

Where Low (Latency) Is High (Value) It's a Technology Arms Race for Securities Firms and Vendors Alike

Securities Industry News

2007-03-05

By Chris Kentouris

Obtaining market data used to be as simple as taking a feed or terminal from an aggregator such as Bloomberg, Reuters or Thomson Financial. But high-frequency electronic trading and direct-market access (DMA) have made it a lot more complicated. Profits and opportunity costs are tied up with latency--the lag time between receiving data and acting upon it, or sending an order and having it executed--which is measured in milliseconds, and even microseconds.

Firms used to talk about how quickly a trade could be executed, says Hugh Hughes, chairman and CEO of low-latency market data feed provider Fixnetix in London. 'Now they are talking about how quickly they must receive the information,' Hughes says.

Accompanying the explosion in trading activity are vast streams of data that require the electronic feeds to add capacity and run at peak performance, lest latency result. Market data volumes rose to new plateaus after the 2001 introduction of decimal pricing of stocks and their narrower quote spreads, but the trade reporting and compliance requirements of Regulation National Market System (NMS) in the U.S. and the Markets in Financial Instruments Directive (MiFID) in Europe--along with the continued growth of algorithmic and other programmatic forms of trading--are poised to take that up to far greater heights. The Options Price Reporting Authority has said that as of January 2008, firms receiving its direct U.S. exchange pricing feeds must handle peak rates of 400 megabits of data per second, an eightfold increase from January 2005.

One solution seems obvious: Stop using data aggregators, which are several times slower than direct market feeds. Well, not so fast.

'Firms seeking to process data directly have the continual burden of adding hardware and bandwidth to maintain the increase in data volume,' says Frank Piasecki, president of New York-based market data feed handler and processor Activ Financial Systems. 'These costs are not sustainable for doing business globally, nor that feasible for the newest exchange feeds delivering depth of book, such as the New York Stock Exchange OpenBook and Nasdaq TotalView.'

To meet the best-execution mandates of Reg NMS and MiFID, firms will have to connect to a widening array of exchanges, electronic communications networks (ECNs), alternative trading systems, crossing networks and other dark pools. Each venue has its own data format and nomenclature, meaning that the data must be normalized and consolidated.

'Firms are being inundated with market data from multiple liquidity pools and might have normalized data from one source, but not necessarily from another,' says **James Leman**, a veteran Wall Street trading and technology executive who now heads the capital markets practice of New York-based

consulting firm Westwater Corp. He notes that because each source of data has different product strengths, normalization may have to take place multiple times.

Feed handlers such as those of Wombat Financial Software, Activ Financial and Fixnetix might do some cleaning and consolidation, but they can't completely solve the problem of downstream distribution, says Leman. Stream and event processors may be needed to analyze data while message buses transport it to the right applications. Then there is the need for storage for the terabytes of historical tick information used in back-testing strategies and generating audit trails for internal and regulatory reporting.

'Downstream data-quality ownership problems abound,' says Piasecki. 'Who is there to assure the quality of the consolidated book of books?'

Solutions and their vendors fall into several categories, and firms often take a mix-and-match or best-of-breed approach. 'One of our broker-dealer clients uses Vhayu Velocity for real-time tick analysis and Wombat for direct exchange feeds, [Progress Software's] Apama for analyzing streaming information and Tibco [Software's] Rendezvous as a message bus,' says John Coulter, VP of high-performance technology vendor Vhayu Technologies in Los Gatos, Calif.

Paying Up

Including data feed handlers and middleware to facilitate the movement of data across applications, the price tag for each feed can climb into six figures. And a global organization can be taking several dozen feeds. A New York hedge fund manager tells Securities Industry News that his firm spent about \$400,000 for each direct exchange feed during its first year of operation. That figure almost doubled within five years as more data feeds and middleware were added.

The costs and administrative burdens don't mean that firms are resigned to using incumbent data aggregators and hoping for the best. They do face tough choices. The ultimate decisions rest on a series of factors including the type of firm, its budget and how much latency it wants to tolerate.

'Market-making firms with auto-quoting capabilities and those in the algorithmic or black-box trading space--typically top-tier investment banks--need near-zero latency, while funds with longer-term strategies might do better with data consolidators,' says Jeromee Johnson, senior analyst with research firm Tabb Group of Westborough, Mass.

'The flexibility to choose direct data feeds for latency-sensitive sources and consolidated feeds is the best choice,' adds Activ's Piasecki, whose company offers both direct and consolidated feed handlers.

Brokerage giant Bear Stearns & Co. historically relied on consolidated feeds and data distribution platforms, but when latency delays mounted to two to three seconds, it began benchmarking data platforms that supported ticker plant software. That latency is outside 'when you are trying to make decisions in a time frame of under 20 milliseconds,' says Jerome Downey, senior managing director in Bear Stearns' equities analytics and systematic trading division in New York, which houses its algorithmic trading systems.

In 2005, Bear Stearns signed a global licensing deal with Wombat of Incline Village, Nev. for direct exchange and ECN feed handlers for equities, futures, options and currencies in the U.S. and overseas. Downey says that Wombat's data latency was half that of other vendors surveyed, which he declined to name.

Wombat, which says its systems are in seven of the top ten investment firms, has signed Merrill Lynch to install low-latency feeds for a new global equity trading platform. Another client is EdgeTrade, an agency brokerage that offers DMA and algorithmic trading strategies.

While JP Morgan Chase & Co. relies on several data consolidators to keep track of interlisted shares, it acknowledges it's taking steps to stay competitive in algorithmic trading by 'building our own data feed handlers to a number of exchanges with the help of Neovest,' says Carl Carrie, head of product development in the electronic client solutions group at JP Morgan, which bought the Neovest DMA business in 2005. 'We are linked to all of the major U.S. markets and are expanding in Europe and Asia-Pacific,' and Neovest offers connections to 15 dark liquidity pools.

Aggregators Respond

The demand for direct exchange feeds is prompting traditional data aggregators to expand their reach and bringing new players into the market. One is the year-old Fixnetix, which says that as of May it will offer direct feeds to 27 European markets with maximum latency of two milliseconds; the firm has 10 clients and interest from an additional 35, says CEO Hughes.

In 2005, Reuters introduced a set of low-latency direct feed services, including full-order book feeds. Reuters Data Feed Direct uses Reuters industry codes as well as its data models and application programming interfaces so that firms do not have to rewrite their legacy apps. Reuters has linked directly to nine U.S. exchanges and to those in London, Sydney, Seoul, Hong Kong and Eurex in Frankfurt. Next up for 2007 are the major Japanese exchanges and Euronext.

But many sell-side firms running legacy data platforms end up with 'parallel structures' in supporting consolidated feeds as well. 'It's about the cost of ownership,' says Mike Powell, global head of real-time enterprise information for Reuters in London. 'Global players will look at each of the strategies of their business units and determine which exchanges they need to connect to directly and which of the end users are less latency-sensitive.'

Whether opting for direct or consolidated feeds, firms are paying closer attention to their internal messaging platforms. The more efficiently the data can move into their applications, the faster decisions can be made. Reuters boasts an array of data distribution channels, as does Wombat, which in December announced a partnership with complex event processing vendor Coral8 of Mountain View, Calif. Wombat also has an agreement with 29West of Warrenville, Ill. to embed its Latency Busters Messaging (LBM) software as a low-latency message transport and alternative to the competing Rendezvous. Including Wombat's clients, LBM has over 100 deployments at exchanges, banks, hedge funds and proprietary trading desks.

Firms also need bigger and faster analysis of tick prices so they can test strategies and simulate trades. Complex algorithms are analyzing vast quantities of streaming market data and executing strategies in multiple asset classes. The Wombat infrastructure, for instance, can feed streaming data into a platform such as Vhayu Velocity to create an intraday tick cache for running fast queries against historical and real-time market data. 'One brokerage client asked us to map every tick of data across 90 exchanges and do VWAP [volume-weighted average price] analysis on a tick-by-tick basis,' recalls Vhayu's Coulter, whose firm recently expanded into fixed-income tick analysis.

The arms race among securities firms trying to keep ahead of the latency and analytics game is replicated among vendors, which periodically boast breakthroughs in benchmarking tests. StreamBase Systems of Lexington, Mass. in April 2006 touted the 'unprecedented performance ... and near-zero latency' of its version 3.0 StreamOptimizer. Aleri Labs of Chicago in June said it 'broke the submillisecond latency barrier for event stream processing.' Others counsel caution in interpreting how such results apply to real-world situations.

To be sure, these are highly advanced technology developers that stay on top of the latest processor innovations from the likes of Advanced Micro Devices and Intel Corp. for incorporation into their platforms. One of them, Kx Systems of Palo Alto, Calif., which supplies the high-performance kdb+ database and related products to a who's who of global financial institutions, added a new dimension to its strategy in January with the appointment of Niall Dalton as chief solutions architect. His background is in extremely large-scale, real-time processing environments, and he's advising Kx customers on the system architectures best suited to their performance requirements.

Event stream processors are often used with tick databases, allowing applications to detect patterns and act upon opportunities in real time. 'Instead of storing the data and querying it, we let an organization specify the strategy up front in terms of when/then rules,' says John Bates, VP of Apama products for Progress Software of Bedford, Mass. 'The faster one pulls in data, the faster analysis needs to occur to get to market first.'

Major clients include JP Morgan, Deutsche Bank and ABN Amro. Eighteen months ago, Progress began selling directly to the buy side and nabbed Aspect Capital, a \$2.4 billion London hedge fund that uses the platform in foreign exchange trading.

But with everyone competing on speed, how fast can any advantage last? At some point, the emphasis may shift to the quality and reliability of the data. 'Ideally, normalization should take place at the start, before prices hit any decisionmaking engines, but it is often occurring as the data simultaneously moves through the streaming process, which is necessary in a quantitative strategy but could lead to errors,' says David Hirschfeld, COO in New York of data management software provider Asset Control.

As the discipline of enterprisewide data management--ensuring a consistent view of data across organizations--becomes more entrenched, tick data needs to be accurate if firms are to link it to static reference data and corporate actions information. With firms retaining different databases for each type of information--often multiple ones for reference data--that is no easy task.

Complicating matters, from a data governance perspective, firms specializing in algorithmic, black-box or quantitative methodologies still separate the market data decisionmaking process from that of static reference data. 'When it comes to real-time market data, traders and IT staffers often carry a lot of clout on which and how many data feeds they need,' says Alan Paris, a director in the financial advisory services group of PricewaterhouseCoopers in New York. 'That is because market data is consumed, whereas other types of data are retained.'

Still, firms are trying to control the use of market data at the desktop level. Hence, there is some room for cooperation between market data specialists and their data management peers in middle and back offices. 'The investment data services group puts tighter restrictions on data, often asking what data the investment management desk can live without,' says John White, VP of global data management at State Street Global Advisors (SSGA) in Boston.

According to White, since his group's formation in 2002, SSGA's data procurement budget has grown by only 4 percent annually, instead of the previous 20 percent, even as SSGA added vendors specializing in over-the-counter derivatives valuation. The total number of vendor relationships remained constant as some contracts were eliminated and others leveraged to add new services.

Joseph Morant, managing director of buy-side consulting firm Citisoft in Boston, recalls a study conducted on behalf of a global fund management firm to evaluate its data expenditures. After a year of analysis, he says, the firm reduced its number of data feeds by more than 30 percent, to 35, which ultimately lowered data procurement costs also by 30 percent annually. In addition, the buy-side firm--which ran mutual fund and separate managed account portfolios--set up ground rules for which branch office was responsible for cleaning each trading venue's market data.

But the fund manager's streamlining efforts never end as it adds more trading venues and asset classes. 'It's a journey, and everything comes with a trade-off,' says Morant. 'All companies need to be vigilant in refining their framework that supports their business and understand they may have to adapt to change.'

(c) 2007 Securities Industry News and SourceMedia, Inc. All Rights Reserved.

<http://www.securitiesindustry.com> <http://www.sourcemedia.com>

++>