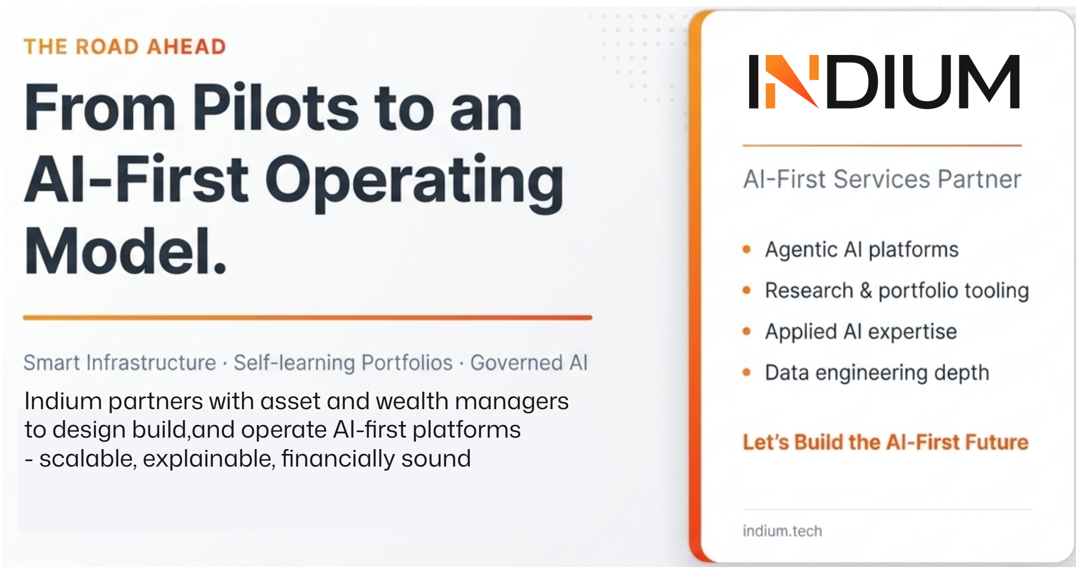

References

1 Buller, Rob. "The Untapped Power of Unstructured Data in Enterprise AI." Forbes, November 24, 2025. https://www.forbes.com/councils/forbestech council/2025/11/24/the-untapped-power-of-unstructured-data-in-enterprise-ai/

2 "Asset Management 2025: The Great Convergence." McKinsey & Company, 2024. https://www.mckinsey.com/industries/financial-services/our-insights/asset-management-2025-the-great-convergence

3 Hoffman, Karen. "US Financial Firms Face Growing Regulatory Fines for Poor Security Management." SC Media, August 4, 2022. https://www.scworld.com/analysis/us-financial-firms-face-growing-regulatory-fines-for-poor-security-management

4 Indium Success Story: "Cost Saving with Striim & Azure Synapse – Global Bank Analytics." https://www.indium.tech/success-stories/cost-saving-with-striim-azure-synapse-global-bank-analytics/

5 "Prediction of Realized Volatility and Implied Volatility Indices using AI and Machine Learning: A Review." International Review of Financial Analysis, 2024. https://doi.org/10.1016/j.irfa.2024.103221

6 "Deep Reinforcement Learning for Optimal Portfolio Allocation: A Comparative Study with Mean-Variance Optimization." ICAPS 2023. https://icaps23.icaps-conference.org/papers/finplan/FinPlan23_paper_4.pdf

7 "Reinforcement Learning for Portfolio Optimization with a Financial Goal and Defined Time Horizons." arXiv:2511.18076. https://arxiv.org/abs/2511.18076

8 EFMA 2022 Annual Meetings, Rome – Paper ID 195. https://www.efmaefm.org/0EFMAMEETINGS/EFMA%20ANNUAL%20MEETINGS/2022-Rome/papers/EFMA%202022_stage-3032_question-Full%20Paper_id-195.pdf

About Indium

We are an AI services company specializing in Agentic AI, Data & Analytics, Application Engineering, and Quality Engineering.

By combining autonomous AI systems, modern engineering, and data intelligence, we build innovative AI solutions that drive measurable business outcomes and long-term value.

With 5,000+ associates worldwide, Indium partners with Fortune 500 companies, Global 2000 enterprises, and leading technology firms in Financial Services, Healthcare, Manufacturing, Retail, and Technology. Our teams operate in North

USA	INDIA	UK	SINGAPORE
Cupertino | Princeton Toll-free: +1-888-207-5969	Chennai | Bengaluru | Mumbai Hyderabad | Pune Toll-free: 1800-123-1191	London Ph: +44 1420 300014	Singapore Ph: +65 6812 7888

www.indium.tech

For Sales Inquiries
sales@indium.tech

For General Inquiries
info@indium.tech