Tabb FORUM
WHERE CAPITAL MARKETS SPEAK

October 05, 2025

The State of AI in the Capital Markets 2025: Beyond the Hype, from 'Cool to Core'

Sean Creamer
Contributing Writer
Follow | Profile | More

Share

Artificial intelligence is reshaping capital markets, but unevenly. Firms have pushed past proofs of concept into live deployments across the front, middle, and back office. But data silos, governance hurdles, and hallucination risks remain persistent obstacles. In this TabbFORUM report, Contributing Writer Sean Creamer explores how AI is moving from cool novelty to core infrastructure, what's working in production today, and where the next wave of adoption will determine lasting competitive advantage.

A year ago, the first edition of this report highlighted a capital markets industry on the brink of transformation. Artificial intelligence had begun to move beyond proofs of concept, edging into production workflows across the front, middle, and back office. But adoption was uneven, and practical challenges. Data fragmentation and hallucinations remained serious roadblocks.

Today, that transformation is no longer theoretical. Firms across the ecosystem are piloting AI tools and scaling them. In 2025, the conversation has shifted from “if” to “how fast.” Buy-side leaders are embedding AI across investment workflows, broker-dealers are accelerating surveillance automation, and infrastructure providers are redesigning architectures to support agentic orchestration and retrieval-augmented generation (RAG). Even the regulatory landscape is evolving to meet these capabilities head-on.

While hype persists, practitioners are increasingly skeptical of general-purpose platforms. They want traceability, explainability, and ROI. And they are rapidly developing internal governance models to match. What emerges from this report is a market no longer wondering whether AI can work. It already does. The question now is: Who will get it right, at scale, and with lasting strategic advantage?

1. Buy-Side vs. Sell-Side Approaches to AI

AI adoption is accelerating across capital markets, but firms are charting different paths depending on internal capabilities, institutional posture, and risk tolerance. While the end goal of speed, insight, and automation remain consistent, the strategies for building vs. buying, data governance, staffing, and control vary significantly.

Institutional Build vs Vendor Stack Buying

Firms are reevaluating their technology strategies in light of AI. The build-versus-buy decision depends on internal expertise, data sensitivity, regulatory posture, and speed-to-market needs. Most describe a hybrid approach.

At AllianceBernstein, Chief AI Officer **Andrew Chin** explained the firm’s decision to create its own AI platform: “At the time, we started early on to build something that we call ABAI, said Chin. “We built it because the options weren’t out there. So it’s basically an internal chatbot that allows all our employees to interact with internal documents as well as the web. We built our own AI platform with four walls around it, to make sure that none of our data gets leaked externally”

For institutions like J.P. Morgan Asset Management, a hybrid model prevails. “The number of vendors and the suppliers that we will be using will probably be less than, say, the hedge funds. That said, our vendors tend to be a much more kind of enterprise level,” noted **Kristian West**, Head of Investment Platform.

Dillon Edwards, AI Strategist at J.P. Morgan Asset Management, further explained that firms with complex workflows and proprietary data are moving beyond off-the-shelf models. “We can’t simply just take an off-the-shelf large language model and think that it’s good enough for our users,” he said. “We have a lot of internal proprietary data coupled with vendor data sets that need to be stitched into the right model and the right interface.” The focus now is on data engineering, ensuring the right data powers the right model, invisibly embedded into daily workflows.

> Others, like **Eugene Grinberg** of SOLVE, a capital markets vendor focused on quote aggregation and AI-driven pricing intelligence, stressed the benefits of relying on external vendors, especially for firms without large internal engineering teams. “I think a lot of firms eventually realize that a lot of the vendor ecosystem is much better positioned to provide these types of solutions… especially firms that don’t have large budgets to invest in specialized talent or infrastructure to build these themselves.” For those firms, partnering with vendors is often the most practical path to competitive parity.

Even smaller units are leveraging scale. **Cesar Orosco**, Head of U.S. Alpha Investments within the Quantitative Equity Group at Vanguard explained, “We didn’t do it alone. That’s the beauty and the advantage of being part of Vanguard,” said Orosco. “QEG is about 30 people, but we have the backing of Vanguard’s full back office. We have a dedicated risk team. We have a dedicated technology team. The AI work was done by QEG with collaboration of AI engineers within Vanguard”

Vendor reliance comes with tradeoffs, especially around control, oversight, and regulatory posture. Orosco warned that external models may sacrifice long-term governance. “You can also access languages that are developed by others, but that forces you to relinquish some sort of control,” he said. “More importantly, because things might be modified through time that can create challenges for people who are just outsourcing things.”

At J.P. Morgan Asset Management, the preferred strategy blends internal development with highly customized use. West explained that embedded models within proprietary tools are delivering the most value. “The bit in the middle, the embedded, is probably the most impactful for us because it’s part of something else, and it just makes it better,” he said. “That doesn’t necessarily lend itself to a vendor solution. A lot of that is proprietary.” This approach allows the firm to retain full control over data, performance, and compliance, while tailoring AI capabilities to the nuances of its investment workflows.

Across firms, the trend is to build proprietary capabilities where differentiation matters most, while using vendors to accelerate standard processes.

Control, Governance, and Proprietary Data Access

With AI now embedded in workflows across all business lines, data governance has become foundational. Firms are emphasizing ownership, internal policy, and the ability to control what data is used and how.

> At CEPRES, the leading private markets data & analytics firm emphasizing data sovereignty and deal-level transparency, CEO **Dr. Daniel Schmidt** envisioned a future where firms dictate model-level permissions: “Each owner of data can decide how much and what data should be fit to other counterparties. So in the future, they also can decide what the AI is able to analyze or not. Since transparency and analysis helps justifying investment allocations and it generates the asset-owners’ trust to increase

> allocations, fund-managers increasingly value the win-win and decide to grant data access via a controlled way."

Inside large institutions, this mindset drives governance frameworks. "It is our data. We need to be very clear on the policies and procedures and make sure that it's really well governed," said **Ashok Krishnan**, Managing Director and Global Head of Electronic Trading and Platform at Bank of America.

Vendors are drawing similar lines. **Jim Kwiatkowski**, CEO of LTX, an AI-powered fixed income trading platform focused on liquidity discovery and data-secure execution, said, "As a Broadridge company, information security is a paramount concern for LTX. We never use client data for training. Client queries into BondGPT use Broadridge's OpenAI enterprise services account, and all data is anonymized before being sent."

Firms like SOLVE use proprietary data as a strategic edge. "We're training our models not just on some of the publicly available data," explained Grinberg, "but on many millions of proprietary quotes that we collect by leveraging other AI techniques to aggregate a tremendous amount of quotes data from unstructured messages between market participants." This approach gives SOLVE a differentiated advantage, enabling it to surface pricing intelligence that other firms, especially those limited to public feeds, may miss.

The takeaway is clear: strong AI governance begins with strong data discipline.

Staffing Models and Team Structure

As firms move toward production-level AI, team structure has become critical. Success now depends on cross-functional teams with both business fluency and technical depth.

At Bank of America, Krishnan described a modular model: "Each team is basically a combination of a couple of people from technology, a couple people from my group, which is platform, and a couple of quants. And we bring them together with a small focus group of SMEs from that area."

Tech firms also emphasize talent. Grinberg from SOLVE explained, "We have a really talented team of AI researchers and engineers that spend a lot of time in feature engineering to make sure that the model is as accurate as possible."

At J.P. Morgan Asset Management, AI isn't new. The firm has invested in a dedicated data science and AI team for nearly a decade. Over the last eight years, Edwards has worked to align advanced tools with the workflows of portfolio managers, sales teams, and operations. JPMorgan Chase allocated $18 billion to technology this year, up from $17 billion last year, and employs more than 63,000 technologists. "You can't necessarily just drop a technologist in the room with an investor and expect them to figure out the right thing to build in a way that would work for our end user," Edwards said. "And so I play that intermediary role and then I'll work with our data scientists, our data engineers, to start to think through how we are going to build that?"

That same model shaped the firm's Smart Monitor system. Edwards and his team interviewed dozens of portfolio managers to map daily routines, data pain points, and unmet needs. One consistent message: they were overwhelmed by data and needed help prioritizing it. Smart Monitor acts as a personalized signal filter, using LLMs and investor-specific inputs to push forward the most relevant companies, events, and research insights. The tool synthesizes broker reports, transcripts, and internal memos against each manager's stated investment thesis. The system identifies what each portfolio manager is monitoring and flags those signals at the right moment.

> What began as technological novelty is now becoming operational necessity. AI, once seen as a shiny tool with niche applications, is being embedded deeper into core investment workflows. "I've been using the phrase *from cool to core* to describe what's been happening in our industry," said AB's Chin. He explained that while GenAI once wowed teams with flashy capabilities, the focus has shifted toward integration. "We're all trying to make AI core to our investment processes, integral to how we run our business." For Chin, that transition means solving real organizational problems, not just experimenting with new tech for its own sake.

TP ICAP built an internal innovation lab to scale experimentation. "We've created a dedicated AI and innovation lab that is effectively the incubator for AI within TP ICAP, helping us to build these new solutions and empower the rest of the organization to take advantage of AI technology," said **Andrew Shannon**, Head of Cloud Infrastructure and Head of AI Adoption at TP ICAP.

Across the board, technical capabilities are being paired with business context to enable scalable execution.

Asset Manager vs. Bank vs. Hedge Fund Posture

Institutional structure drives AI strategy. Regulatory burden, internal resources, and risk posture differ widely across banks, asset managers, and hedge funds.

At J.P. Morgan Asset Management, West explained how compliance requirements shape strategy: "Being a bank-owned asset manager, we have very high expectations and requirements on everything we do. We will be exposed to regulations that many of our peers won't be covered by."

These constraints shape everything from vendor evaluation to rollout speed. Grinberg reinforced the point, noting wide variation in AI maturity: "The gap between the most advanced firms compared to the firms that are falling behind is quite significant, both on the buy and the sell side. But that gap is probably wider on the buy side."

In short, market structure and culture continue to influence how, and how fast, AI is being adopted across the industry

2. AI Use Cases by Office Function

AI is no longer confined to R&D labs or sandbox environments. In capital markets, it has moved decisively into front, middle, and back-office operations. From personalized trading tools to reconciliation engines and compliance copilots, firms are embedding AI into real-world workflows. But progress varies as some are scaling hundreds of use cases, while others are still laying governance and infrastructure foundations.

A. Front Office

Earnings, Sentiment, and the AI Research Edge

Natural language tools are reshaping how analysts evaluate potential investments. At AllianceBernstein, Chin described how assessing a new company often involves weeks of manual effort, gathering descriptions, analyzing suppliers and customers, and understanding the competitive landscape. “That’s a process where these tools can have a lot of impact,” he said, “because it can help pull the data together, the untapped, unstructured data sets, and synthesize information.” By automating the labor-intensive steps of early-stage company research, GenAI has the potential to dramatically compress time-to-insight.

> This speed-to-context advantage is now being embedded directly into trading platforms. At LTX, AI is being placed directly into the trader workflow. BondGPT Intelligence anticipates user needs and surfaces relevant answers automatically. “It’s built into the trading platform in a way that it’s anticipating the questions that users might ask,” said Kwiatkowski. “It provides answers depending on where they are in the application without them needing to know how to use the application or even type in the question.”

Buy-side firms are building proprietary research assistants that use internal data and retrieval-augmented generation (RAG) to deliver firm-specific insights. “We’re hearing about how buy-side firms are using GenAI to generate fairly accurate recommendations for reallocation to specific sectors or even stocks months before they moved,” said **Irene Galperin**, Senior Advisor at InterSystems, a data infrastructure provider enabling real-time analytics and AI capabilities. These assistants, she added, become truly proprietary “when they’re leveraging the firm’s existing research and data.” By grounding outputs in internal sources, firms retain analytical edge without dilution, delivering insights that are accurate, contextual, and uniquely their own.

In the private markets, research acceleration offers outsized benefits due to the scarcity of structured data. At CEPRES, Dr. Schmidt explained how his team optimizes its AI models on granular, asset-level datasets to replicate the work of a live investment consultant to enable deeper portfolio analysis and smarter allocation strategies. “It’s not a simple AI, it’s a full investment consultant who helps to make investment decisions, to analyze portfolios smarter,” he said. “We can deploy investment market-data back to 1970’s on that AI from the development step and grow it smarter than others who have not had data access and just build a shelf and feed the data afterwards in.”

Even as AI matures, firms remain cautious about crossing into fully autonomous decision-making. **Alex Schlesinger**, Director of AI Product Management at SS&C, a provider of AI-driven decision-support tools, highlighted the regulatory sensitivities: “Trade recommendations get pretty tricky from a regulatory standpoint. Where is our place as a company, as a product provider, in offering the recommendations?” SS&C draws a firm line that empowers human decision-makers without assuming their responsibilities.

AI is being used to sharpen and accelerate investment insight, with firms firmly retaining control over decision-making.

Trading and Execution (AI Trade Suggestions, Co-Pilot Tools)

AI is advancing fastest in trading domains still reliant on manual effort and voice execution. While equity markets are largely electronic, fixed income remains fragmented and opaque, making it ripe for transformation.

> **Mark Katz**, CTO for Financial Services at Hitachi Vantara, a data infrastructure and analytics firm supporting AI adoption in trading and compliance, pointed to fixed income as a rapidly evolving frontier. Unlike equities, which he described as “well understood and well studied,”AI usage in fixed income trading is “growing by leaps and bounds” as firms apply AI to the “dark art of predicting long and short term rates and when those yield curves cross.” By training models on a hundred years of market data, Katz explained, firms can detect patterns that suggest an inflection point may be near, helping AI tools “divine some sense of that.”

Residual mismatches from block trades, especially in dark pools, have long posed a latency challenge for brokers. TP ICAP developed an AI-driven tool to address this gap by analyzing leftover orders in real time and identifying likely counterparties based on historical trading behavior. “We wrote an AI program which effectively was able to look at those residuals in real time and then match them,” said Shannon of TP ICAP. What once took hours is now executed in minutes.

As credit markets evolve, with greater electronic trading and new liquidity sources like ETFs and portfolio trades, traders are managing more complex data streams. LTX’s Kwiatkowski noted that AI is arriving at a critical moment: “These new tools and the use of AI will help the traders to consume more information to make more timely and informed decisions. They come at exactly the right time to help them as this market’s evolving.”

Automation is reshaping quoting behavior, blurring traditional lines between dealer and client. As Grinberg of SOLVE explained, “There’s a lot more systematic quoting across the buy and sell side. The more advanced firms are already relying on a combination of people, data, automation, and algos, and they’re trading much higher volumes than in past years.” This shift is driving greater liquidity, transparency, and tighter bid-ask spreads.

Together, these tools are reducing latency, improving pattern recognition, and beginning to operate in a more autonomous capacity, though human oversight still plays a key role.

Client Personalization and Engagement Tools

Client personalization and engagement tools in capital markets have rapidly evolved from flashy virtual assistants into deeply embedded systems that anticipate user needs and support high-value interactions. As firms move past the early generative AI hype, they are beginning to focus on integrating intelligence into real-time workflows, especially for sales and trading professionals.

The novelty phase of GenAI is giving way to more utilitarian applications, with virtual assistants now embedded into nearly every public interface. “Every public interface has a virtual assistant at this point,” said Katz at Hitachi Vantara, reflecting on how firms are moving beyond surface-level experimentation toward tools that drive real operational value.

This shift from novelty to function reflects a growing preference for AI systems that deliver value in precise, contextual moments. At Bank of America, Krishnan shared an example of how AI is being used to streamline sales meeting preparation. “We are trying to create, at the touch of a button, a document that sweeps in every possible piece of information that a salesperson would potentially need before he or she walks into that client meeting.” Rather than simply offering assistance, the tool consolidates key data to enable more impactful client conversations.

That same principle is being applied on the trading floor. Personalization has taken the form of intelligent nudges and targeted insights delivered during the execution process, according to Kwiatkowski from LTX. “Traders on the platform are receiving proactive insights based on where they are in their workflow, ensuring they get the information that matters most to them exactly when they need it.” Whether preparing for a meeting or executing a trade, these AI-powered touchpoints are increasingly designed to support human action.

B. Middle Office

Reconciliation Automation

Reconciliation remains a persistent source of operational friction, particularly in trade break resolution, document matching, and exception handling. AI is reducing manual effort and unlocking faster resolutions.

> Machine learning is beginning to ease the complexity of reconciliation by reducing the need for hand-coded rules across highly variable client environments. “One of the tricky things with reconciliation is that there’s enough variance where it’s hard for us as a product provider to just say, ‘we can resolve these breaks this way for everybody,’ because that’s just not how it works,” said Schlesinger at SS&C. “But getting somebody to sit down and write all the rules, that’s pretty laborious.” By layering AI on top of

> machine learning, firms can automate more of the matching logic without sacrificing flexibility.

In post-trade operations, the impact of AI is increasingly centered on speed and efficiency. “Reconciliation, settlements, the focus is going to be on the efficiency side,” said Chin at AllianceBernstein. “How do you reconcile better, how do you settle better?” While AI can help reduce error rates, the greater value lies in accelerating these back-office processes and minimizing latency.

BNP Paribas is one of the largest custodians and fund administrators in the world, and its Securities Services division has moved to operationalize AI at scale. **Jeffrey Zoller**, Director and Head of US Fund Services at BNP Paribas Securities Services, explained how the firm has embedded intelligent document processing into core workflows to streamline reconciliation and reduce manual burden across its global client base. “We also leverage intelligent document processing tools across several functions within Securities Services in order to read and extract key information from documents, add controls and simplify reconciliation within core systems. As a result, we have realized a reduction in manual processing volumes, an improvement in the quality of data received into our systems, and improved efficiency and productivity,” he said.

Pricing Engines

In fixed income and credit markets, AI is enabling firms to anticipate market shifts and respond dynamically. These tools ingest real-time and historical data to inform pricing strategies.

SOLVE began its predictive pricing rollout in municipal bonds, where data fragmentation and pricing opacity presented an immediate opportunity. “Step one was munis,” said Grinberg. “Step two is corporate bond predictive pricing, which we launched in June, and enhanced in September with greater coverage and improved accuracy.” The firm plans to extend the model across fixed-income asset classes using its platform’s daily intake of millions of quotes.

By fusing proprietary and traditional data, firms like SOLVE are modernizing price discovery across fragmented credit markets.

Corporate Actions Processing

Corporate actions processing has long been one of the most manual and error-prone areas of post-trade operations, especially when dealing with multilingual data from global issuers.

At BNP Paribas, that challenge is being met with internally developed AI tools. “An in-house translation tool into corporate action processing allows us to process corporate action notices with speed and accuracy across more than 30 source languages,” said Zoller at BNP Paribas Securities Services. “This capability performed nearly 45 million translations in 2024.”

Data quality is proving the biggest bottleneck for scaling AI. “The data that you train your models on is not Oracle databases and structured data. It’s files, objects, streams of social media feeds, and in some large firms, billions of objects,” said Katz of Hitachi Vantara. He noted that firms like Zataris, Hitachi’s Pentaho division, and Data Dynamics are emerging to help categorize unstructured streams, tag sensitive fields like NPI or PII, and make sense of sprawling data estates.

In QEG at Vanguard, AI scaled quickly because it was built on strong infrastructure. “We have a data strategy that has been consistently applied for a long time, so the AI is just an extension of that,” said Orosco. “We began the research in 2020, greenlighted implementation in 2022, and within six to nine months had a full MLOps capability.”

Other firms are experimenting with workflow transformation. At TP ICAP, AI was used to extract 60 key metadata fields from tens of thousands of legacy contracts, improving downstream procurement. At CEPRES, Dr. Schmidt described internal systems for AI-driven data structuring and verification: “The reporting of the future is an AI, you ask your questions into the AI, and the AI answers.”

C. Back Office

AI is reshaping the operational backbone of capital markets firms, streamlining everything from regulatory response and compliance to document processing and internal knowledge access. Where traditional automation was rules-based and brittle, today’s AI systems prioritize explainability, risk alignment, and tight integration with upstream workflows.

Regulatory Reporting

High-pressure regulatory reviews often require rapid retrieval and reconciliation of vast data sets, a process that can overwhelm internal teams. When audit timelines shrink to hours, AI offers a path to faster response with less manual strain.

“When a regulator or an auditor comes calling, they often make requests with extremely short turnarounds measured in hours,” said Katz at Hitachi Vantara. “Gathering that data is a huge effort internally for the firms. But if you’ve got an AI trained in those areas, you can produce that very quickly without a lot of human effort.”

BNP Paribas is piloting AI models to enhance anomaly detection in post-trade workflows. “Pilot use-cases are underway in several areas, including the detection of unusual or exceptional transactions against expected behavior,” said Zoller at BNP Paribas Securities Services. These models are designed to flag attention points based on historical data patterns that may have gone unnoticed by traditional rule-based systems. To ensure these AI tools meet regulatory expectations, Zoller emphasized the need to “extend core control frameworks to the AI capabilities, including model validation, testing and monitoring.”

Compliance Automation (KYC, AML, eComms)

Compliance remains one of the most natural fits for AI in the back office, as it helps in reducing false positives, identifying true risks, and accelerating reviews.

Machine learning is already improving the efficiency of anti-money laundering (AML) and sanctions workflows. “The industry standard [false positive rate] sits around 95%,” noted **Tod McKenna**, Head of AI Client Engagement & Solutions at Citi Services. “Our focus is on improving the true positive detection rate. By leveraging better data collection, enhanced labeling, and machine learning, we’re aiming to increase accuracy and efficiency, interconnect financial crime disciplines, and improve detection of complex and evolving financial crime patterns.”

Beyond flagging alerts, AI is starting to power deeper diagnostic insight. “A.I. is being used in compliance to just do what a human would do, but there’s a much more exciting aspect where we can take machine learning to understand why there are frequent compliance breaks,” said Schlesinger at SS&C.

Emerging use cases in financial services are less about generating text and more about enabling action, according to Grinberg from SOLVE. “They’re not looking for language model generative things. They’re looking for action, fixed income trading or KYC are the sorts of use cases we discussed. AI moves beyond machine learning’s pattern matching to take systemic, systems-based action without necessarily producing a verbal output,” he said. “This is basically next-gen machine learning, and these are the use cases that are actually demonstrating ROI now.”

Strong governance remains non-negotiable, especially at regulated entities. “At LTX, as a registered broker-dealer operating within a comprehensive compliance framework, we viewed compliance as a critical requirement from day one and we knew our clients felt the same way,” said Kwiatkowski of LTX.

RAG, Audit Trails, and Hallucinations

As AI models increasingly influence financial decisions, explainability and auditability have become non-negotiable. To meet regulatory and internal governance demands, firms are embedding mechanisms like Retrieval-augmented generation (RAG) to ensure AIs can only pull from certain databases, and audit trail tools to demonstrate how predictions are made and how reliable they are.

At SOLVE, infrastructure was built with transparency in mind. “We designed our machine learning model to produce a mathematical artifact that essentially allows us to compute the confidence score, calibrated on a scale from one to 10, so that our clients, or their models, are able to ascertain how confident the model was for a particular prediction,” said Grinberg.

CEPRES is applying AI at a broader scale to support macro-level risk assessment—particularly for public-private investment programs that demand real-time accountability. “AI can do that much more complex analysis, much faster, and help decision makers to

understand the risk at the end," said Dr. Schmidt. He noted that governments want forward-looking insight into economic outcomes, not just retrospective reports.

Robust governance is essential to help mitigate against hallucinations from slipping into production. At Citi Services, this starts with traceability. "AI components and their outputs are meticulously logged and tracked," said McKenna. "This rigorous documentation helps us identify and address instances of fabricated or misleading content produced by LLMs."

Edwards from J.P. Morgan Asset Management highlighted the challenges posed by hallucination and inconsistency in model behavior. "The problem with some of these upgrades was the same prompts in different versions yield a slightly different output," he explained. "That difference raises uncertainty in the mind of our users." To mitigate these risks, his team takes a conservative approach: "We stress to the models, do not look any further."

Retrieval-augmented generation (RAG) is now a cornerstone for hallucination mitigation, but its value depends entirely on provenance. Without clear sourcing, grounded outputs are indistinguishable from guesses. Katz at Hitachi Vantara, stressed that trust begins with upstream controls: "Do you know what data was fed into this? Have you taken care to ensure that NPI and PII data was not part of what fed a public-facing model?" he asked.

At LTX, citation transparency is built into the product. "We call it 'Show your work,'" said Kwiatkowski. "It's not enough to be compliant and accurate. You actually have to demonstrate your accuracy with this explainability." The patented feature walks users through how results were generated, what data was gathered, how it was processed, and how answers were derived, to build trust in real-time outputs.

3. Human-in-the-Loop vs. Full Automation

AI is transforming operations across capital markets, but firms are not relinquishing human control. Even as automation scales, human supervision remains a critical layer, especially in areas involving compliance, risk, or investment judgment. Firms are actively defining boundaries between where machines can act alone and where oversight is mandatory.

Where Firms Draw the Line Today

AI is transforming operations across capital markets, but firms are not ceding control. While automation has improved efficiency across functions like reconciliation, compliance, and reporting, there is clear consensus that human oversight remains indispensable where accuracy, accountability, and regulatory exposure are at stake.

In compliance and financial reporting, firms are drawing hard boundaries around autonomy. McKenna at Citi Services, emphasized that AI-generated output also passes through human verification. "We firmly believe in a 'human-in-the-loop' philosophy, where human verification of AI-generated outputs is a crucial control point," he said. "This is vital

in areas like compliance and financial reporting, where accuracy and accountability are critically important."

Operational functions like reconciliation follow similar logic. SS&C's Schlesinger described how the firm avoids full delegation of decision-making by integrating human approval into the workflow. "The last thing we want is for the AI to make a decision in its entirety," he said. "You have a prompt that says, here's what I'm about to do and you can approve it. The AI takes care of the rest."

TP ICAP has hardcoded this philosophy into its automation design. AI is allowed to propose actions, but not to act. "We're not letting the AI just go off and do its own thing," said Shannon. "We're creating the outcomes for a human to check, and then the person has to have accountability."

This view of human review as an active triage function, not a passive sign-off, is echoed elsewhere. Galperin of InterSystems pointed out that exceptions, edge cases, and low-confidence outputs still require domain expertise to resolve. "There will be some exceptions," she noted. "This is where a human can jump in and quickly take action with precision."

Even as models improve and use cases scale, firms are embedding human controls into every AI-enabled workflow. Across compliance, operations, and reporting, human-in-the-loop systems are not a temporary guardrail so much as they are the default architecture.

What Can and Can't be Automated

Firms are embracing automation when it clearly reduces workload, but the threshold for full autonomy depends on risk. At AllianceBernstein, Chin applies a scale of consequence to decide where AI can act independently and where it cannot. "If I'm making an investment decision, I do not want it to be autonomous," he explained. In contrast, for tasks like summarizing daily emails, the stakes are low enough to permit full delegation.

That same logic is shaping deployment strategies elsewhere. At TP ICAP, Shannon noted that automation is beginning in the back office, where processes tend to be more structured and less risky. "We're starting off in the back office, but over time, we'll see [AI] in all sorts of different areas," he said. These tools, in his view, function like virtual humans, able to extend productivity without displacing core expertise.

Rather than pursue full replacement of human roles, firms are using automation to relieve teams of low-value or repetitive tasks. That frees up time for higher-order responsibilities involving discretion, compliance, or client outcomes. All areas where the cost of error remains too high to entrust to machines.

PMs and Compliance Analysts Still In Control

Even as AI systems mature, final accountability still belongs to the human expert. Across capital markets, there's clear consensus: portfolio managers, compliance analysts, and operations leads must remain responsible not only for decisions but also for understanding how those decisions are made.

Human oversight at Vanguard is embedded across every phase of the modeling process, from initial data validation to final trade execution. Orosco emphasized that this begins well before deployment: “Every step of the way, we need to make sure that we are properly vetting, validating the data. not including forward-looking biases,” he said. To eliminate lookahead bias, the team-built infrastructure allows models to “walk in time,” training only on data available at each historical point.

Even once a model produces trade recommendations, a human remains in control. “That trade list is also going to be vetted by a portfolio manager,” Orosco explained, “not because the portfolio manager is going to override the bet of the model, but because we’re constantly making sure that the implementation process, the portfolio construction process, the trade list generation, all of that is aligned and that the process is reflecting our views.”

The same dynamic is at play in compliance. Kwiatkowski of LTX stressed that “accuracy is job number one,” but it’s not enough to simply be right. “You actually have to demonstrate your accuracy with this explainability so that users have confidence relying on the results to make critical decisions,” he said. Even features like dealer selection scoring, which streamline workflows, leave ultimate discretion to the human: The patented scoring technology uses data science to help identify the optimal subset of dealers for an RFQ, while leaving the final counterparty selection entirely with the trader.”

At AllianceBernstein, Chin framed the future of investment roles through the lens of what he calls the “Iron-person” persona, a skilled subject matter expert equipped with what he called “AI armor.” “First and foremost, you have a human who is accountable for the decision at the end of the day,” he said. That individual must “know what risk management means or what equity research means,” and possess the raw talent to perform the job. But they also need to use the tools well. “This person needs to use the armor,” Chin explained, meaning they must be capable of prompting, tuning, and directing AI effectively to sharpen their work.

Whether in investment strategy, compliance, or operations, the message is consistent: AI can reshape the process, but responsibility and control stay with the human.

4. Data Architecture & Infrastructure

As generative AI moves deeper into production, data infrastructure is emerging as a key competitive edge. Firms are upgrading pipelines, shifting to cloud-native or hybrid environments, and implementing retrieval-augmented generation (RAG) frameworks that demand precise data tagging, lineage tracking, and latency control. To scale responsibly, institutions are investing in governance overlays, smart data fabrics, and low-code tools that make AI both efficient and auditable. The objective is clear: create infrastructure that supports rapid AI deployment without sacrificing risk controls or performance.

Migration to Cloud or Smart Data Fabrics

Cloud migration and modern data architectures are reshaping AI readiness across the capital markets. At Citi, a wholesale infrastructure change has underpinned the goal of achieving the safe and ethical scaling of AI. “We’ve been making substantial investments to build a robust data technology structure that is designed for the safe, ethical, and

effective use of AI across the firm," said McKenna, citing the firm's proprietary AI gateway as a key element for architectural governance and compliance enforcement.

A similar transformation is underway at TP ICAP. Shannon led the firm's migration to AWS and the launch of an AI innovation lab, accelerating solution delivery across the enterprise. "The speed in which we were able to deliver these solutions was considerably faster than traditional software engineering," he said. "We've moved from early prototypes to 22 AI solutions now being used within the firm, and that was only possible once we moved to the AWS cloud."

> At InterSystems, the trend is moving beyond simple cloud lift-and-shift. Galperin explained how their smart data fabric architecture embeds AI directly into the data layer, eliminating the need for wholesale migrations. "Instead of modernizing by moving all of your data into a new cloud-based data lake or warehouse, which often just creates another silo, a smart data fabric connects to data at the source," she said. The approach allows firms to unify legacy systems, cloud platforms, and new lakehouses into a non-disruptive architecture that enables AI use cases without replatforming.

Katz at Hitachi sees an inflection point: "It's going to be very difficult to outfit any more infrastructure on-prem." Yet financial firms remain wary of sending sensitive data to public cloud hyperscalers, creating friction between performance goals and regulatory obligations.

Infrastructure cost control is a key concern for SOLVE. Grinberg pointed to two core challenges: elasticity and runaway costs. "The first challenge is architecting the system so that it could support auto-scaling, given our growing coverage and the ability to interface with our models via APIs, but maintaining reasonable costs in the cloud becomes difficult," he said. "So the second challenge is optimizing our processes so that our costs don't scale with our volumes of data."

Security, Power Usage, and Latency Concerns

AI performance gains come with serious infrastructure and security tradeoffs. Katz at Hitachi noted that most data centers are underprepared for today's GPU power demands: "No data center was engineered for the power consumption of current computing needs people are running partially filled data centers," because traditional setups can't handle GPU-intensive racks.

That strain on infrastructure is mirrored by rising security threats. Shannon at TP ICAP warned of the risk of increasingly realistic deepfake attacks, where impersonation over phone or video could bypass trust safeguards. "You could hit a scenario where you're being called by your boss, but actually, it's not them," he said, urging firms to train employees the same way they did around phishing.

Mitigating these risks requires infrastructure that enforces data governance by design. At SS&C, Schlesinger described how their internal AI gateway supports redaction, logging, and PII filtering to ensure safe handling of client information. "We have the assurance that we are handling our clients' data in the safest manner possible."

5. AI Governance, Regulatory Trends & Risk

As AI becomes embedded in capital markets infrastructure, regulatory scrutiny is intensifying. Firms are focused on how much visibility, explainability, and oversight is enough. From cross-functional risk reviews to real-time model tracing and jurisdiction-specific compliance strategies, a strong governance posture is now the baseline.

Governance Frameworks and Regulatory Divergence

At SS&C, Schlesinger pointed to regional divergence as a complicating factor for firms trying to comply with developing U.S. regulations versus more clearly defined frameworks abroad. “There’s due caution because some of the regulations, particularly in the U.S., are in development. They are very regionalized and in some cases misunderstood,” he said. “In Europe, it’s a little more straightforward. They’ve provided some useful checklists. It’s a paperwork exercise at that point. But on the U.S. side, with explicit AI laws still undefined at the federal-level and there’s only a handful at state-level, there’s caution because people are unsure which guidelines may become law.”

Citi, meanwhile, has taken a proactive approach to internal governance. Tod McKenna described the creation of a dedicated AI Risk Management Committee to help ensure consistent oversight. “Over the past year, we’ve made significant progress in advancing our internal AI governance structure, firmly embedding artificial intelligence as a core tenet of our transformation efforts,” he said. “This committee provides enterprise-wide oversight for all AI-related risks, helping to ensure a consistent, responsible, and client-centric approach across Citi’s diverse operations and external engagements.”

Global Data Privacy Adherence (GDPR, etc.)

Complying with global data privacy and AI regulations requires region-specific deployments. “We are not able to deploy the same AI systems as in each region due to different regulations, AI solutions will need to be configured differently than we can do for our clients in Europe, the Middle East, in Asia, or in the US,” said Dr. Schmidt of CEPRES. “We cannot say, just because in Europe maybe strong regulations will become applied, that we cannot provide the strongest technology to the global market. In other countries, we have different rules, and so we can deliver it according to the rules of the specific country.”

InterSystems takes a client-by-client approach, mapping AI access controls to internal governance. “Every firm has their own unique internal organizational structures and their own data governance framework that dictates who can see which data,” said **Galperin**. “If, based on your role, you can only access market data, you will only access market data. The same applies to PII data. If it needs to be scrubbed for any reason, whether for internal use or for regulatory reporting, then we make sure that the steps are in place to anonymize the data.”

At BNP Paribas, Zoller positioned data privacy within broader AI governance. “Firms must ensure that their data management strategy and capabilities are aligned to their AI objectives,” he said. “Risks around data privacy, accuracy and bias highlight the need to establish robust governance and oversight.” This emphasis on data privacy as a core AI risk reinforces the need for firms to treat governance as a foundational element of AI system design and deployment.

To navigate global compliance demands, TP ICAP has built internal coordination into its AI oversight. “We have created an AI governance framework to make sure that we’re operating in a safe way within the kind of regulations around the world that we have to adhere to,” said Shannon. “We have different people across the business that are involved and help us assess those types of projects.”

XAI (Explainable AI) Design Goals

As generative models gain traction, transparency is no longer optional. Firms must ensure accuracy and demonstrate it to regulators, clients, and internal teams.

Today’s black-box models still leave institutions guessing how outputs are generated, an unacceptable risk in regulated environments. Katz at Hitachi Vantara, noted that “you have an input and you get an output and nobody really knows how that output derived from that input.” To address this, firms are increasingly embracing Explainable AI and Responsible AI (RAI) to enforce traceability. These frameworks aim to create “a breadcrumb trail or some telemetry or observability about how you arrived here,” Katz explained, along with standardized tests to detect bias in model behavior.

Transparency also plays a critical role in honoring fiduciary obligations to clients, according to AB’s Chin. “If there’s a prompt that says, ‘Hey, what is the appropriate asset allocation for a particular type of investor?’ we have to be careful how we answer that.” Understanding investor profiles, preferences, and eligibility requirements is essential to ensure responsible output.

6. Agentic AI and Autonomous Systems

As capital markets move beyond process automation and into orchestration, agentic AI is becoming the scaffolding for intelligent workflows. These systems interpret intent, interface across domains, and execute autonomously within guardrails. While adoption is early, foundational patterns are emerging across investment, operations, and infrastructure teams.

Work Orchestration Agents

Agentic AI is best understood not as a singular model, but as a system of task-specific agents coordinated under a central layer. It is a network of AI agents, composed into a network of neural networks, overall, it takes a network of networks.

J.P. Morgan Asset Management’s Spectrum IQ platform is designed as a unified access point for firmwide data. “If you ask a question about a particular account, it will pull together operational data, investment data and sales data,” said West. “That is the start of the agentic framework, an orchestration that sits across those agents that can do things and respond to your requests and instructions.

Schlesinger of SS&C pointed to the release of Anthropic’s MCP standard as a tipping point for orchestration. “It really opened the door for us to work with more data sources and partners, because now we as a collective AI community have the standard to build

against," he said. "Agentic lets us truly use the best tools for the job rather than being stuck trying to shoehorn a bunch of features into one LLM."

Chin of AllianceBernstein raised the importance of inter-agent oversight, especially as firms scale model count. "If I get a stock analysis back, how do I make sure it's the right one? I need an agent that challenges that first agent," he said, noting that second-order validation will be critical as complexity grows.

Autonomous vs. Co-Pilot Boundaries

Agentic AI handles repeatable tasks well, but exceptions still require human oversight. "The more repeatable they are, the easier it is for agentic AI to perform them," said Galperin of InterSystems. "But that human in the loop will be a true subject matter expert for triaging situations where there are anomalies." In cases like trade reconciliation, most steps can be automated, but data mismatches or validation issues still need a human to step in and resolve them quickly.

At Bank of America, Krishnan described how an agentic framework has collapsed latency between operations and trading. "You come in on a Monday, there's a new bond that's not showing up," he said. "In the past, ops or tech would have to set it up. Now, with an agentic framework, we click a button and the bond is tradeable in two minutes, when it used to take hours." This change allows agents to span traditional silos, accelerating tasks once gated by functional separation.

Governance and Overreliance Warnings

Unchecked autonomy could create new forms of systemic risk, warned Grinberg of SOLVE. "There will probably be another Knight Capital," he said, "not for the same reason, but because of this massive overreliance on algorithms. That will probably force other firms to rethink how they should be using humans and AI as co-pilots. It will just bring a lot more scrutiny into it, which isn't a bad thing"

While the full vision of agentic AI remains aspirational, these early deployments are shaping the contours of what's possible, and what's still essential to control.

7. Market Impact and Alpha Opportunities

Firms across the ecosystem are beginning to harness AI's pattern-recognition capabilities not just for operational lift, but for alpha generation. From uncovering small-cap opportunities to cross-asset optimization and cleaner signal extraction, the search for a predictive edge is increasingly AI-enabled.

Small-cap discovery through AI monitoring

At J.P. Morgan Asset Management, early thinking around tools like Smart Monitor suggests that AI could help uncover opportunities among lesser-known companies. "I'm hoping that some of the smaller companies at the tail end of the investment universe will be monitored, and might start to differentiate our portfolios over time," said Edwards. By systematizing the detection of growth indicators and risks across this less-covered

segment, J.P. Morgan Asset Management hopes to sharpen its view on emerging investment opportunities.

Signal extraction from noisy/unstructured sources

Discovery is only useful if signals can be separated from noise, a core challenge in investment AI. “My first priority was can we use these tools to make better forecasts, to create better models?” said AB’s Chin. But he acknowledged the difficulty: “The signal-to-noise ratio in the industry is very low. There’s a lot of noisy data out there.” With non-stationary data and frequent market regime shifts, traditional models struggle. Still, the search for cleaner signals, especially across unstructured sources, remains central to Chin’s mandate.

Enriched analytics are also shaping post-trade workflows. “There are multiple data sources, both external and internal pricing, ratings, liquidity data, counterparty selection data, holdings, issuer filings recent trades, TRACE,” said Kwiatkowski from LTX. “These new capabilities have the potential right now to make users better at selecting bonds, better at portfolio optimization, and better at counterparty selection.” As AI tools surface insights from diverse datasets, pre-trade decision-making is becoming faster, more informed, and increasingly strategic.

Multi-Asset Strategy and Portfolio Optimization

AI is helping reshape how firms design cross-asset portfolios and tailor strategies to investor needs. “Construct the most efficacious model and look at performance over time,” said Katz at Hitachi Vantara, describing user-directed optimization based on inputs like risk tolerance or investment horizon. Portfolio optimization, he added, is “definitely one” of the most active alpha-seeking use cases on the buy side.

Dr. Schmidt of CEPRES also pointed to AI’s role in expanding the opportunity set itself. As private markets become more data-driven, Schmidt believes AI will guide capital toward less traditional assets. “AI will help market capital [be] guided into the private markets, non-traded markets. That wouldn’t be possible without AI since the complexity to solve for multiple targets can now be superiorly handled.” As transparency improves and models evolve, AI may become a key driver in unlocking capital flows into previously overlooked or capital-limited segments of the market.

Conclusion

Capital markets firms have moved past the era of AI demos and slide decks. They’re deploying scalable tools, training internal teams, and embedding AI into the fabric of operations. Across reconciliation, pricing, compliance, and portfolio design, AI is increasingly part of the daily workflow as a systemic enabler.

But the path forward remains steep. Fragmented data, unclear governance, and organizational misalignment can derail many projects. AI is not a silver bullet. But for the capital markets, it is becoming a strategic differentiator. The next wave of winners will be not only those who adopt AI, but embed it in their operations and workflows

• • • •

Sean Creamer *is a writer focusing on AI, financial, and marketing technology. He's worked as a journalist for HFMTechnology, Wall Street Letter, and Operational Risk and Regulation. As a marketing tech writer, he penned stories for EMARKETER. In addition to writing about real estate, Sean is an outdoor enthusiast who manages publicity for the Kayak and Canoe Club of New York, a whitewater enthusiast club.*

Photo Credit: ChatGPT generated image

CEPRES

www.cepres.com

CEPRES

CEPRES is the leader in private markets technology, trusted by institutional investors and asset managers to make smarter, faster decisions across the entire deal lifecycle. From fundraising and due diligence to portfolio monitoring and risk management, CEPRES delivers AI-powered insights driven by the industry's Gold Standard Data – verified, granular intelligence sourced directly from GPs, never scraped or estimated. With a foundation of trusted data from 16,500+ funds and 143,000+ deals, CEPRES transforms private market complexity into actionable intelligence, making it the partner of choice for investment professionals worldwide.

Dr. Daniel Schmidt
Founder & CEO at CEPRES
CEPRES

LTx

www.ltxtrading.com

LTX is an e-trading platform that helps corporate bond market participants trade smarter, combining powerful artificial intelligence with innovative trading protocols to improve liquidity, efficiency, and execution.

Jim Kwiatkowski
Chief Executive Officer
LTX, a Broadridge Company

Hitachi Vantara

www.hitachivantara.com

Hitachi Vantara powers innovation with resilient data infrastructure and real-world solutions that drive business outcomes. We help customers solve digital challenges and move from now to next by leveraging unmatched industrial and digital expertise. Trusted by over 80% of the Fortune 100, we unlock new revenue streams, enhance competitiveness, reduce costs, and improve customer experiences, all while delivering social and environmental impact.

Mark Katz

Client Strategy and Technology Officer for Financial Services

Hitachi Vantara

InterSystems

www.InterSystems.com/Financial

Established in 1978, InterSystems is the leading provider of next-generation solutions for enterprise digital transformations in the finance, healthcare, manufacturing, and supply chain sectors, including production applications at most of the top global banks. Its cloud-first data platforms solve interoperability, speed, and scalability problems for large organizations around the globe. InterSystems is committed to excellence through its award-winning, 24×7 support for customers and partners in more than 80 countries. Privately held and headquartered in Cambridge, Massachusetts, InterSystems has 36 offices in 25 countries worldwide. For more information, please visit **InterSystems.com/Financial**

Irene Galperin
Senior Advisor, Financial Services
InterSystems

SOLVE

www.solvefixedincome.com

Hundreds of the world's most sophisticated financial firms rely on SOLVE fixed income data, analytics, and tools to transform the way they bring new securities to market, trade on secondary markets, and value highly illiquid securities. Both buy-side and sell-side participants trust the AI-driven SOLVE Market Data Platform to provide unparalleled market transparency, reduce risk, and save hundreds of hours across front-office workflows. Founded in 2011, SOLVE has developed the largest datasets of real-time bids, offers, and market color across Securitized Products, Municipal Bonds, Corporate Bonds, Syndicated Bank Loans, Convertible Bonds, CDS, and Private Credit. With locations in North and South America, Europe, India, the Caribbean, and APAC, SOLVE is the leading provider of market pricing in fixed-income markets. The SOLVE team has deep industry expertise and a passion for financial markets and developing innovative technology. More information about SOLVE can be found at **https://solvefixedincome.com**

Eugene Grinberg
Chief Executive Officer
SOLVE

SS&C Advent

www.advent.com

SS&C Advent helps over 4,300 investment firms in more than 50 countries—from established global institutions to small start-up practices—to grow their businesses, minimize risk and thrive. We have been delivering unparalleled precision and ahead-of-the-curve solutions for more than 30 years, working together with our clients to help shape the future of investment management. To learn more about Advent's solutions and services, visit advent.com or contact **info@advent.com**

Alex Schlesinger

Director, AI Product Management

SS&C Advent