Improving TCA with Kinetica

Transaction Cost Analysis (TCA) in capital markets refers to a comprehensive evaluation of all costs associated with executing a trade.

This analysis takes into account not just the direct costs of the trade (such as the spread or commissions), but also indirect costs such as market impact, liquidity costs, and opportunity costs. The goal of TCA is to provide a more accurate picture of the actual cost of a trade for both regulators and traders.

TCA uses a range of complex financial metrics to accurately determine the total cost of a trade. Some of the metrics used in TCA include Implied Volatility, Market Impact, Liquidity Costs, Opportunity Costs, Price Slippage, Notional Value, Volume Weighted Average Price (VWAP), and Time Weighted Average Price (TWAP). These metrics are used in combination when calculating the total cost of a trade.

Individual Trade Metrics versus Roll-up Metrics

When calculating Transaction Cost Analysis (TCA) metrics, there are two types of metrics that can be used: individual trade metrics and rollup metrics.

Metrics at the individual trade level provide detailed information about each trade, including its execution price, size, and the time it was executed. These metrics allow traders to analyze the performance of each trade individually, and to identify specific areas for improvement in their execution strategy. Examples of metrics at the individual trade level include slippage, fill rate, and latency.

Rollup metrics are summary statistics that aggregate trade-level data across a specified time period or across multiple trades. Rollup metrics provide a high-level overview of trading performance, allowing traders to quickly identify trends and patterns in their execution. Examples of rollup metrics include VWAP (Volume-Weighted Average Price), Implementation Shortfall, and Arrival Price. These comparative metrics can be used to evaluate and improve performance across trading desks.

TCA Rollup Metrics Challenges

Calculating TCA rollup metrics on continuously changing trade data can be technically challenging , especially when it comes to issues related to latency, high cardinality joins, and operationalizing insights in near real-time.

Latency

In order to calculate TCA metrics in near real-time, it is necessary to have low latency data processing capabilities. High latency can result in missed trading opportunities and delayed insights. This requires efficient data ingestion, processing, and analysis pipelines, which can be technically complex and require a significant amount of computational resources.

Metric calculation complexity given high cardinality joins

TCA rollup metrics often require aggregating data across multiple dimensions, such as time, asset class, and trading venue. This can result in high cardinality joins, which can be computationally expensive to calculate. This requires efficient algorithms and data structures, as well as specialized data processing and storage systems that can handle large amounts of data.

Operationalizing insights in near real-time

Once TCA rollup metrics have been calculated, it is important to operationalize the insights to get the information in front of traders in near real-time. This requires having a robust and scalable infrastructure that can handle the high volume of continuously updated data generated by the TCA analysis. This infrastructure should include real-time dashboards and alerts that provide traders with actionable insights and facilitate quick decision-making.

Current Approaches Are Falling Short

Tools like Spark and Redis are popular distributed data processing tools that can be used for a variety of tasks, and have been used to calculate TCA rollup metrics. However, banks and hedge funds are learning they are not ideal for all TCA use cases due to several weaknesses:

Latency

Spark can introduce latency when processing extreme volumes of data. This can make it challenging to perform TCA analysis in near real-time, which is often required for trading applications. While there are mechanisms to reduce latency, they do not scale linearly, which ultimately results in performance issues.

Cost

Spark can be expensive to deploy and maintain, especially in large-scale TCA use cases that require joining real-time streams with historical data and aggregating within moving windows. These workloads require significant hardware resources and extensive data engineering.

Complexity

Spark data pipelines begin to get complex given the recursive nature of time-series analytics required for TCA rollup metrics. Further, in order to deliver fast reads to downstream trading desks, the results need to be moved into a tool like Redis or other key-value database within the microservices layer, further adding complexity and restricting agility.

Ultimately, institutions are finding that current approaches are not providing TCA rollup metrics in a timely and agile manner despite extensive investments.

Kinetica

Kinetica is a real-time analytic database with unique capabilities to drive differentiated TCA applications. Kinetica is used for Transaction Cost Analysis (TCA) in capital markets by banks like Citi and hedge funds like Point 72.

Use Case: TCA Rollup Metrics

	With Spark	With Kinetica
SLAs	Daily and Weekly	< 5 Minutes
H/W Environment	700 Nodes	16 Nodes

Kinetica uses a native vectorized database architecture, ideal for TCA rollup metrics that involve high cardinality joins and time series windowing functions due to its ability to handle complex data operations with low latency. The vectorized architecture requires less tuning and up-front data engineering, thereby improving developer productivity and making it easier to implement and iterate on TCA rollup metrics.

In a vectorized query engine, data is stored in fixed-size blocks called vectors, and query operations are performed on these vectors in parallel, rather than on individual data elements. This allows the query engine to process multiple data elements simultaneously, resulting in faster query execution and improved performance. In addition to improving query performance, this vectorized approach can also reduce the amount of compute and data engineering required, making them more efficient and cost-effective. This is in contrast to conventional distributed analytic platforms – like Spark - which process data on a row-by-row basis, which is usually much slower and requires more computational resources.

Key Features:

- *Fast ingest through a distributed, headless, lockless architecture with native connectivity with streaming tools like Kafka is used to rapidly ingest and process extreme amounts of market data, trade information, and pricing data, necessary for accurate TCA calculations.*
- *In-database analytics capabilities perform continuously re-aggregated complex time-series calculations that form the basis of TCA metrics.*
- *Converged columnar and key-value store database makes it easy to enable high-speed reads of freshly calculated TCA rollup metrics.*
- *Kinetica can be integrated with other tools, such as trade cost estimation tools and execution management systems, to provide a complete TCA solution.*

By leveraging Kinetica's high-performance database and analytics capabilities, investors can perform accurate and efficient TCA rollup calculations, allowing them to make informed decisions and optimize their trading strategies.

Contact Kinetica to learn more.