databricks

Modernizing the Trade Lifecycle in Capital Markets

From fragmented data
to AI-driven intelligence

Contents

databricks

SECTION 1

Why the status quo is no longer an option

Capital markets firms have spent decades layering point solutions on top of aging infrastructure. The result is familiar: siloed data warehouses, custom-built grids and fragile spreadsheets that slow every stage of the trade lifecycle, from alpha generation to settlement.

The pressure to change is no longer gradual. It's acute.

Trading and volatility
Intraday swings remain elevated as geopolitics, AI-linked themes and policy uncertainty drive sharp repricings. AI is now large enough to influence GDP, earnings and credit markets, not just technology indices, with nearly $3 trillion in AI-related data-center investment projected through 2028 and more than 80% of that spend still ahead.

Compressed timelines
The move to T+1 in North America and pilots toward T+0 in other markets are squeezing operational windows that were already tight. European authorities are simultaneously advancing work on shortening settlement cycles, putting additional pressure on post-trade infrastructure.

AI as a macro variable
Major sell-side research desks now treat AI as a macro driver, not a niche theme, a sustained industrial build-out that can support growth, but also trigger sharp valuation resets when expectations shift.

Systemic risk focus
Central banks and standard-setters are explicit: AI and digital finance can increase the intensity, speed and complexity of risk transmission, especially when firms converge on shared models, data sources and third-party platforms.

Regulatory convergence on AI and trading
Supervisors are sharpening expectations around algorithmic trading and AI-enabled execution, with new guidance on pre-trade controls, governance and testing frameworks.

Capital markets executives are navigating real vendor noise from foundation model providers, cloud hyperscalers and financial data terminals all claiming to offer a complete AI solution.

The distinction that matters: large language models are powerful, but blind to your firm's proprietary order flow, internal research and risk positions. Cloud providers require stitching together separate services across storage, compute, orchestration and governance. And terminal extensions, however sophisticated, cannot reason across your internal systems. The competitive edge in capital markets AI is not in the model. It is in the proprietary data layer, and that is precisely where most firms remain unprepared.

Across global strategy firms and analyst houses, the conclusion is consistent: firms that defer modernization of their data and AI foundations will face compounding disadvantages in speed, cost, resilience and regulatory readiness.

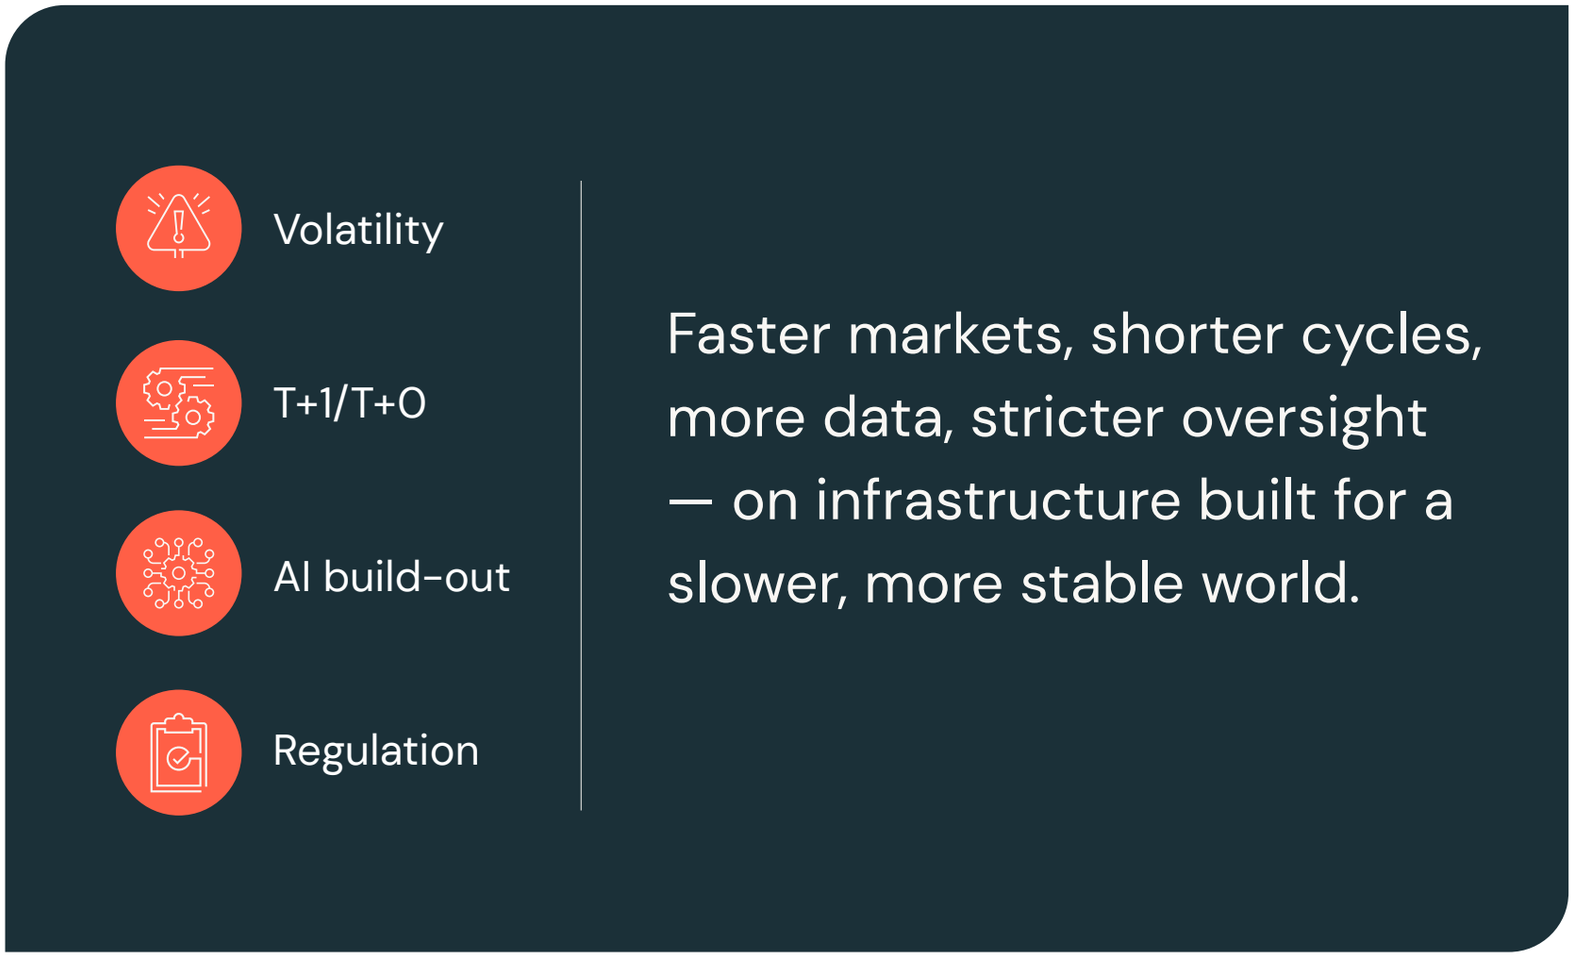

databricks

SECTION 2

The data and AI forces reshaping capital markets

The alternative data and provider shift

The last five years have transformed not just what data capital markets consume, but how they consume it.

A 2026 survey of 100 portfolio managers and analysts (managing roughly $2 trillion in assets) found that:

94% already use AI and/or machine learning in their alternative data research process.

71% say combining data from different sources is the most frustrating workflow challenge.

43% cite data evaluation as the hardest stage.

57% believe employment data will provide an outsized informational edge in the near future.

The 2026 Alternative Data Market – Insights & Trends from Portfolio Managers and Analysts (2026).

In parallel, traditional terminals remain critical, but come with a growing cost and complexity burden. Financial data platforms typically cost $10,000 to $30,000+ per user annually, with pricing generally increasing alongside platform complexity, breadth of data and depth of analytics. Higher-cost offerings tend to support more advanced workflows, broader market coverage and tighter enterprise integration.

At the same time, industry analysts suggest that AI-powered, API-first and open data platforms could capture a meaningful share of the terminal and data-licensing market over the next decade, particularly where regulators push for greater data transparency and interoperability.

The net effect is clear:

- **Access is no longer the bottleneck.** Most firms have more data and more providers than ever before.
- **Integration, governance and time-to-insight are.** Without a unified data platform, each new dataset or provider increases complexity, cost and operational risk.

AI as both edge and risk

Leading institutions and governing bodies are converging on a dual message: AI is simultaneously the biggest source of competitive edge and a new channel for systemic risk.

Morgan Stanley describes AI as a "macro variable" that now influences GDP, earnings, credit spreads and geopolitics. Its research estimates nearly $2.9 trillion in data-center construction through 2028, with AI-related investment contributing roughly 25% of expected US GDP growth in the near term. Early adopters with measurable AI impact are already seeing cash-flow margin expansion at about 2× the global average.

The World Economic Forum, in an article authored by DTCC's CEO, argues that AI will "fundamentally reshape all aspects of business and financial markets" and highlights post-trade as a key area where AI can shift infrastructure from reactive to predictive. DTCC's own GenAI-powered Risk Calculator assistant has achieved around 97% accuracy with no observed hallucinations post-launch, significantly compressing the time clients spend on manual documentation review.

The Bank for International Settlements (BIS) warns that AI and digital finance can alter market functioning, liquidity and risk propagation. Its 2026 speech on AI and digital finance emphasizes that these technologies increase the intensity of shocks (through concentration and shared dependencies), the speed of transmission (shorter decision and settlement cycles) and the complexity of interconnections, making risks harder to observe and manage.

The Financial Stability Board (FSB) has elevated digital innovation and AI to a core pillar of its 2026 work program, with explicit plans to publish sound practices for AI adoption, use and innovation, alongside its ongoing monitoring of vulnerabilities and non-bank intermediation.

ESMA, in its 2026 supervisory briefing on algorithmic trading, calls out "the extended use of artificial intelligence in algorithmic trading" and sets expectations for pre-trade controls, governance arrangements, testing frameworks and outsourcing oversight for AI-enabled trading systems.

Taken together, these perspectives crystallize the challenge: AI requires a data and governance foundation strong enough to deliver competitive edge without undermining resilience.

OPPORTUNITY

- Alpha and signals
- Execution quality
- Predictive surveillance

RISK

- Model herding & procyclicality
- Operational dependencies
- Opaque decisioning
- Regulatory scrutiny
- Intraday risk

The deciding factor is not the model alone, but the data architecture and governance behind it.

SECTION 3

What “modernized” looks like: the Databricks Data + AI Platform

Modernizing the trade lifecycle is less about adding a new analytics tool and more about changing the underlying assumptions of the architecture. The most valuable data in capital markets is proprietary: order flow, internal research, client relationships, risk model outputs. The firms winning with AI are the ones who figured out how to bring that data off the sidelines and into a unified environment alongside their market data. That combination is where the edge lives.

The Databricks Data + AI Platform is built for this shift:

- A lakehouse foundation (Delta Lake) that unifies real-time and historical data, market, reference, alternative and risk, in a single, governed platform.
- Unity Catalog for centralized governance, granular access controls and end-to-end lineage across data, models, dashboards and AI applications.
- Agent Bricks for building agentic applications, from research copilots to trading analytics and surveillance assistants, with retrieval-augmented generation (RAG), model serving and evaluation.
- AI/BI Dashboards and Databricks SQL that allow PMs, traders, sales and risk teams to ask questions in natural language or via familiar dashboards, without leaving a governed environment.
- Serverless compute and workflows that simplify running large-scale risk, analytics and reporting workloads without maintaining complex infrastructure.

Instead of separate stacks for research, trading, risk, finance and compliance, firms get a single, governed platform that supports end-to-end trade lifecycle intelligence.

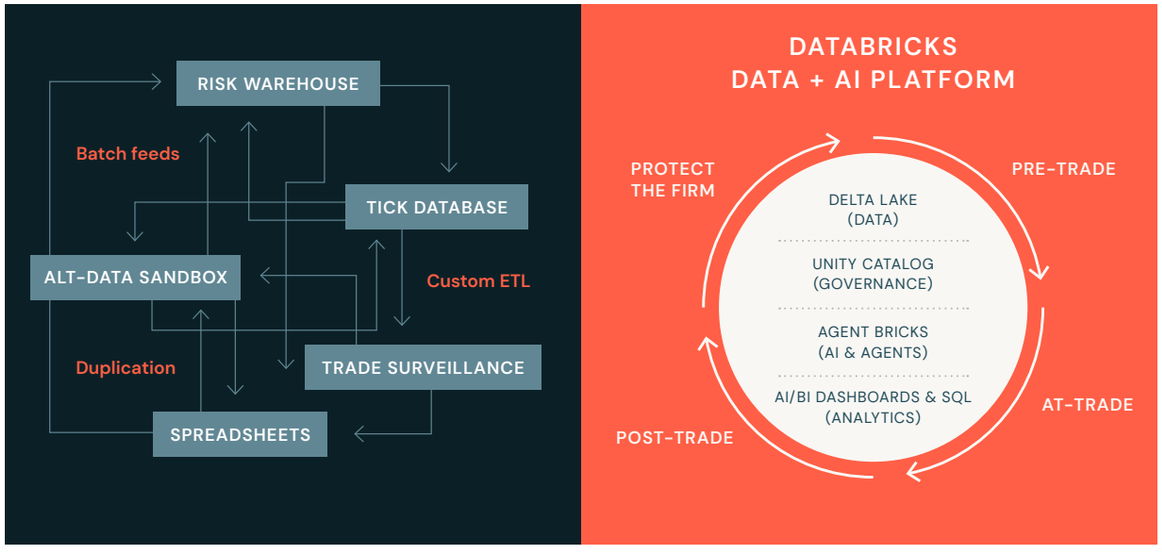

SECTION 4

Modern use cases across the trade lifecycle

Pre-trade: turning data abundance into alpha

CHALLENGE

Quant and fundamental teams ingest more data than ever including market data, fundamentals, ESG scores, news, research, transactions, geospatial and alternative data. The blockers are unification, governance and experimentation speed, not access.

Exabel's 2026 survey underscores the operational pain: 71% of managers say combining data from different sources is their most frustrating alternative-data challenge, and 43% cite data evaluation as the hardest step.

WHAT DATABRICKS ENABLES

On the Databricks Data + AI Platform, firms can:

- Consolidate market, fundamental, ESG and alternative data into a single governed lakehouse, integrating licensed feeds (e.g., FactSet, S&P, Deutsche Börse) with proprietary and alt-data listings available via Databricks Marketplace.
- Use Agent Bricks and RAG to build research assistants that summarize earnings transcripts and filings, answer analyst questions in natural language, and propose relevant peers, events or factors to explore further.
- Run signal research and factor backtests with experimentation with Databricks-managed MLflow, ensuring reproducibility and version control from notebook to production.
- Expose insights to PMs, sales and trading desks via AI/BI Dashboards, governed end-to-end by Unity Catalog.

CUSTOMER SPOTLIGHT: FRANKLIN TEMPLETON

Franklin Templeton built **SIGNALS**, an AI-powered investment analysis platform on Databricks Agent Bricks, to scale portfolio insights and fund research:

- Expanded AI-generated investment commentary from **200 manually covered products to hundreds**.
- **2+ hours saved per analyst per week**, reclaiming more than **15 collective analyst hours weekly**.
- **$15M in new product opportunities** identified during the early rollout by enabling faster portfolio comparisons and fund recommendations.

By grounding AI agents in proprietary fund scoring models and enterprise documents governed by Unity Catalog, Franklin Templeton delivers compliant, analyst-quality portfolio insights in seconds, helping financial professionals evaluate investments and identify opportunities faster.

At-trade: smarter execution in real time

CHALLENGE

Execution quality is increasingly defined by technology. Smart order routing, transaction cost analysis (TCA) and intraday liquidity management all depend on processing tick-level data at scale, combining it with historical context and surfacing useful signals in near real time.

ESMA's 2026 supervisory briefing on algorithmic trading reinforces that firms must have robust pre-trade controls, governance and testing in place for algorithmic and AI-driven execution systems. That raises the bar for both performance and transparency.

WHAT DATABRICKS ENABLES

On Databricks, trading and execution teams can:

- Ingest and process real-time tick and order-book data into Delta Lake, alongside multi-year historical data for backtesting.
- Run TCA and execution analytics that blend venue, strategy, client and market context to highlight slippage drivers and best-execution performance.
- Support algo research and deployment within one system — from prototyping and backtesting in notebooks to model serving and monitoring in production.
- Give desk heads and risk managers intraday views of P&L, flow and risk, accessible via dashboards or conversational interfaces.

CUSTOMER SPOTLIGHT: OPTIVER

Real-Time Trading Dashboards with Databricks Apps and Dash

In trading, speed and insight go hand in hand. At the 2025 Databricks Data + AI Summit, Optiver shared how they built low-latency, self-serve dashboards by integrating Databricks Apps with Dash. The result: real-time market visibility that scales with both data and decision-making.

From Fragmented Systems to Live Insights at Scale

In the session "Real-Time Market Insights — Powering Optiver's Live Trading Dashboard with Databricks Apps and Dash," attendees got an inside look at how they:

- Replaced siloed, on-prem systems with a unified, petabyte-scale Databricks Data + AI Platform
- Cut end-to-end latency from minutes to seconds through Spark and Structured Streaming optimization
- Empowered traders to build and iterate on dashboards without engineering support
- Solved scale and reliability challenges with smart caching, version control, and modular dashboard generation

By hosting Dash apps directly on Databricks, they eliminated infrastructure overhead while enforcing fine-grained access controls through Unity Catalog. The result is a dashboarding ecosystem that balances trader autonomy with engineering performance — built for impact at speed and scale.

Post-trade: risk and reporting that scale

CHALLENGE

Post-trade is where technical debt and regulatory pressure collide: VaR and xVA running on overnight batches, reconciliations glued together by spreadsheets, and regulatory reporting processes that cannot easily adapt to new rules or intraday expectations.

The BIS warns that as AI and tokenization shorten decision and settlement cycles, the speed at which stress propagates increases, leaving less time for human or supervisory intervention. That makes modern, scalable post-trade infrastructure not simply a cost issue, but a stability issue.

WHAT DATABRICKS ENABLES

On Databricks, firms can:

- Run VaR, stress testing and scenario analysis on unified, high-fidelity data with elastic compute, reducing calculation times while improving transparency and auditability.
- Automate NAV, P&L attribution and reconciliations, replacing networks of spreadsheets and manual checks with governed pipelines.
- Standardize and simplify regulatory and internal reporting by feeding all functions (risk, finance, treasury, compliance) from the same curated datasets.
- Use AI to detect anomalies in trade, position, collateral and cash data, surfacing potential breaks before they escalate.

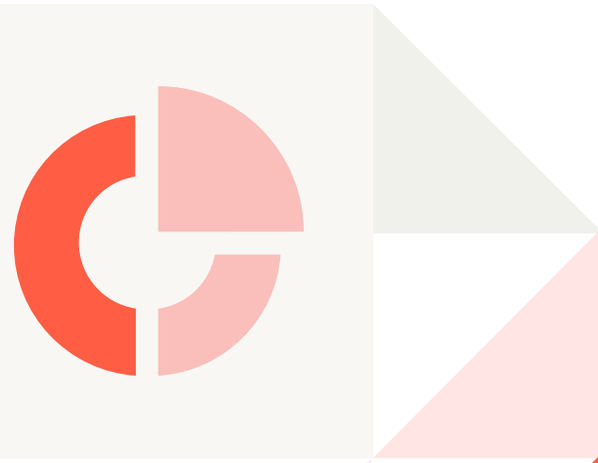

Protecting the firm: surveillance, fraud and systemic resilience

CHALLENGE

Surveillance, fraud prevention and conduct teams must monitor increasingly complex activity spanning:

- Trades and orders across venues.
- Communications and collaboration tools.
- Reference data, sanctions lists and beneficial-ownership information.
- Market, news and even social-media signals.

At the same time, global bodies such as the FSB and BIS are examining how AI might affect liquidity, interconnectedness and procyclicality, emphasizing the need for robust governance and operational resilience.

Stage	Primary users	Example questions AI can answer	Databricks outcomes
Pre-trade	Research, PMs, quants	"How did spreads move around last earnings surprise?"	Faster research, richer signals, explainable model
At-trade	Trading, execution, sales	"Where is execution cost diverging from expectations today?"	Intraday TCA, smarter routing, liquidity analytics
Post-trade	Risk, finance, operations	"What is our VaR under Scenario 3 by desk and region?"	Faster, transparent risk and reporting
Protect	Compliance, surveillance, AML	"Which traders show unusual patterns over the past month?"	Fewer false positives, better coverage, auditability

WHAT DATABRICKS ENABLES

On Databricks, firms can:

- Combine streaming transaction data (trades and orders) with communications, reference data and alternative signals, including listings such as FactSet sanctions data, Deutsche Börse flow insights, S&P Global Commodity Insights and Stocktwits sentiment available via Databricks Marketplace.
- Train machine-learning and anomaly-detection models on historical patterns and deploy them into production with full lineage and audit trails.
- Provide surveillance teams with conversational tools that summarize unusual patterns and generate case narratives, while still allowing analysts to drill into raw data when needed.
- Demonstrate transparent AI governance to internal model risk teams and external supervisors, with consistent controls over data access, versioning and evaluation.

This aligns with the World Economic Forum's view that AI can help post-trade infrastructure, including central utilities such as DTCC, move from reactive processes to predictive resilience, but only if organizations invest in clean, well-governed data and robust oversight.

SECTION 5

The governance layer that makes AI trustworthy

External research is clear: AI without governance is a risk multiplier.

- The BIS highlights that shared AI models and platforms can increase correlated behavior, intensify liquidity strains and compress the time window for corrective action, while making risks harder to observe due to model opacity and complex dependencies.
- The FSB's 2026 work program includes a dedicated report on sound practices for AI adoption, use and innovation, emphasizing governance, oversight and cross-border consistency.
- ESMA's algorithmic-trading briefing sets expectations for governance, testing and controls around AI-enabled execution systems, with an explicit focus on avoiding new technology risks.
- The World Economic Forum / DTCC article stresses that successful AI deployments require data that is "highly available, discoverable, clean and clearly labeled," backed by strong data governance and security.

The Databricks governance stack is designed to meet these expectations:

- Unity Catalog provides a single control plane for access, lineage and auditing across data and AI assets.
- Business semantics in Unity Catalog allow institutions to define key measures — such as VaR, expected shortfall, liquidity buffers or exposure metrics — once and reuse them consistently across reports, dashboards and applications.
- MLflow on Databricks adds structured model governance, including versioning, experiment tracking, evaluation and controlled deployment into production.
- Comprehensive logging captures who ran what queries or AI interactions on which data, supporting model risk management, internal audit and supervisory reviews.

This is the difference between experimental AI and AI that can stand up to board-level and regulatory scrutiny.

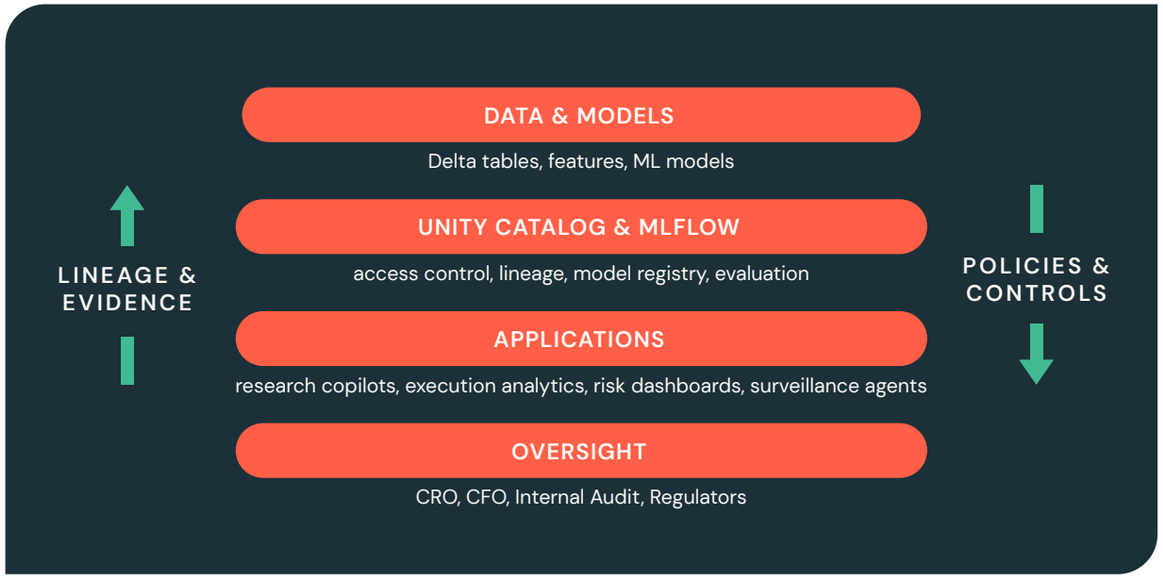

SECTION 6

How to get started, in a way that lasts

Modernizing the trade lifecycle does not require a "big bang" rewrite. The firms moving fastest tend to follow a practical pattern.

Start from the conversations that matter

Identify recurring questions that matter to your leadership:

- *Heads of trading and execution:* "Which strategies and venues contributed most to execution slippage this week?"
- *CRO and risk committees:* "How would a specific volatility or rate shock affect our VaR and liquidity profile by desk and region?"
- *COOs and heads of operations:* "Where are we seeing the highest break and fail rates by asset class?"

These questions define the first workspaces and AI applications to build.

Unify the minimum viable data

Use Databricks to bring together the critical datasets behind those conversations:

- Market, reference and alt-data (including data-provider listings on Databricks Marketplace).
- Positions, trades, orders, collateral and risk factors.

Register them in Unity Catalog, define ownership and apply consistent semantics so that different teams can work from one version of the truth.

Deliver two or three lighthouse workflows

Rather than trying to modernize every function at once, focus on a small set of high-impact, visible workflows, such as:

- A research copilot for one priority asset class.
- An execution analytics cockpit for one major desk.
- A trade surveillance uplift in one business line or region.

Measure improvements in:

- Cycle time (from question to answer).
- User coverage (how many front-office or risk users are self-serving).
- Risk/operations metrics (e.g., fewer false positives, fewer breaks).

Build governance into the rollout

Treat governance and evaluation as accelerators, not obstacles:

- Reuse existing access policies so AI and analytics cannot bypass entitlements.
- Turn on lineage and detailed logging from day one.
- Define simple evaluation criteria for AI applications including accuracy targets, escalation thresholds, review cadences, all aligned to your model risk framework.

This approach keeps you aligned with the direction set by the FSB, BIS and regional supervisors while avoiding costly retrofitting later.

Scale by domain, not by tool

Once the initial workflows demonstrate value, expand by domain including more desks, asset classes, regions or risk categories, using the same Databricks platform, semantics and governance.

This avoids the trap of adding yet another siloed “AI tool” and instead builds a durable data and AI fabric for capital markets.

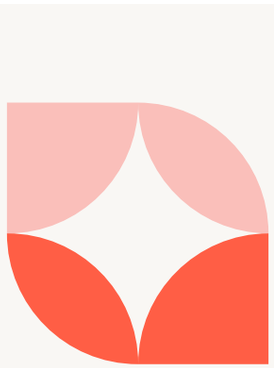

databricks

SECTION 7

The architecture decision that defines the next decade

Global institutions and governing bodies have reached a shared conclusion from different angles:

- AI is now a structural force in the macro outlook, not a temporary theme.
- Data fragmentation, opaque models and fragile infrastructure amplify both operational and systemic risk.
- Regulatory expectations around AI, algorithmic trading, resilience and data governance will keep rising.

In that context, the question for capital markets leaders is not whether to modernize the trade lifecycle, but how quickly they can move from:

- Fragmented data, bespoke pipelines and experimental AI, to
- A unified, governed Databricks Data + AI Platform that supports research, trading, risk, finance and compliance end-to-end.

The Databricks Data + AI Platform already underpins this transformation for firms such as RBC Capital Markets, State Street, Nasdaq and ABN AMRO — helping them compress research and risk cycles, reinvent trading intelligence and launch new data products.

> In an environment where AI is a macro variable, alternative data and provider models are shifting, and supervisors are watching data and models more closely, the firms that will lead in capital markets understand that the model is not the strategy. The platform is.

Models will keep changing, and the firms positioned to win can toggle to the latest and greatest without disruption, because their proprietary data foundation is already in place, already governed, and already working for them. That same foundation gives them the tools to train AI on their own data, run the evaluations to ensure it performs, and orchestrate multi-agent workflows that turn proprietary knowledge into production intelligence. Every new model that enters the market becomes more powerful in their hands than in a competitor's. The firms building that foundation now are the ones that will spend the next decade compounding that advantage. Not catching up.

References

Morgan Stanley Research, *AI Market Trends 2026: Global Investment, Risks, and Buildout* (2026).

ESMA, *"Shortening the settlement cycle to T+1 in the EU" — Markets and Infrastructure work program* (2026).

Tao Zhang (BIS), *The financial stability implications of artificial intelligence and digital finance* — remarks at International Financial Week, Hong Kong (26 January 2026).

Financial Stability Board, FSB Work Program for 2026 (3 February 2026).

ESMA, *Supervisory briefing on algorithmic trading in the EU* (26 February 2026).

Exabel, *The 2026 Alternative Data Market – Insights & Trends from Portfolio Managers and Analysts* (2026).

Wall Street Prep, *Bloomberg vs. Capital IQ vs. FactSet vs. Refinitiv – Comparison of the Costs and Features of the Four Leading Financial Data Providers* (updated April 2025).

Sparkco, *Bloomberg Industry Disruption Forecasts and Strategic Playbook 2025* (2025).

Frank La Salla (DTCC), "How the power of AI can revolutionize the financial markets," *World Economic Forum Annual Meeting* article (2026).

Databricks, Financial Services Outcome Maps and Databricks Marketplace partner listings for FactSet, S&P Global, Deutsche Börse, Stocktwits and others (2025–2026).

Databricks, *RBC Capital Markets uses AI to transform research* — customer story and Capital Markets GTM collateral (2025–2026).

Databricks, *State Street GlobalLINK uses Databricks for trading analytics and AI-driven liquidity management* — customer story (2025–2026).

Databricks, *Capital Markets Data Intelligence Outcome Map* — anonymized case studies on trade surveillance, credit risk calculators and trade repositories (2025).

Databricks, *Nasdaq reinvents finance with Databricks* — customer story (2026).

Databricks, *Data Intelligence for Financial Services* — customer evidence including ABN AMRO and other financial institutions (2025–2026).

databricks

Authors

ANDREA DESOSA
Global Head of Capital Markets GTM, Databricks

JUNTA NAKAI
VP — Global Head of Financial Services, Cybersecurity and Public Sector GTM, Databricks

KIM HATTON
Global Head of Financial Services Marketing, Databricks

About Databricks

Databricks is the Data and AI company. More than 20,000 organizations worldwide — including adidas, AT&T, Bayer, Block, Mastercard, Rivian, Unilever, and 70% of the Fortune 500 — rely on Databricks to build and scale data and AI apps, analytics and agents. Headquartered in San Francisco with 30+ offices around the globe, Databricks offers the Databricks Data + AI Platform, which includes Lakebase, Genie, Agent Bricks, Lakeflow, Lakehouse, and Unity Catalog. To learn more, follow Databricks on LinkedIn, X, YouTube, and Instagram.

databricks