

[SEARCH](#)[Advanced Search](#)[HOME](#) [PRINT EDITION JULY 20, 2009](#)

To Reduce Delays, Firms Filter-Or Cut Down-The Data Coming In

July 20, 2009
Chris Kentouris

Although low-latency--and even ultra-low latency--have become Wall Street's latest mantra, simply reducing the time it takes for tick data to move through pipelines isn't enough to ensure the best trades are made.

Financial firms specializing in high-frequency trading or even those relying on out-of-the-box algorithms are now turning to an array of high-speed analytic and query tools such as complex event processing engines to filter and analyze the surging tide of messages flooding their servers with market data. Other chief information officers are going further and trying to cut down on the amount of data they choose to analyze.

[Printer Friendly](#)[Email](#)[Reprints](#)[Reader Comments](#)

Securities Industry News Research Vault

Your New Source
for Comprehensive
News, Research, and
Industry Intelligence

Unchecked, New York research firm Tabb Group estimates that firms will have to double the number of equity ticks analyzed every second from 150,000 to 300,000 by mid-2010.

And that's not all firms have to worry about. They must also combine their analysis of tick data with historical data to make a final decision on when to execute an order and in which venue.

How to tackle the data overload is becoming a critical concern for top executives. "Many CIOs recognize there is no point in trying to combine a firehose of market data, when the actionable data is just a fractional subset of what is coming in," said Vijay Oddiraju, president of New York-based Volante Technologies, a low-latency integration

service provider. "The extra data just burdens the system and adds latency, whether it is a complex event processing engine or a consuming application."

A large European investment bank, which Oddiraju declined to name, has just installed Volante's Designer software to provide two layers of filtering for tick data from a large data vendor. The first "bulk" filter reduces data volume by 80 percent by removing unneeded fields, as well as equities types in which the bank does not trade. The second level of filtering customizes the data for the algorithms of different trading units, stripping off an average of another 90 percent of the remaining data.

Firms offering CEP software tout its ability to analyze real time and historic data combined to develop and power pricing and execution algorithms within seconds. In doing so, some filtering takes place when the data is cleansed, validated and normalized to prevent any erroneous information from being factored into the analysis.

PhaseCapital, a Boston-based high frequency trading shop, says it has implemented a CEP engine from rival Boston-headquartered StreamBase to scrub and normalize consolidated and direct market data feeds and to help manage high-speed routing and trade execution across data centers in Boston and Jersey City, N.J.

The cleansed data from the consolidated feed in PhaseCapital's Boston data center prices the firm's portfolio holdings so that profit and value at risk (VAR) can be calculated in real time. The data also feeds proprietary algorithms which generate what PhaseCapital calls trade signals, normalized target portfolios and parent orders which are then broken down further by a proprietary execution management system into executable orders. Those orders, which rely on algorithms running on the StreamBase platform, are traded and monitored using direct market access (DMA) sessions with a variety of execution venues such as NYSE Arca, Bats, EdgeA, EdgeX and Inet through a compressed version of the FIX protocol.

Complex event processing tools often come with their own dashboards and are used with third-party data visualization software which firms to drill down to specific data points.

Such "tools are great for creating and executing algorithms if certain conditions are met. However, they don't provide a trader with enough analysis of how the algorithm is performing or how the market is doing on a real-time basis," says Brian O'Keefe, president of the Americas for Panopticon, a data visualization software firm based in Stockholm.



While Wall Street firms will increasingly embrace the combination of CEP and data visualization technologies as they become more familiar with their benefits, they aren't the solution for every company.

According to some buy-side firms, complex event processing--which is used to predict the market impact of economic, political and other events or scenarios--is best-suited for only the most basic algorithms and not customized models. "We wanted to base our algorithms on fundamental mathematical and scientific modeling rather than micro and probabilistic solution sets delivered through CEP," said Paul Roland, managing principal of Theoriem, LLC, a quantitative trading

firm in Philadelphia.

Theoriem has installed Xenomorph's TimeScape software to analyze large sets of real-time and tick-data on an intraday basis rather than rely on Microsoft Excel spreadsheets--a time consuming process. "Storing the data in TimeScape allowed quant analysts to reduce the time to formulate trading models to several minutes rather than several hours or days," said Roland. Theoriem also can add new data vendors and asset classes on the fly, freeing up IT staff to focus on broader strategic initiatives rather than day-to-day tactical work