TradingTech Insight
From A-TEAMGROUP

The Buy-side Trading Desk of the Future:

How Multi-Asset Convergence and the Agentic Overlay Are Reshaping Institutional Trading

Commissioned by:

LSEG

Table of Contents

Executive Summary

The buy-side trading desk is facing a structural reckoning. Decades of siloed systems — separate teams, tools, platforms and protocols for equities, FX, and fixed income trading — have imposed a compounding operational burden that COOs can no longer justify. Call it the silo tax: duplicated headcount, fragmented data, and systems that resist the joined-up workflows AI requires.

The good news is that 30 years of market evolution have quietly built the antidote. Every asset class now shares common core workflows — similar processes for monitoring markets, sourcing liquidity, pre-trade analysis, and executing via algo or RFQ. As firms look to deploy artificial intelligence in their trading workflows, this functional convergence is emerging as the foundation stone from which AI will take the trading process to the next level.

The marketplace has come full-circle: trading began with voice on the floor, evolved into highly structured deal tickets delivered electronically, and is now devolving back to natural language — this time powered by natural language process (NLP) and AI. An agentic overlay, sitting on top of a clean, unified infrastructure, translates a trader's chat instruction into the precise structured workflow beneath. The silo tax's days are numbered.

Firms that rationalise their technology stacks now — by establishing common layers across asset classes — will be well positioned to deploy the agentic overlay and capture the efficiency and performance gains ahead. Those that delay will find the silo tax compounding.

I. Why the Buy-Side Trading Desk Is Ready for AI

AI has made early inroads on the trading desk in research and idea generation, where the absence of heavy infrastructure constraints makes GenAI copilots and large language models (LLMs) relatively straightforward to deploy. But the execution workflow — governed by regulation, compliance, and deeply entrenched connectivity — has proved more resistant.

Things are changing, however. Practitioners are realising that AI's moment in execution is approaching, not because AI has become more powerful, but because the market infrastructure beneath it has finally converged. AI is not the primary driver of the desk's impending transformation; it is a tool that can now thrive because common workflows and functionality have emerged across asset classes after three decades of fragmented development.

The vehicle for that transformation is the agentic overlay: an AI layer that sits on top of unified, multi-asset workflows and translates natural language instructions into the structured execution processes beneath. A trader types a chat instruction — "buy $50m USDMXN spot, minimise market impact" — and the agentic overlay autonomously determines whether to route via RFQ or algorithm, selects the appropriate venue, and captures the full post-trade record. The underlying infrastructure is unchanged; only the interface has evolved.

II. The Great Convergence: 30 Years in the Making

Trading has evolved in three distinct phases. In the first, all asset classes operated on floors and in physical pits, mediated entirely by voice. In the second, migrating upstairs during the 1980s, voice gave way to electronic messaging, then to highly structured deal tickets as vendors added precision execution and post-trade data capture. In the third — now underway — the industry is devolving back to natural language, but with AI translating intent into the structured workflows underneath.

This evolution was not uniform. Between 1995 and 2025, equities, FX, and fixed income each adopted different protocols at different speeds. Equities led in algorithms — VWAP, implementation shortfall, POV — and Transaction Cost Analysis. Fixed income built its liquidity model around RFQ and axes. FX gravitated toward streaming prices on multi- and single-dealer platforms.

Post-2008, 'Big Reg' — Dodd-Frank, MiFID II — fostered cross-pollination. RFQ migrated to less-liquid equities. Central clearing extended from listed futures to OTC derivatives. FIX, originally an equities protocol, became the universal messaging standard. The result: every asset class now runs on the same functional playbook.

SIDEBAR: Lessons from History — Three Standards That Crossed Asset-Class Boundaries

FIX Protocol. Developed by Salomon Brothers and Fidelity in the 1990s to standardise equity order messaging, FIX now defines order-routing and execution across FX, fixed income, exchange-traded derivatives, and commodities. Its standardised message types — NewOrderSingle, ExecutionReport — enabled straight-through processing and opened the multi-asset OMS/EMS market.

Central Counterparty (CCP) Clearing. The CCP model — placing a clearing house between buyer and seller to manage counterparty risk — began in listed futures. Dodd-Frank and EMIR mandated its extension to OTC derivatives post-2008. Interest-rate swaps, credit default swaps, and other instruments now operate within a standardised post-trade infrastructure, reducing bilateral risk across derivatives markets.

Request for Quote (RFQ). Originating in fixed income, where institutional investors requested competing dealer quotes for less liquid securities, RFQ has been adopted by platforms across asset classes for products including FX options, interest-rate swaps, structured products, and listed options block trades. It bridged the gap between voice trading and order-book markets, creating electronic price discovery in previously opaque segments.

The lesson: when a convention solves a genuine market problem, it travels. The agentic overlay is poised to be the next such convention.

The Silo Tax

Despite this convergence at the workflow level, many firms retain asset-class-specific teams and systems — a structural inheritance that now constitutes a direct drag on operating performance. Traders across equities, FX, and fixed income perform identical core functions: monitor markets, source liquidity, conduct pre-trade analysis, execute, and evaluate post-trade execution costs and quality. Yet asset managers maintain separate headcount, separate systems, separate operational budgets for each.

COOs increasingly frame this as the silo tax: the cumulative cost of duplication that was justifiable when market structures genuinely differed but is now a strategic liability. The business case for consolidation — common OMS/EMS infrastructure, unified data, integrated teams — has become compelling. A 100% poll result at a recent A-Team Group trading technology conference, where all practitioners expected full consolidation into integrated multi-asset trading architectures within the year, reflects how widely that case is now accepted.

Eliminating the silo tax is the COO mandate. The agentic overlay is the mechanism for executing it.

III. Drivers of Change: Business, Technology, and Regulation

Three converging forces are accelerating the desk's transformation.

Business imperative: The silo tax is no longer defensible when workflows are functionally identical across asset classes. Common OMS/EMS and integrated OEMS platforms now exist to support the rationalisation COOs require. Earlier technology cycles — the shift to algorithms in equities — demonstrated that redeploying traders from routine execution to high-value roles improves both cost efficiency and performance quality.

Technology maturity: Workflow convergence across asset classes provides the clean, unified foundation that AI agents require. A fragmented, asset-specific technology stack cannot support an agentic overlay; a common layer can. LSEG's platforms are architected on precisely this premise — providing the joined-up multi-asset infrastructure upon which AI tools can operate reliably.

Regulatory environment: Regulation has historically driven standardisation — FIX adoption, OTC clearing mandates — and is likely to do so again as AI in execution moves from experimental to mainstream. Firms that invest in compliant, auditable, unified infrastructure today are building the regulatory resilience that AI deployment will eventually require.

IV. AI in Trading: Inevitable but Not Immediate

The trading desk is regulation-dense and infrastructure-heavy by design. OMS, EMS, and OEMS platforms are deeply connected to regulated execution and investment workflows — they cannot be replicated by a generic LLM in minutes. This complexity functions as a moat: it protects incumbents and prevents unqualified AI tools from disrupting execution without proper integration.

Another recent A-Team Group conference poll crystallises the market's position: one-third of practitioners expected AI to take over core trading functions within five years; two-thirds said it would never fully happen, citing regulatory constraints and the irreducible value of human expertise. Both camps are right, in different senses. The complete displacement of trading infrastructure is implausible. The deployment of an AI layer on top of that infrastructure is not only feasible but is already underway.

Trading platforms today are characterised by fixed connectivity to thousands of market participants — relationships built over years that no AI model can replicate on the fly. This sunk infrastructure investment is a structural advantage for established providers. It also means that change will proceed as an agentic overlay atop incumbent systems, not as a wholesale replacement.

The timeline is a question of when, not whether. Change will be incremental: humans moving up the value chain toward deep analysis and high-value trade ideas while AI handles routine execution. This trajectory is feasible within five years, and likely sooner.

V. The Agentic Overlay in Practice

The agentic overlay is not a theoretical construct. It is a definable architecture: an AI agent layer that accepts natural language input from the trader, interprets intent, and routes instructions into the structured workflows beneath — selecting protocols, venues, and counterparties without requiring the trader to navigate asset-class-specific forms or ticket structures.

In practical terms, a trader issues a single chat instruction spanning multiple asset classes. The agentic overlay determines the appropriate execution method — algo or RFQ, based on liquidity profile, urgency, and cost — selects the venue, and processes the instruction through the underlying workflow. The trader never fills a structured ticket; the AI does, invisibly, using the precision the infrastructure

demands. Each firm then has the option to configure how much the 'human trader in the loop' is required based on instrument, size, and market conditions. As execution heads become comfortable with agentic workflows, they can 'loosen the reins' by relaxing parameters, thus driving ever-increasing efficiency and productivity.

This is the full-circle story of trading's interface evolution. Voice on the floor was natural and intuitive but imprecise. Structured electronic tickets were precise but rigid, demanding traders adapt to machine logic. NLP-powered chat is both: intuitive for humans, precise for machines. The AI translates between the two.

LSEG's multi-asset infrastructure, for example — spanning OMS/EMS capabilities, pre-trade analytics, TCA, and liquidity sourcing tools — provides the structured layer on which agentic functionality can be deployed. Clean, joined-up workflows are not merely a benefit, they are a prerequisite. Firms still operating siloed, asset-specific stacks will find the agentic overlay cannot function reliably on fragmented foundations.

VI. Where AI Adds Value Across the Trade Lifecycle

AI's entry points into the trade lifecycle are sequential, progressing from information-light to infrastructure-heavy functions as integration matures.

Idea generation and research: The most immediate value, and the most widely adopted today. AI copilots and LLMs accelerate analyst and portfolio manager workflows with minimal infrastructure dependency — querying market data, summarising research, generating trade hypotheses at speed and scale.

Pre-trade analysis: Natural language replaces rigid TCA forms. A trader asking for specific execution strategies for a $50m position receives scenario analysis drawn from structured underlying data, without manually configuring asset-class-specific parameters. The agentic overlay handles the translation.

Execution support: The same workflow convergence that enables common pre-trade analysis enables AI to recommend execution protocols, venues, and algorithms in execution. The agent identifies whether an algo or RFQ is optimal — based on liquidity, spread, urgency, and counterparty axes — by operating across the unified multi-asset workflow layer.

Post-trade and analytics: Common infrastructure enables AI to ingest and analyse data from the entire trade lifecycle, 24/7. Cross-asset rationalisation and performance attribution — impossible when data was trapped in asset-specific silos — becomes routine. The silo tax is eliminated across the board.

VII. The Trading Desk of the Future

According to this vision, the buy-side trading desk of the future is asset-agnostic, AI-augmented, and voice-and-chat-first. Its defining characteristic is not the technology deployed but the architecture beneath it: common, consistent workflows across equities, FX, and fixed income that allow an agentic overlay to function reliably.

Structurally, highly manual market-specific desks will consolidate into multi-asset human teams supported by AI agents handling routine execution. Human traders will focus on oversight, exception management, high-value analytical and relationship-based work that AI cannot replicate: deep cross-asset strategy, relationship management with both customers and the sell side, and complex trade structuring.

The trader's interface will be natural language. The underlying execution will be structured, regulated, and precise. This combination — intuitive input, controlled output — is what the agentic overlay uniquely enables, and what 30 years of market infrastructure evolution has made possible.

VIII. Preparing for 2026 and Beyond

The silo tax is the problem. The agentic overlay is the solution. Common infrastructure is the prerequisite.

For COOs, the immediate priority is establishing a capability matrix: auditing whether foundational functions — pre-trade analysis, liquidity sourcing, execution, post-trade data capture — are consistent across all products. Where gaps exist, they represent both operational risk and barriers to AI deployment.

Once common layers are in place, the agentic overlay can be positioned on top: a natural language interface that expedites workflows, simplifies multi-asset execution, and begins systematically eliminating the silo tax. Firms that establish these foundations now will be best positioned as AI capability in execution accelerates over the next three to five years.

Trading has come full circle — from voice, to structured electronic tickets, to voice again. The desk of the future speaks naturally and executes precisely. The infrastructure to support it already exists. The question is whether firms choose to use it.

Introducing LSEG

LSEG is one of the world's leading providers of financial markets infrastructure and delivers financial data, analytics, news and index products to more than 44,000 customers in over 170 countries.

From Insight to Execution

Our next-generation data and analytics workflow solution, LSEG Workspace, offers a combination of cutting-edge technology with market-leading content to help you save time, reduce errors and ultimately get things done through a seamless experience.

Beneath this interface lies LSEG's execution stack, bringing together LSEG TORA and REDI on Workspace for execution management, alongside LSEG's portfolio management capabilities as the portfolio backbone. Together, they deliver consistent, multi-asset workflows across pre-trade, execution and post-trade processes. This architecture enables you to move seamlessly from insight to execution within a unified operating model, providing the foundation for the multi-asset convergence and AI-driven workflows described in this paper.

Setting the Standard for Trusted AI in Finance

Through the three pillars of our AI Strategy, LSEG Everywhere, we are setting the standard for trusted AI in finance.

- **Trusted data** – our open, partnership approach fuels AI in financial services with world-leading, high-quality, accurate data
- **Transformative products** – delivering new AI products that enable our customers to work faster, smarter and with confidence
- **Intelligent enterprise** – using AI to drive efficiency, innovation and growth across our operations

We help organisations fund innovation, manage risk and create jobs by partnering with customers at every point in the trade lifecycle: from informing their pre-trade decisions and executing trades to raising capital, clearing and optimisation.

Backed by more than three centuries of experience, innovative technologies and a team of over 26,000 people in 65 countries, we are driving financial stability, empowering economies and enabling you to grow sustainably.

Find out more: lseg.com

A-TEAMGROUP

For Marketing: **www.a-teamgroup.com**

For Content: **www.a-teaminsight.com**