

AI in Asset Management

Most managers have AI. Far fewer have operationalised it — and the gap is where the advantage sits.

Having it vs operationalising it

Almost every asset manager now says it uses artificial intelligence. Far fewer can show that they operate it. The interesting question has moved on from whether a firm uses AI to how well that AI actually runs once the demo is over and a portfolio is on the line.

This page sets out what AI in asset management means, where it is applied across an investment firm, how quickly adoption is moving from chatbots to agentic systems, and how an allocator can tell real capability from "AI washing". One frame organises the rest: having AI and operationalising it are different achievements, and the distance between them is where most of the value, and most of the risk, sits.

The view is that of an operator rather than a vendor or an advisor. IMS Group is a technology-forward private markets investment group and partnership of family offices that runs its own AI lab, IMS Labs. That vantage point shapes the emphasis here, away from the familiar "AI is transformative" narrative and toward the harder, more useful question of what it takes to make AI work inside a real investment process, in public and private markets alike.

What "AI in Asset Management" Actually Means (and What It Doesn't)

AI in asset management is the use of machine learning and, increasingly, agentic systems to support investment research, portfolio construction, risk management and operations across an investment firm. That is the working definition, and it is deliberately narrow: it is about applying these techniques to the business of investing money, not about managing physical equipment.

The distinction matters because the same three words describe two unrelated fields. In investment asset management, "assets" are securities, funds, private holdings and the portfolios built from them. The job is allocation, selection and risk. AI here reads filings, extracts signals from unstructured data, stress-tests portfolios and automates reporting.

Most of what searchers want sits in the first field, even though the second one ranks for the same phrase. The rest of this page stays firmly in investment asset management: how an investment firm uses AI to research, build, monitor and operate portfolios, and what separates a firm that runs that capability from one that merely advertises it.

Investment AM vs enterprise AM: the same words, two different fields

Enterprise asset management is a different discipline entirely. There, "assets" are physical: pumps, turbines, vehicles, plant. AI is used for predictive maintenance, sensor analysis and asset-lifecycle planning on the factory floor or across an infrastructure estate. It is a legitimate and valuable use of the technology, but it has nothing to do with managing capital.

Confusing the two leads readers somewhere unhelpful, into industrial maintenance content when they came to understand investing. For clarity: when this page says asset management, it means the management of investment portfolios. The machine-learning techniques may rhyme across the two fields, but the objectives, the data and the people involved do not.

Where Asset Managers Actually Use AI (The Use-Case Map)

AI is not one capability inside an investment firm; it is several, applied to different parts of the process. The cleanest way to see the field is to map it to the work itself. Five areas account for most of what is in production today, and a private-markets line runs through each of them.

- **Investment Research.** This is where AI earns its keep first. Models read and summarise filings, transcripts, news and other unstructured data at a scale no research desk could match, extracting signals and surfacing the document that matters from a pile of thousands. In manager research and due diligence, the same techniques compress weeks of reading into a working brief. The constraint is not capability but judgement: the output is a faster starting point, not a decision.
- **Portfolio Optimisation.** AI supports construction, rebalancing and scenario modelling, testing how a portfolio behaves across many possible futures rather than a handful of hand-built cases. Used well, it widens the range of scenarios a team can consider and sharpens the trade-offs; it does not replace the mandate, the constraints or the allocator's call on risk.
- **Risk Management.** Here the work is monitoring and measurement: supporting value-at-risk estimates, flagging concentration and exposure drift, and watching positions continuously rather than at month-end. Model risk itself becomes a live concern, because a risk system built on a model is only as trustworthy as that model and the oversight around it.
- **Operational Efficiency.** The least glamorous use case is often the most immediate. Reporting, reconciliation and other back-office workflows are heavy with repeatable, rules-based tasks, and that is precisely where automation removes cost and error. Much of the early, measurable productivity gain in the industry comes from here, not from the research desk.
- **Personalisation.** On the client side, AI tailors reporting and communication, shaping how a portfolio is explained to the person who owns it. For an institutional allocator this matters less than research or risk, but it is a real and growing part of how managers differentiate their service.

Across all five, the same pattern holds. AI compresses the work and widens what a team can consider; people still exercise the decision. Technology produces the package; judgement signs it off.

How Fast Is Adoption Moving? From Chatbots to Agentic AI

The shape of adoption has changed in the last two years. The first wave was conversational: chatbots and copilots that answered questions and drafted text, useful but bounded, a demo of what the technology could say. The second wave is agentic: systems that do not just answer but act, executing multi-step tasks across tools and data with limited supervision. The shift toward [agentic AI as an operating system](#) rather than a chat window is the single most important development in how investment firms use AI, and it is the development the industry is now racing to absorb.

The headline figures point in the same direction, though they should be read as a range rather than a single number. McKinsey has estimated that AI and agentic adoption could improve the cost base of asset management by roughly 25% to 40% over time (McKinsey, 2025); the figure is an estimate of potential, not a measured result, and it engages the scale of the opportunity without prescribing a methodology. For a wider frame, J.P. Morgan Asset Management estimates that generative AI could lift developed-market labour productivity by roughly 1.4% to 2.7% per year over a decade (J.P. Morgan Asset Management, 2025); this is a whole-economy estimate rather than a figure measured inside asset-management workflows, and it is included to set the macroeconomic backdrop, not as a directly comparable asset-management result. Grant Thornton's 2025 survey of around 500 senior executives globally (Grant Thornton, 2025) provides the breadth: at the level of firm-wide adoption, AI in investment management is broad in intent, uneven in execution, and concentrated for now in operations and research rather than in core investment decisions.

The numbers support two readings at once. Direction is clear: adoption is accelerating, and the agentic shift is real, not hype. Magnitude is contested: McKinsey's 25% to 40% is an estimate of potential, and the distance between that potential and what most firms have so far made AI do is precisely the gap that matters. That gap is the subject of the next section.

Having It vs Operationalising It — The Execution Gap

This is the distinction that matters most, and it is rarely named. Almost every firm now has AI ambition; far fewer have AI capability that runs. The difference between appetite and operating reality is the defining feature of where the industry stands today.

The reason is structural. A model is not an operating system. A capable model, bought or accessed off the shelf, is the easy part; making it useful inside an investment firm requires four things the model itself does not provide. It needs data: clean, connected, proprietary information for the model to work on, which most firms do not have in usable shape. It needs workflow: the model has to sit inside how the firm actually researches, decides and reports, not alongside it as a novelty. It needs oversight: human review, model-risk governance and the authority to overrule the machine, built in rather than bolted on. And it needs integration: the connective engineering that turns a clever demo into a system the firm relies on every day. Assemble those four and AI operates; leave any of them out and the firm has a capability it owns but cannot run.

This is why "we use AI" and "we operate AI" are not the same claim, and why the distance between them tends to be measured in years of engineering rather than a procurement decision. Industry research has repeatedly found that only a small minority of organisations have scaled an AI agent into genuine production in any function; the ambition is near-universal, the operating capability is not. Building it is a question of appetite versus capability, examined directly in [appetite vs capability — the AI skills gap](#), and of why a firm might choose to engineer the capability in-house at all, the subject of [why an investment group runs its own AI lab](#). The point to hold at the category level is simpler: production beats demonstration, and the gap between the two is where credibility is won or lost.

Why "we use AI" and "we operate AI" are different claims

A firm that uses AI has access to a tool. A firm that operates AI has rebuilt part of how it works around that tool, with the data, workflow, oversight and integration to match. The first is a statement about software the firm can reach; the second is a statement about how the firm runs. Most marketing language blurs the two deliberately, which is exactly why an allocator needs a way to tell them apart.

The Credibility Test: Telling Real AI Capability From "AI Washing"

"AI washing" is the practice of marketing AI capability a firm does not genuinely have, or materially overstating what it does. It is the predictable consequence of a moment when claiming AI is rewarded and verifying it is hard, and it is a live enough risk that the term now appears in the industry's own list of challenges and in the attention of financial regulators. The useful response is not cynicism but a test.

An allocator can apply four practical checks before taking an AI claim at face value. The first is production evidence: not "we are exploring AI" but what is running today, in which workflow, with what observable effect. A firm that operates AI can describe the system; a firm that washes can only describe the ambition. The second is human oversight: a credible operator can explain exactly where a person reviews, overrules and signs off, because unsupervised models in an investment process are a risk to be governed, not a feature to be advertised. The third is governance: model-risk controls, validation and documentation, the unglamorous scaffolding that separates a system you can rely on from one you merely hope works. The fourth is measurable workflow change: a real capability shows up as a process that demonstrably runs differently, faster diligence, continuous monitoring, automated reporting, not as a logo on a slide.

This test sits comfortably alongside the governance stance taken by bodies such as the CFA Institute, whose work on AI in asset management stresses transparency, oversight and the management of model risk as the conditions for using these tools responsibly (CFA Institute, 2025). The credibility floor is the same one a careful allocator would set anyway: human oversight and model-risk governance are not optional extras on top of an AI capability; they are the evidence that the capability is real. No fee comparison and no named competitor is required to apply the test. The questions do the work.

What This Means for Private-Markets Allocators

For an allocator in private markets, the operationalisation gap is not an abstraction; it is a selection criterion. Most of the published guidance on AI in asset management is written for public-markets managers, so the questions that matter to private capital fall to the allocator to ask directly.

The places AI changes the private-markets process are specific. Deal sourcing widens the aperture on off-market opportunities. Diligence and underwriting compress weeks of document review toward hours, which changes not just cost but the number of opportunities a team can credibly assess. And monitoring turns illiquid holdings from a quarterly valuation exercise into a continuous read of operating signals. At each of these points, a firm that operates AI does the work differently from one that does not, and an allocator evaluating managers can ask to see exactly that difference. The same dynamic runs through adjacent private-markets strategies, including [private credit](#), where underwriting quality and speed are the whole competitive question. It also connects to the infrastructure layer, where [tokenised assets need an authentication layer](#) and AI sits behind the verification of what an on-chain instrument represents. All of this is one part of the firm's wider reading of [the next generation of private capital](#).

Where This Is Heading

Over the next cycle, the dividing line in asset management is unlikely to be whether a firm uses AI; that question is already settled. It will be whether a firm operationalises it, turning a capability it owns into a system it runs, with the data, workflow, oversight and integration that distinction requires. Adoption will keep accelerating and the agentic shift will keep maturing, though the realised gains should still be read as a range rather than a single promised number.

The firms that close the gap between having AI and operating it will compound that advantage across cycles, particularly in private markets, where the published field has barely begun to look. To see how an operator approaches this in practice, explore how [IMS Labs](#) builds and runs the firm's AI capability in production ([imslabs.ai](#)), and read across the cluster, from the agentic shift to the skills gap, for the detail beneath this category view.

This article is provided by IMS Group for information purposes only and does not constitute investment advice, an offer, or a solicitation. Figures are sourced and dated where cited; forecasts are projections and stated as ranges.

Frequently asked questions

What is AI in asset management?

AI in asset management is the use of machine learning and, increasingly, agentic systems to support investment research, portfolio construction, risk management and operations across an investment firm. It applies these techniques to the business of investing capital, distinct from enterprise asset management, which uses AI to maintain physical equipment. The technology compresses and supports the work; people still make the decisions.

What is the role of AI in asset management?

AI's role is to augment the investment process, not to run it. It reads and synthesises research at scale, supports portfolio construction and scenario modelling, monitors risk continuously, and automates reporting and reconciliation. In private markets it also helps source deals and compress diligence. Across every use case the pattern holds: technology produces the analysis, and human judgement exercises the decision under proper oversight.

How are asset managers using AI in 2026?

Adoption is broad in intent and uneven in execution. Most measurable gains so far sit in operations and research rather than core investment decisions, and the field is shifting from conversational chatbots toward agentic systems that execute multi-step tasks. McKinsey estimates a potential 25% to 40% cost-base improvement for asset management over time (McKinsey, 2025), while J.P. Morgan Asset Management's roughly 1.4% to 2.7% figure is a whole-economy developed-market productivity estimate rather than an asset-management result. The wide gap between such potential and what firms have so far operationalised is the point.

What are the risks and where is human oversight still needed?

The principal risks are model risk, over-reliance on unvalidated outputs, and "AI washing", where capability is marketed but not genuinely operated. Human oversight remains essential at the point of decision: a person must review, validate and retain the authority to overrule the model. Governance, model-risk controls and transparency, consistent with CFA Institute guidance, are the conditions for using these tools responsibly rather than optional additions.

Sources & important information

1. CFA Institute (2025). AI in Asset Management: Tools, Applications, and Frontiers — governance, transparency, human oversight and model-risk management as the conditions for responsible use; entity co-citation source for the concept. CFA Institute.
2. McKinsey & Company (2025). Analysis of AI and agentic adoption in asset management — estimated potential cost-base efficiency improvement of approximately 25% to 40% over time. Figure is an estimate of potential, engaged here but not reproduced as methodology. McKinsey & Company.
3. J.P. Morgan Asset Management (2025). Estimate that generative AI could lift developed-market labour productivity by approximately 1.4% to 2.7% per year over roughly a decade. A macroeconomic, whole-economy forecast, not a measure of productivity inside asset-management workflows; attribution qualitative. J.P. Morgan Asset Management.
4. Grant Thornton / ThoughtLab (2025). Global survey of approximately 500 senior executives on AI in asset management (conducted Q3 2025) — adoption breadth, concentrated for now in operations and research. Used as an adoption anchor. Grant Thornton.
5. ION Group (2024). How AI Is Shaping Asset Management — structural/use-case reference for the explainer scaffold (not authority-by-citation). ION Group.
6. Russell Investments (2024). The Value of AI in Asset Management — practitioner framing of AI for operational efficiency and manager research (structural reference). Russell Investments.

Point-in-time figures are volatile and should be re-verified on the date of publication.

Important information

This document is provided by IMS Group ("the Group") for general informational and educational purposes only. It does not constitute investment advice, an offer to sell, or a solicitation of an offer to buy, any security or investment product, and should not be relied upon as the basis for any investment decision.

Forecasts and projections are illustrative only and are not guarantees of future results. Past performance does not predict future returns. Point-in-time figures cited are volatile and should be re-verified on the date of publication. © 2026 IMS Group. All rights reserved.

AI, read from the buy-
side — operating layer,
not chatbot.

IMS Group

info@groupims.com

groupims.com

IMS Group operates across investment management, venture building, sports finance and proprietary technology — IMS Capital Management, IMS Digital Ventures, IMS Labs and IMS Sports Capital. This brochure is for information purposes only and does not constitute investment advice or an offer. © 2026 IMS Group.