

Low latency has long been viewed as a race to zero - measured in nanoseconds and microseconds and defined by how quickly market data can be delivered and acted upon. That paradigm is now shifting. In modern electronic markets, while speed remains critical, it is no longer the sole determinant of competitive advantage.

Instead, low latency has become a broader, multi-dimensional capability. Firms are increasingly focused on accessing precise, high-quality, and trusted datasets instantly. Data that is not only fast, but also usable, consistent, and ready for analysis at scale. The ability to convert a market event into actionable insight, with minimal friction, is now the true differentiator.

This shift is being driven by several converging forces. Market structures have become more fragmented and data-intensive, while the evolution of algorithmic and quantitative trading has increased reliance on both real-time and historical datasets. As a result, expectations have evolved.

Firms are no longer asking how quickly they can receive data, but how quickly they can trust it, analyse it, and act on it across increasingly complex workflows.

Data scale meets cloud accessibility

The scale of market data has grown exponentially, creating both challenges and opportunities. With more than 100 petabytes of data spanning more than 400 venues, London Stock Exchange Group's low latency offering exemplifies the industry's shift towards vast, interconnected data ecosystems.

However, scale alone does not create value. Historically, firms have spent significant time and resources acquiring, transporting, and managing data before any meaningful analysis could begin. This operational burden often delayed insight and limited agility.

Cloud-based delivery models are now fundamentally changing this dynamic. By making large datasets readily accessible within cloud environments, firms can bring their analytics directly to the data. This eliminates the need for complex data pipelines and significantly reduces time-to-analysis.

As a result, workflows are becoming more fluid and iterative. Research, trading, and risk teams can access shared datasets, isolate relevant subsets, and test ideas in near real time. This not only accelerates decision-making but also enhances resilience, enabling firms to respond more effectively to volatile market conditions. Increasingly, this data is made available natively within cloud environments, enabling clients to access it where they already operate. Whether through direct feeds, cloud-native analytics, or API-based delivery, this flexibility allows firms to integrate market data seamlessly into their existing workflows - effectively meeting them where they are.

The new low latency equation: Speed, scale, and certainty

Pramod Nayak, director of product management, London Stock Exchange Group, examines how low latency strategies are being redefined by data scale, cloud delivery, and advanced capture technologies - highlighting why speed alone is no longer enough to compete in increasingly complex, data-driven global markets today.

LSEG DATA & ANALYTICS

Shared storage and a new operating model

A key enabler of this transformation is the adoption of shared storage models. Traditional approaches required firms to replicate large datasets within their own infrastructure, leading to duplication, high storage costs, and fragmented workflows.

Shared storage fundamentally reshapes this model. Data remains centrally managed and maintained by LSEG within the cloud, while clients can access, query, and utilise it as if it were their own. This enables firms to extract only what they need without the burden of replication or infrastructure overhead, dra-

matically reducing time between data access and analysis.

The impact extends beyond efficiency. By enabling multiple teams to work from a single, consistent dataset, shared storage enhances collaboration and ensures alignment across functions such as trading, quantitative research, and compliance. It also improves auditability, as all users operate from the same underlying source of truth.

Crucially, this model reflects a broader shift in how firms view data - not as a static asset to be moved and stored, but as a dynamic resource to be accessed and analysed in place.



Pramod Nayak

Capture technology: The 'invisible engine'

While advances in cloud and storage are transforming accessibility, the integrity of any low latency strategy ultimately depends on how data is captured. This transformation has been underpinned by sustained investment in capture, normalisation, and distribution capabilities, enabling data to be delivered with both precision and consistency at global scale.

Capture technology operates as the 'invisible engine' behind market data, ensuring that events are recorded accurately, completely, and with precise timing. Without this foundation, it is impossible to reconstruct the true state of the market or rely on data for critical decision-making.

High-quality capture involves collecting data as close to the source as possible, applying high-precision timestamping, and ensuring lossless transmission. These capabilities are essential for supporting a wide range of use cases, from order book reconstruction and backtesting to market microstructure analysis and validation of execution strategies.

In parallel, data normalisation plays a critical role in making this data usable.

With hundreds of venues generating data in different formats, standardising this information into a consistent structure allows

firms to analyse cross-market activity more efficiently. This reduces complexity and enables faster, more scalable analysis across global markets.

Convergence and the future of low latency

Looking ahead, it is becoming less about the difference between real-time and historical data, and more about consistency between the two. Firms want to know that when they go back and analyse data, they are seeing exactly what they saw in real time - without gaps or distortion. That is what gives them confidence when they are testing strategies or trying to understand how the market behaved. Achieving this consistently at scale is where modern, cloud-based architectures start to play a critical role.

Cloud-native architectures will continue to accelerate this convergence. By decoupling storage from compute and enabling flexible, on-demand access to data, these models allow firms to scale their operations while controlling costs. They also support more iterative workflows, where strategies can be tested, refined, and redeployed with minimal delay.

At the same time, data quality will remain a central focus. Precision, depth, breadth, and consistency are becoming essential attributes, particularly as trading strategies grow more

sophisticated and sensitive to market microstructure.

For firms seeking to modernise their low latency capabilities, the priority is no longer simply reducing latency metrics. Instead, the focus shifts to building a holistic data strategy that aligns with specific business objectives based on a fully deterministic latency profile per market.

This begins with understanding the decisions that need to be optimised - whether in execution, risk management, or research - and ensuring that the data supporting those decisions meets the required standards of accuracy, completeness, and timeliness.

Equally important is the adoption of scalable infrastructure. Cloud delivery, shared storage, and integrated data models can significantly reduce operational complexity, allowing firms to focus on generating insight rather than managing data logistics.

Ultimately, the evolution of low latency reflects a broader transformation in financial markets. Speed remains essential, but it is the combination of speed, scale, precision, determinism and seamless accessibility - delivered through cloud-native, globally distributed infrastructure - that defines success. Firms that embrace this shift will be best positioned to extract value from their data and compete in an increasingly dynamic trading landscape.