

# Survival of the Fittest: Modernizing Capital Markets Infrastructure

## Introduction

Data is the core of a successful digital transformation strategy. But the extent of that success depends on how easily – and reliably – financial institutions can access and leverage relevant data from across the firm. As the breadth and depth of those requirements grow, it is clear that a modern data architecture is required for survival. While mission-critical systems built on legacy technology cannot be rebuilt or replaced easily, however, simply integrating point solutions only adds to the complexity and cost.

This is especially true in the capital markets, where success increasingly requires incorporating artificial intelligence and

machine learning approaches to provide proactive business insight as firms leverage data as a competitive weapon. This, however, can be challenging if firms' legacy data and compute architectures stand in the way.

Current demands for high-performance computing and big data analytics extend across traditionally siloed asset classes, technologies, and workflows (see *Exhibit 1, below*).

### Key Points

- Legacy data and compute architectures are standing in the way of capital markets firms' need to incorporate AI and ML approaches in order to provide proactive business insights as firms leverage data as a competitive weapon.
- The demands of regulatory (e.g. CAT, Reg SCI, MiFID II, FRTB) outstrip the compute capabilities of many existing applications and solution architectures
- It is time to look at new options to meet firms' most challenging data requirements including handling different data types while combining advanced in-flight analytics with transactional processing.
- The fees and friction of today will continue to decline but the data and complexity of finding alpha will only increase
- New offerings are combing the best of: in-memory performance; reliability; durability; AI/ML; and, the cloud

**Exhibit 1: Compute and analytics demands span diverse workflows and functions**



Source: TABB Group

To maximize competitive advantage, firms need a simplified, cost efficient, elastic, and dynamic compute environment. However, the industry’s historical reliance on a collection of point solutions can hinder a financial institution’s ability to meet ever-increasing business and regulatory demands. To break down these historic silos, new approaches are needed to increase scalability and performance while providing access to accurate and real-time data for insights and maintaining operational resiliency.

While financial institutions have invested billions of dollars, pounds and euros on implementing once-state-of-the-art systems, demands from both regulators and clients for ever-greater transparency into data and its analysis can mean tenfold increases – or more – in the amount of data generated. In addition, relying on “start of the day” or “end of day” data often is no longer acceptable. Instead, more functions are requiring near- or fully real-time access to data and analytics, which compounds the data architecture challenge.

On the regulatory front, compute requirements from initiatives such as the Consolidated Audit Trail (CAT) and Regulation Systems Compliance and Integrity (Reg SCI) in the US and, globally, the revised Markets in Financial Instruments Directive/Regulation (MiFID II/MiFIR) and the Fundamental Review of the Trading Book (FRTB) will outstrip the compute capabilities of many existing

applications and solution architectures. This, coupled with business pressures, means that financial institutions’ legacy data and high-performance compute architectures no longer are providing a competitive edge.

**Common approaches**

Financial institutions are increasingly focused on improving data availability, transparency and quality; sourcing more data across different types; and leveraging this cleaner, enriched data both to improve efficiency and find alpha. The problem is that past data architectures and infrastructures were not designed for today’s investment world, which has shifted toward multi-asset and multi-regional approaches and demands real-time decision support and risk management. In addition, the push for greater transparency, efficiency and regulatory oversight has fueled the spread of digitization throughout all asset classes, resulting in the increased proliferation of data, which of course has to be processed by any and all market participants simply to survive. Together, these drivers are leading firms to re-assess their data architectures for both transactional and analytical workloads (*see Exhibit 2, below*).

**Exhibit 2: New requirements are driving firms to re-assess data architectures**



Source: TABB Group

While financial firms have always leveraged data and technology to their competitive advantage, capital markets' technology tends to be fragmented and driven by its close alignment to organically grown businesses or, even more problematic, by companies built through a series of mergers and acquisitions. In either case the result is increased organizational complexity with an infrastructure highly dependent on duplicated data and operations. This creates an overly complex and costly operational structure reliant on excessive amounts of reconciliation and operational redundancy. This is not conducive to the technical and operational agility needed to compete in today's markets.

After a decade of constrained technology budgets dominated by crisis-instigated regulation, financial organizations' IT resources have been focused largely on "run the bank" initiatives, not "change the bank" projects. And, while the business and innovation climate is changing, financial infrastructure doesn't change that quickly. Today's financial data resides in multi-dimensional silos, meaning data can be aggregated by business line, data type, computational requirement, and/or dozens of other dimensions that need to be constructed at a moment's notice. Unfortunately for most firms, their data resides within a mix of mainframe, client-server and, lately, cloud applications over both local and wide area networks and is integrated through the use of backend-oriented middleware tools.

As a result, financial firms have been accumulating a technical debt, especially with regard to their data infrastructures. New requirements are often met with a point solution that is integrated with a pre-existing, poorly documented and brittle system. This is a pragmatic approach, as there are very few opportunities for greenfield developments within capital markets; critical systems cannot be taken offline for any length of time, and the risks to reputation and the bottom line are deemed too great.

### Data is the catalyst for success

As investment managers move toward more complicated investment strategies – e.g., multi-asset, multi-region, and highly efficient execution – they are increasingly relying on sophisticated platforms for portfolio management, order management, execution management and compliance. Again, data is the catalyst for success. The buy-side front office no longer can rely only on batch data. Instead, positions, cash, open orders and risk need to be calculated continuously throughout the day, ideally in near real time. This can only be achieved with reliable data architectures,

with very high performance, that scale seamlessly and ensure data quality and data governance are not compromised for the sake of speed or agility.

A number of trends have encouraged the buy side to ramp up their data infrastructures to support advanced analytics, including cognitive approaches incorporating machine learning (ML), natural language processing (NLP) and other artificial intelligence (AI) approaches (see *Exhibit 3, below*). It is no wonder the buy side has been an early adopter of cloud technology, as AI approaches rely on the ability to process and analyze vast quantities of both structured and unstructured data types.

**Exhibit 3: FS industry trends supporting move to modern data architectures on buy side**

Trend	Impact	Outcome
Low interest rate environment	Investors seeking higher returns, scrutinizing fund costs more closely. Passive investment gains in popularity.	Move to multi-region, multi-asset class, alternative and passive investments. Data analytics requirements across entire workflow.
Sell side facing increasing capital requirements from regulators	Risk trading on sell side decreases; focus on agency trading.	Buy side adopts technology-led approaches to find liquidity, support price formation.
Cloud computing adoption	Democratization of analytics; accessibility of compute & data storage.	A virtual technology arms race as firms adopt quantitative approaches and vie to model and execute market-beating strategies.
Regulatory-mandated push to greater transparency and increased surveillance	Firms faced with more comprehensive and real-time reporting requirements across enterprise.	Move to data architectures supporting holistic, real-time and predictive analytics.

Source: TABB Group

The sell side sales and trading functions meanwhile have drivers of their own, resulting in numerous pain points (see *Exhibit 4, below*).

**Exhibit 4: FS industry trends supporting move to modern data architectures on sell side**

Trend	Impact	Outcome
Cost of capital increased	Brokers decrease risk trading, in favor of agency model.	Need to enhance risk management capabilities to free up precious capital.
Regulatory-mandated push to greater transparency and increased surveillance	Ramp up in compliance and reporting spending.	Opportunity to recoup spending by monetizing resulting collected data and/or leveraging for competitive advantage and alpha generation.
Buy side under cost pressures	Commission wallet decreasing; move to passive and low-touch trading.	Customer relationship management key to focus on most profitable clients.
Cloud computing adoption at enterprise level	Lower costs, faster time to market.	Experimentation and advanced analytics approaches increasingly accessible.

Source: TABB Group

In addition to financial industry trends, there are drivers common to all businesses encouraging firms to re-architect their data solutions:

- The democratization of data analytics is fueling a rise of “self-service” approaches.
- Digital transformation at a corporate-wide level is creating hyper-connected enterprises that have access to more data than ever before.
- Desire for proactive (predictive and prescriptive), as opposed to reactive, analytics means firms are looking to cognitive analytic approaches that require access to both historical and real-time data as well as deeper and wider data sets across the enterprise.

Financial institutions hobbled by legacy technology will always struggle to keep pace with changes in their market or profit from the savings and opportunities new technology offers. Legacy architectures are being pushed beyond their capabilities, falling short on business service-level agreements (SLAs), while increasing both operational support costs and business costs. This lack of flexibility impacts the time to market for new ideas and competitive positioning in the market

## The future is now

In today’s fast-changing financial businesses, there is no patience for multi-year technology refresh programs or monolithic replacement approaches. Financial firms have embraced DevOps and highly modular architectures in order to be as nimble and agile as possible, while also meeting high standards for reliability and uptime. They are looking to augment and rationalize their technology stacks, with a preference for commodity hardware that can be elastically scaled and that also offers lower total cost of ownership. With a further need to rein in operational costs while reacting to changes in markets and demands, financial firms are looking for data architectures that can quickly adapt, while utilizing commodity hardware that can be quickly scaled.

Modern data architectures need to support greater business agility, interoperability, flexibility and cost efficiency by transitioning to elastic, dynamic compute infrastructures that can support both analytics and transactional workloads, and allow data to be managed in multiple ways, including supporting multiple types of relational and non-relational access. A modern architecture employs an elastic infrastructure, leverages microservices and containers, and

provides unified and secure access to all data from across the organization.

A common approach to data architecture in financial services has been to utilize a data technology stack comprising a collection of approaches – for example, historic data in a relational database, recent data in a columnar database, and real-time data in some sort of event processing software engine. To obtain a consolidated view, connections must be made directly to various systems; an approach that gets increasingly complex over time.

While technologies would love to rethink their architectures with a clean-slate approach, a “rip and replace” strategy is not typically practical or indeed realistic. Ideally, firms, when integrating or adding new features to existing systems, try to reduce the reliance on existing monolithic systems over time. In this way, legacy systems near end-of-life do not need to be ripped out, but rather can be left to retire gracefully as new functionality and features are added incrementally via newer approaches.

One-way firms are solving these problems is through new cloud infrastructures. Firms are embracing elastic computing approaches. For Tier 1 front-office functions, this is likely to mean private cloud; however, Tier 2 firms are thinking more about public cloud, although there will continue to be workloads with hosted or managed solutions in a co-location facility or third-party data center. Firms therefore are still faced with trade-offs when deciding how to scale solutions.

## AI and ML

Front-office transactions and analytics, which increasingly span diverse use cases (*see Exhibit 5, below*) are increasingly incorporating artificial intelligence (AI) and machine learning (ML). These technologies are often supported through specialized software and hardware for performance reasons and include in-memory databases. In-memory databases typically save time by avoiding the need to write/read from a physical disk. While firms increasingly are leveraging these platforms, they are not without their own challenges, especially around data persistence, data duplication, scalability and reliability if memory is exceeded. Data stored in-memory can be lost when the system closes or goes down unexpectedly, so strategies for handling overflow data must be implemented. Accommodating unexpected spikes in traffic in a reliable and cost-effective way is also challenging with in-memory architectures, as the cost of supporting rarely used resources has to be weighed

against the potential cost of performance and reliance issues to the business.

In-memory solutions which handle both transactional and analytical data workloads allow advanced analytics to be run “in flight” on transaction data (as opposed to after the fact, once data is written to disk). However, these still require the use of a considerable amount of expensive specialist hardware. Balancing the expense of full persistence and reliability against the performance of in-memory has been challenging for firms.

**Exhibit 5: Sample front-office use cases for high-performance data architectures**

Use case	Challenge
New/expanding algorithmic infrastructures	Trading in established electronic markets, such as equities, is ramping up use of algorithms for decision support and execution, incorporating more data and new cognitive analytic techniques. A regulatory push for more transparency is driving more markets to go electronic – resulting in algos spreading to other asset classes, such as fixed income.
Alternative data use in trading	With a prolonged low-interest-rate environment amplifying focus on costs, financial firms are incorporating alternative data into their trading models in order to find alpha or best execution in the market. Legacy data systems were not designed with the vast quantities of structured and unstructured data this entails.
Central risk books	Regulatory-mandated changes to how banks manage risk mean firms must be able to access an ever-growing set of internal and external data in order to perform highly computational-intensive risk modelling more frequently than before – and increasingly this needs to be done on a business line or even enterprise level.

Source: TABB Group

Lately, though, new technologies have become available that combine transactional and analytical workloads, AI and ML, and cloud and commodity compute resources, but without the cost and reliability issues of traditional in-memory approaches. It comes at a good time, with firms evaluating their compute requirements with a view to managing a transition to elastic compute infrastructures.

## Conclusion

The asset management and capital markets industries are facing an existential challenge as fee, revenue and regulatory issues seem to be pressuring virtually all aspects of the business. But all is not lost; there are significant opportunities as well. It is not hyperbolic to say, however, that realizing these opportunities will depend largely on a firm’s data capabilities.

There are various data architecture approaches available to address the most challenging data requirements. While the point solution has been to look at in-memory database approaches, it is time to look afresh at new options in the space that offer greater efficiency at lower operational costs, meeting the need to handle different data types while combining advanced in-flight analytics with transactional processing. Firms need to approach new offerings combining the best of in-memory performance (without the limits imposed by memory), reliability and durability, AI and ML, and the flexibility and scalability of cloud, as the industry grapples with managing in a world tipped toward lower costs and increased agility.

While the industry focus has been to respond to post-crisis regulation, the world has changed. Customers want innovation, new products, and new services. Unfortunately, these new demands are difficult if not impossible to provide with today’s legacy and limiting heterogenous data technology infrastructures. While firms can survive with outdated approaches for a limited time, the clock is ticking.

If financial firms do not start rethinking their technology and data architecture plans now, it will be even more challenging to meet the demands of the future. The fees and friction of today will continue to decline but the data and complexity of finding alpha will only increase. This will force a restructuring of virtually all financial businesses. To survive, firms need to get started today.



*Distributed by InterSystems*



**About TABB Group:**

TABB Group is a financial markets research and strategic advisory firm focused exclusively on capital markets. Founded in 2003 and based on the methodology of first-person knowledge, TABB Group analyzes and quantifies the investing value chain, from the fiduciary, investment manager and broker, to the exchange and custodian. Our goal is to help senior business leaders gain a truer understanding of financial market issues and trends, so they can grow their businesses. The press regularly cites TABB Group members, and analysts routinely speak at industry conferences and gatherings.

For more information about TABB Group, visit [www.tabbgroup.com](http://www.tabbgroup.com)