AI ON THE BUY SIDE: THE ACTIVATION GAP

SimCorp's 2026 survey found 70% of buy-side firms now deploy AI for front-office functions, up from 10% the prior year. AIMA puts the number even higher, with 95% of hedge fund managers using generative AI in some capacity.

The activation numbers tell a different story. Carne Group found that only 18% of firms utilize AI in core investment operations. SimCorp's survey found that 63% of firms lack unified data across front, middle, and back offices, meaning the majority are deploying AI in environments where the model can't see the full picture. A December 2025 survey by OriginTrail and Fortune's found only 14% of CFOs reported measurable ROI from AI investments.

So what does "using AI" actually mean at most buy-side firms right now? Enterprise LLM licenses, RFP generation (34%, according to Carne Group), marketing materials (42%), proxy voting summaries, and email copilots. Billions in aggregate spending—and the median outcome is still a chatbot that summarizes earnings calls.

The four of us have spent a lot of time in the weeds on this, together, roughly 125 years of industry experience. Dan Royal spent most of his 40-year career running institutional trading desks and brings a practitioner's perspective on exactly where the plumbing tends to break. DJ Adams spent 25 years on the buy side before moving into technology advisory, where he now works on initiatives such as enterprise platform migrations and agentic trading co-pilots. Glenn Lesko and Nilesh Nanavati have spent the past five years building OPCO Advisory, advising across the capital markets fintech ecosystem and working closely with companies tackling these problems from the technology side.

Between us, we've sat on every side of the table — buying, building, and now helping firms close the gap between what AI can do and what their infrastructure will actually support.

Most of the commentary on AI and the buy side comes from technologists looking in. We're offering the opposite: an inside-out view from practitioners who have lived through decades of change in this industry. One thing that stands out in nearly every conversation is that people running desks and managing teams don't have the bandwidth to track the dozens of companies now building solutions to their specific problems. At the same time, many of the companies building those solutions often don't fully understand the operational constraints they're solving for. We sit in that gap — which is why we're writing this.

Here's what we're seeing: the bottleneck is infrastructure, not models; a new generation of domain-specific companies is emerging to address it; and firms are rethinking both the architecture and the organizational shifts needed to close the activation gap.

THE CONSTRAINT IS RARELY THE MODEL

The first thing we see consistently is this: the bottleneck is the plumbing.

FRAGMENTED PLUMBING: Buy-side firms operate across multiple interconnected systems — EDM, PMS, OMS, EMS, risk engines, compliance platforms, reporting, fund accounting, and reconciliation — each with its own data model, reconciliation cycle, and audit requirements. Much of the data feeding front-office decisions originates in back and middle-office platforms built for end-of-day batch processing, designed for settlement cycles and regulatory reporting rather than real-time analytics. As Carne Group puts it, "Digitising a broken process just digitises inefficiency." That matches what we hear from nearly every firm we speak with.

STALE AND CONFLICTING DATA: The quality and structure of that data matters enormously. A portfolio manager may want to stress-test a position, but the risk system, OMS, and compliance platform each hold a different view of the current holdings. Someone pulls three CSVs, reconciles them in a spreadsheet, and by the time the analysis is complete, the data is already stale. Dan lived this at Janus Henderson for years: as an example, the time spent pulling together reports for trustees, best execution committees, and compliance — assembling TCA data, commission data, and qualitative data from separate systems, formatting it into a unified presentation, double-checking accuracy, inevitably finding discrepancies to trace back through the chain — was enormous. And that's just reporting on historical activity. For real-time optimization, the challenge grows by an order of magnitude.

CLIENT CONFIDENTIALITY: This compounds everything. Buy-side firms handle sensitive data under strict regulatory and contractual obligations, and sharing it with external AI providers raises questions that compliance and legal teams are still working through. The hesitation is entirely rational: these are the kinds of risks that can cost people their jobs and firms their licenses.

THE INNOVATION-PRACTITIONER GAP: Then there's the gap between the people building AI solutions and the people who would use them. At every large firm, innovation teams and practitioners tackle the same problems from opposite directions and often speak different languages. Innovation teams think in terms of capability and platform transformation; practitioners on the desk are focused on ensuring trades clear and best execution is documented. Even at the largest asset managers, the connection between innovation teams and practitioners falls well short of what's needed.

Bridging that gap requires both courage and leadership, as well as a new way of thinking. It takes someone senior enough to declare: this is the direction we're going, and we will invest in figuring it out, even without perfect clarity on the outcome. But it also requires a shared language: innovation teams need to understand what keeps traders up at night, and traders need to see that the manual workarounds they've normalized are not permanent features of the industry.

The last point is harder than it sounds, because change management in a large buy-side organization is a discipline, not a pep talk. Manual workarounds become tribal knowledge — the spreadsheet that reconciles three systems, the morning routine of 20 phone calls, the compliance check that lives in someone's head — and the people who built their careers around those workarounds have rational reasons to resist replacing them. We've seen firms succeed when leadership treats adoption as a measured outcome: not "did we deploy?" but "did workflows actually change, and can we see it in

the numbers?" The firms that fail typically announce a technology initiative and assume behavior will follow. It doesn't. Incentives, training, and structured feedback loops between the desk and the technology team are what separate a deployment from true activation.

GENERIC AI ISN'T BUILT FOR FINANCIAL MARKETS

In an industry that demands deterministic results — for both regulatory and competitive reasons — the out-of-the-box AI solutions from the large providers are often not the answer.

When we speak with firms, we rarely hear that they're getting tremendous value from the features that come automatically with their enterprise AI licenses. The major LLM providers have built extraordinary general-purpose models, but financial markets require a level of precision those tools weren't designed to deliver. You can't be almost right on a compliance check or approximately correct on a reconciliation. And you certainly can't have a model hallucinate a restriction that doesn't exist or miss one that does.

Even the largest banks — institutions with dedicated AI labs and nine-figure technology budgets — are finding that integrating across siloed systems is harder than building the models themselves. The models work. Getting them to see the same data at the same time, across platforms that were never designed to talk to each other, is where things break down.

There are two fundamentally different types of AI work on the buy side, and missing that distinction is part of the reason the activation gap exists. First are processes that must be deterministic — compliance, accounting, performance reporting — where the answer must be right to the third decimal place. In these cases, AI functions more like a calculator operating across unified data: powerful, but precise. Second are processes that benefit from probabilistic reasoning — venue selection, portfolio construction, alpha generation — where the goal is to surface the insight nobody else saw, what DJ calls the "AlphaGo Move 37," the move that breaks from conventional patterns and turns out to be right.

These require fundamentally different approaches. General-purpose LLMs tend to blur the distinction, and most firms are discovering that the deterministic work is both more urgent and more tractable. It hits operating margins directly — and that's what gets the C-suite's attention. The market needs domain-specific solutions built by people who understand financial markets — and those solutions are starting to emerge.

WHAT'S EMERGING

A new generation of finance-focused AI companies has matured over the past several years, and the pace is accelerating. Most were built by people who came out of the industry — portfolio managers, quant analysts, compliance officers — who understood what general-purpose models couldn't do and built around those gaps. We know many of these companies personally, which is why we can speak to their capabilities — and it's worth noting that upfront.

Document parsing and data readiness have been among the most established use cases, having developed over roughly five years. Companies in this space can extract key data points from complex legal documents — such as loan

agreements and CLO indentures — that would take a human analyst hours to process. Some of the most promising firms are using AI to prepare data so it can actually be used by AI. In other words, the data readiness problem is itself becoming a market of its own.

Execution analytics and trading intelligence are gaining traction for firms willing to look beyond the big names. A growing number of companies are applying AI to trading analytics — taking regulatory data on venue performance and market microstructure and making it actionable, so a trader's morning begins with an intelligent briefing rather than a stack of broker commentary. Others monitor FIX message flows in real time, catching anomalies and compliance issues that might otherwise surface hours later in a reconciliation break. The most advanced are moving beyond detection to resolution, building agentic capabilities that can auto-route exceptions, correct allocation mismatches, and escalate only what truly requires human judgment. These are targeted solutions solving specific, well-defined problems, and that specificity is their advantage.

Full-stack AI platforms represent the most ambitious end of the spectrum — companies working across the capital markets technology stack to build domain-specific intelligence that not only surfaces problems but also resolves them autonomously within defined guardrails. Increasingly, these models are no longer static systems trained once and then deployed; they incorporate reinforcement learning to adapt behavior based on feedback and real-world outcomes.

None have reached full scale yet, but the capability is maturing quickly. Some of the early adoption signals are counterintuitive: one firm with trillions in AUM is running a proof-of-concept with a company whose entire engineering team fits in a conference room, because the startup's domain-specific capability outperformed the enterprise solution the firm was already paying for. A bet like that from an institution of that size tells you something about where the capability edge sits right now.

HOW WE THINK ABOUT THE ARCHITECTURE

We think about AI activation as an architectural problem. What follows is the framework we use when evaluating it with firms.

CONTEXTUAL DATA: THE SOURCE OF TRUTH. This is where the problem concentrates for most firms. Portfolios, holdings, transactions, commissions, restrictions, performance records — all of it needs to be unified, normalized, and accessible in something approaching real time. If the real bottleneck is fragmented data, adding a copilot on top doesn't address the constraint. The plumbing comes first.

PLATFORM: THE SYSTEM OF ACTION. Order management, execution, risk, operations. This is where the buy-versus-build-versus-outsource decision tends to concentrate, and the answer depends on where a firm sits. Greenfield firms can move to a single platform with AI embedded from the start. Firms with open APIs and reasonable contract timelines can consolidate incrementally. Firms locked into long-term licenses — a seven-year deal from a COVID-era deployment — can still make targeted improvements within the existing architecture. Which of these three matters more to your firm than the specific platform decision that follows?

REFERENCE AND MARKET DATA: Major data providers are building for AI consumption — structured feeds, clean APIs, machine-readable formats. This layer is resolving relatively quickly because the providers are mature and the investment is incremental rather than architectural.

SIGNAL, AGENTS, AND INTELLIGENCE: WHERE DIFFERENTIATION LIVES. This is where proprietary analytical capability sits. The general principle is simple: buy breadth, build depth. The competitive advantage of a $25 billion fund has never been its order management system. It's the investment process — and the people running it. That's where proprietary investment should concentrate.

WHAT IT LOOKS LIKE WHEN THINGS WORK

A portfolio manager wants to increase European industrial exposure by 200 basis points. Compliance checks run against current holdings, regulatory constraints, and the client's investment policy — all drawn from the same data. The execution algorithm selects counterparties based on the commission wallet, market-impact estimates, PM intent, counterparty-exposure limits, and historical fill rates. Settlement, FX execution, allocation, and reporting flow from that same unified data layer. The PM's judgment drives the decision; everything between judgment and settlement happens automatically.

All the work that currently sits in between — the OMS, the risk engine, three operations people, hours of manual reconciliation — becomes automated, auditable, and continuous. It becomes infrastructure.

That sequence works because every step draws from the same data. If compliance, execution, and settlement each see a different version of current positions — the operating reality at most firms — the process breaks at the first handoff.

That scenario also leaves things out. No firm operates in isolation. Counterparty systems, custodian platforms, regulatory reporting infrastructure, and settlement cycles governed by market utilities all remain outside the control of any single organization. And the technology change alone isn't sufficient: roles, decision rights, and accountability structures built around manual processes don't automatically translate to automated ones. The end state arrives unevenly, process by process, not as a single switch — which means the operating model has to evolve alongside the architecture.

If you read that scenario and think "That's impossible in my current setup," you've identified something worth examining. The destination is becoming technically achievable.

WHAT THIS MEANS

70% deploying AI, 24% focused on revenue-generating workflows, 14% reporting measurable ROI. The distance between headline adoption and operational activation is where the opportunity concentrates. The firms that close that gap first gain compounding advantages. The organizational learning that comes from actually using AI in production workflows becomes the moat, and it accrues to the firms that start earliest.

The pattern has a historical parallel. Electronic trading reshaped equities in the 1990s, and the transition was slower than the evangelists predicted and faster than the incumbents expected. This wave may be more disruptive because the technology is an equalizer. It doesn't require the capital scale that previous infrastructure shifts demanded. A $3 billion

firm with clean, unified data and a modern platform can activate capabilities that a much larger firm carrying 15 years of architectural tech debt cannot. The binding constraint is now architectural, and architectural constraints don't correlate with AUM. They correlate with the technology decisions firms made years ago and their willingness to revisit them now.

Asset managers are historically conservative and resistant to change. They can't simply throw out the old and install the new. That creates a two-sided imperative. Established platform vendors need to open their architectures — build APIs, create integration layers, and enable domain-specific AI to plug in without a 12-month implementation cycle. Emerging AI companies need to meet them halfway by packaging their capabilities, so they fit into existing infrastructure rather than demanding firms rip and replace. Neither side can solve this alone.

The vendors that treat integration as a strategic priority rather than a roadmap item will capture disproportionate value. The AI companies that make themselves easy to adopt within existing stacks will scale faster than those waiting for greenfield opportunities. The asset managers that see results first will be the ones where that convergence happens — whether through full-stack transformation or targeted improvements layered onto existing infrastructure.

For the people running desks — for everyone whose work centers on decisions rather than on the manual processes that bridge systems that should already talk to each other — the operational drag is real, and it's addressable. The gap between what's possible and what most firms are doing today is the sharpest it has ever been.

DJ Adams is an independent advisor focused on digital transformation and agentic workflow implementation for asset managers.

Dan Royal is a retired senior trading executive with 40 years of institutional trading experience, most recently as Global Head of Trading at Janus Henderson.

Glenn Lesko and Nilesh Nanavati are principals at OPCO Advisory, advising across the capital markets fintech ecosystem.