Inside Market Data

July 2013

waterstechnology.com/imd

DATAFEEDS

SPECIAL REPORT

Sponsored by:

S&P CAPITAL IQ
McGRAW HILL FINANCIAL

S&P CAPITAL IQ
McGRAW HILL FINANCIAL

Real-Time Delivery at the Speed You Need

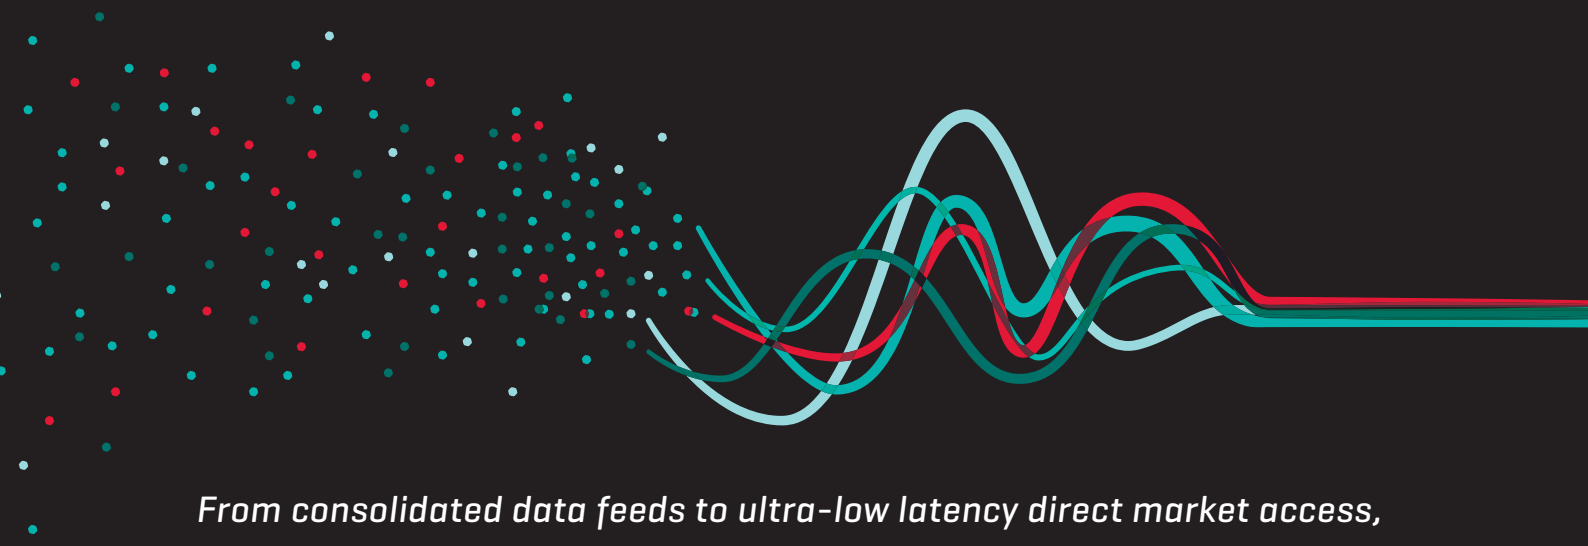

From consolidated data feeds to ultra-low latency direct market access, our delivery options align to your firm's specific data consumption needs.

- Cost-effective hosted solution that can service multiple user profiles
- Robust delivery options including feeds, workstations and APIs that allow for seamless integration into real-time workflows
- Comprehensive direct and value-added data covering over 150 global exchanges

Discover the ways we can enhance your performance.
For more information, contact us at info@capitaliq.com.

www.spcapitaliq.com

Inside Market Data | Awards 2013
Acquisition of the Year – as featured in Inside Market Data
WINNER S&P Capital IQ

S&P Capital IQ, a business line of McGraw-Hill Financial, is a leading provider of multi-asset class and real time data, research and analytics to institutional investors, investment and commercial banks, investment advisors and wealth managers, corporations and universities around the world.

LETTER FROM THE EDITOR

Feed Me, See More!

I remember it like yesterday: my first story about banks bypassing consolidated feed providers and deploying internal ticker plants to capture and process raw feeds of exchange data. It was May 2004, and the story assembled comments from *Inside Market Data*'s New York conference, where a panel predicted more firms would consolidate their own feeds using technology from specialist ticker plant vendors.

Alas, many of those vendors—including CMS Webview, Infodyne and Wombat Financial Software—are now gone, swallowed up by larger vendors looking for an edge, or out of business after struggling to stay afloat in an increasingly competitive market. And as traditional data vendors realized the potential impact of being displaced by exchanges, they built or acquired low-latency technologies to keep pace with their clients' needs and to exploit opportunities to insert themselves into the low-latency data flow.

Back in 2004, data consolidators combated disintermediation by pointing out the value they added through normalization and data quality. Little did I realize that the seed planted by those first direct feed pioneers would become a Little Shop of Horrors of competition over microseconds, and battles to handle high-volume data microbursts and traffic peaks in the millions of messages per second.

Now, these processes can be performed in split seconds by feed handlers and switches, using hardware-accelerated processors initially deployed to handle rapidly-rising market data rates, but which can also be used to perform high-volume repetitive processes.

Certainly, consolidated feeds provide more visibility than direct feeds from single venues, even if not as fast. Ultimately, the value of consolidated feeds is less in the process of consolidation—which anyone can do, should their developers have nothing more important to do—and more in the sheer array of exchanges, over-the-counter sources and proprietary and third-party datasets they can combine and distribute by connecting once to the source, and making its content available to their entire client base, leveraging their economies of scale and pricing their services accordingly, rather than each client connecting to every source themselves. And the potential for a raft of new venues emerging in the form of swap execution facilities—which will trade traditionally low-frequency over-the-counter instruments on exchange-like platforms—could present a bigger connectivity burden for potential participants, and could grant consolidated feeds a new lease of life. n

Max Bowie
Editor, *Inside Market Data*

Inside Market Data

Max Bowie, **Editor**
Tel: +1 646 490 3966
max.bowie@incisivemedia.com

Faye Kilburn, **US Reporter**
Tel: +1 646 490 3967
faye.kilburn@incisivemedia.com

Lee Hartt, **Group Publishing Director**
Tel: +44 (0)20 7316 9443
lee.hartt@incisivemedia.com

Jo Webb, **Global Commercial Director**
Tel: +44 (0)20 7316 9474
jo.webb@incisivemedia.com

Bene Archbold, **US Commercial Manager**
Tel: +1 646 736 1892
bene.archbold@incisivemedia.com

Elina Patler, **Head of Editorial Operations**
Claire Light, **Senior Marketing Manager**
Natasha Gordon-Douglas, **Marketing Manager**

Incisive Media
55 Broad Street, 22nd Floor
New York, NY 10004
Tel: +1 646 736 1888

Incisive Media
32-34 Broadwick Street
London W1A 2HG
Tel: +44 (0)20 7316 9000
Fax: +44 (0)20 7316 9250
E-mail: customerservices@incisivemedia.com

Incisive Media
14th Floor (Unit 1401-3), Devon House, Taikoo Place
979 King's Road Quarry Bay, Hong Kong
Tel: +852 3411 4900

Subscription Sales
Hussein Shirwa Tel: +44 (0)20 7004 7477
waters.subscriptions@incisivemedia.com

Incisive Media Customer Services
Haymarket House
28–29 Haymarket
London SW1Y 4RX
Tel (UK): 0870 787 6822
Tel (International): +44 (0)1858 438421

An incisivemedia publication

NEWS ROUNDUP

ASX Opens Full-Depth ITCH Feed to Vendors

The Australian Securities Exchange is in the test phase of connecting market data vendors to its ASX 24 ITCH feed of futures and options data, to enable users accessing ASX via vendor terminals and trading screens to see full market depth data from its ASX 24 derivatives market.

A large unnamed global data vendor and five global futures data providers began the process of connecting to the ASX 24 feed in May, after an initial rollout to trading members in March. Frank Hoer, market data manager at ASX, says he expects the development process to be complete in around a month, after which, users will be able to access full market depth for ASX 24 data via their vendor screens."

ASX 24 ITCH provides data on contracts for difference, and all tradable contracts on ASX's Trade24 near 24-hour derivatives market, including benchmark futures and options contracts on Treasury bonds, bank bills, cash rates, S&P/ASX equity indexes, energy and commodities.

Previously, ASX offered two feeds for its derivatives market: a feed commonly used by data vendors based on the Inter-Exchange Technical Committee (ITC) standard, which provided full market depth, but with multiple orders aggregated at each price level, and did not break out that depth into individual orders; and a FIX Protocol-based feed used by independent software vendors (ISVs) and some vendors that want to augment the ITC feed, as well as trading participants, which provides the best bid and offer plus five levels of depth, broken out to show individual order details.

However, the introduction of the new feed—which is based on the ITCH protocol and touted to reduce latency of market data from the exchange by 50 percent—means that vendors can now offer full-depth data to customers, Hoer says. "The new ITCH feed contains every single order in the book, allowing vendors to show both depth and details on a screen and toggle between the two, as has been the case for our equities market for many years," he says. n

NYSE Preps Arca Feed Migration to XDP Platform

NYSE Technologies, the data and technology arm of NYSE Euronext, is migrating its Arca BBO and Arca Trade data products to new Arca XDP BBO and Arca XDP Trades products, which use the vendor's Exchange Data Publisher (XDP) market data distribution platform for disseminating price and order data, and which will become the standard market data feed format for all NYSE, NYSE MKT and NYSE Arca data services, giving clients a single framework of message types, formats and delivery mechanisms for all existing and future NYSE data products. n

Barchart Connects to SFTI Network

Chicago-based data provider Barchart has connected to NYSE Euronext's SFTI (Secure Financial Transaction Infrastructure) network, to provide the vendor with access to additional market centers among North American and international exchanges connected to SFTI, broadening Barchart's datafeed services, while also providing another resilient network option for clients to connect to the vendor.

"Our connectivity to SFTI... will allow us to grow into new markets and build a more diverse client base. In addition, with SFTI's numerous access points, it is easier for us to serve clients with low-latency requirements," says Barchart president Eero Pikat. n

Selerity Adds Low-Latency Earnings Date Feed

Low-latency events data provider Selerity has released a datafeed of earnings date announcements for over 3,000 US corporations, to support equity options traders wanting to mitigate risk or exploit short-term price volatility caused by the impact of companies setting or changing the date for announcing financial results.

Jeff Otten, executive vice president of sales and business development at Selerity in Chicago, says that after introducing a feed of dividend announcements in the first half of last year, options traders began asking for a similar feed of earnings date announcements to factor into their pricing models, so they can mitigate risk by setting new prices quickly in response to date announcements, and to generate alpha based on the effect of an earnings date on monthly volume in a contract.

Daunting Task

"In the past, traders would look at press release wires and data terminals, and manually create alerts off those. But if you're looking at a lot of companies, that becomes a daunting task," Otten says. "So to have an automated feed to get that data into their pricing engines is important to them."

The vendor began developing the earnings date feed late last year, and tested data for an initial selection of about 250 companies with highly-traded options with a group of beta clients in the first quarter of this year, then—after validating the accuracy of the data—scaled it to support more than 3,000 names, Otten says.

Selerity sources the earnings date announcements from press releases, companies' websites, and Edgar filings—all of which the vendor already captures to generate its earnings and dividends feeds—and extracts relevant information, such as company name, ticker symbol and announcement date, then delivers that to clients in the form of a low-latency feed. n

Warsaw Stock Exchange Launches New Datafeeds

The Warsaw Stock Exchange has launched a series of new market data feeds in conjunction with its recent migration to NYSE Euronext's Universal Trading Platform (UTP), which officials hope will attract automated trading firms to its equity and over-the counter bonds markets.

After switching to the UTP platform in April, WSE now offers four datafeeds—one that provides the five best bid and offer prices; one that delivers full order-book data for equities, bonds, derivatives, other exchange-traded products and indexes traded on its own markets; and two feeds of data from WSE majority-owned OTC bond market BondSpot.

As part of NYSE's UTP system requirements, the feeds will be available in NYSE's Exchange Data Publisher (XDP) specification, which leverages the UDP (User Datagram Protocol) multicast protocol, allowing WSE to disseminate market data to multiple local and global servers. Previously, WSE's market data feeds used a propriety protocol associated with the exchange's legacy WARSET trading system, which had been in place since 2000.

"Our feeds will be used by existing members and vendors, but the main new target group is the companies interested in automated trading," says Andrzej Grzywacz, director of WSE's information product department. "This protocol is known for its fast data delivery to users... so due to the low latency of WSE XDP feeds, trading on WSE will be much more attractive for such firms."

Growing automated trading on WSE's markets was a driver behind its collaboration with NYSE Euronext in 2010, when the exchange decided to implement NYSE's UTP system. In May last year, the collaboration also resulted in WSE's market data being made available via NYSE Technologies' Secure Financial Transaction Infrastructure (SFTI) network, providing access to its data to greater numbers of small and medium-sized vendors. n

Burgundy Readies Millennium Feed Switch

Stockholm-based exchange and multilateral trading facility Burgundy has switched to a new market data feed after migrating to a matching engine and data platform provided by the London Stock Exchange's MillenniumIT technology subsidiary on June 3, following its acquisition last year by Norwegian exchange Oslo Børs, which already uses the MillenniumIT platform.

In line with the MillenniumIT specifications, Burgundy is disseminating data via a new FIX Protocol-based feed, though the content and performance of the feed remains unchanged. Burgundy previously offered a Level 2 real-time feed based on Swedish exchange trading and data technology provider Cinnober's proprietary EMAPI multicast protocol, and a unicast feed for less latency-sensitive trading firms and vendors.

The data is distributed by Oslo Børs alongside feeds covering Norwegian equities from its own markets, which were previously based on the LSE's legacy Infolect market data protocol before being replaced by ITCH and FIX FAST-based feeds when it migrated to MillenniumIT in November.

Burgundy chief executive Olof Neiglick says he expects the move will increase Burgundy's member base, specifically among traditional large order flow providers and investment banks, as well as among automated trading firms based in London that are already familiar with MillenniumIT's technology and protocols.

"I have 34 trading members in Scandinavia and London today on [the] Cinnober [platform], but I have five more [that joined] from London in June, simply because we [are] part of the much larger Olso Børs-London Stock Exchange network, so absolutely I expect membership to grow," Neiglick says. n

Interactive Data Bows Hosted Consolidated Feed API

Interactive Data has made its Consolidated Feed (formerly known as PlusFeed) of equities, derivatives, commodities, fixed income and foreign exchange data from more than 450 exchanges and over-the-counter sources worldwide, available via a new API, as an alternative to its existing feed delivery mechanism.

The vendor will host the API in its points of presence in Europe, Asia and the US, to enable users to consume customizable sets of market data from the feed without having to deploy hardware-intensive infrastructure, officials say.

IDC previously leveraged the API delivery mechanism to feed its own desktop products, but has now repositioned it as a standalone datafeed after building out the content over the past 18 months to ensure all venues covered on the vendor's equivalent wire-based feed are supported by the API alternative.

The typical use cases for the API will be among financial institutions and third-party software developers that want to display market data in desktop applications for price discovery, trading, portfolio pricing, risk management and other analytical functions. Firms can customize which datasets are received by individual applications and subscribers, and how often the data is updated without the costs of managing exchange data licensing themselves.

"If you're dealing with a client consuming 100,000 symbols into black boxes at the lowest latency, the wire is the best fit. But if an institution is looking at pushing data to hundreds of terminals in a fairly customized way so that each terminal can consume exactly what it wants, than the API is a far better fit," says Henk D'Hoore, head of product for datafeed services at the vendor's Trading Solutions division. n

SURVEY

Latency Drivers Hit Cost Barriers to Feed Adoption

As the latency of direct and consolidated feeds diminishes—along, in some cases, with the opportunities they create—firms that consume datafeeds face a stark choice: how much are they willing to pay to gain an advantage, and will that advantage come from their ability to capture and respond to data fastest, or the ability to process and analyze large volumes of data? The results of a poll by *Inside Market Data* reveal users' key concerns.

Cost remains a key factor for firms selecting and deploying new consolidated datafeeds, with ongoing budget-tightening resulting from continued economic uncertainty and competition from direct exchange feeds, with end-users demanding services that combine low-latency data delivery with broader consolidated datasets and value-added content.

According to a poll of *Inside Market Data* readers, 42 percent of those surveyed cited cost as their single primary concern when choosing a consolidated feed over other data sources, reflecting ongoing expense management programs at financial firms, with market data departments still under pressure to not increase—and in many cases, actively reduce—their data costs.

In addition to outright purchase cost being a consideration when selecting feeds, respondents also cited economic and budgetary issues as barriers to adoption of new feeds, with 48 percent of those surveyed identifying the cost of deploying and supporting feeds as the greatest barrier to adoption, and 20 percent citing the ability to demonstrate return on investment and justify new spend as the primary barrier. Many cited these as their secondary and tertiary challenges as well, with a total of 89 percent of respondents placing cost of deployment and support in their top three challenges, and 82 percent placing justification and demonstrating ROI among the top three barriers to adopting new feeds, demonstrating that in this budget-conscious environment, every penny of new spend—including projected spend on costs such as integration, maintenance and support—is subject to scrutiny.

Lower on the priority lists were integration issues, and the performance and capacity of firms' in-house infrastructures to handle consolidated feeds, suggesting that data professionals are satisfied with the ease of integration offered by consolidated feeds—50 percent of respondents rated more use of open models and standards for integration as the least important feature to add to datafeeds in future—and both with their own infrastructures and with services offered by consolidators to perform the "heavy lifting" of feed management by normalizing and aggregating raw feeds and proprietary data into an easy-to-digest format without the volume, volatility and microbursts associated with direct exchange feeds—although the ability to deliver exchange data remains a key content concern.

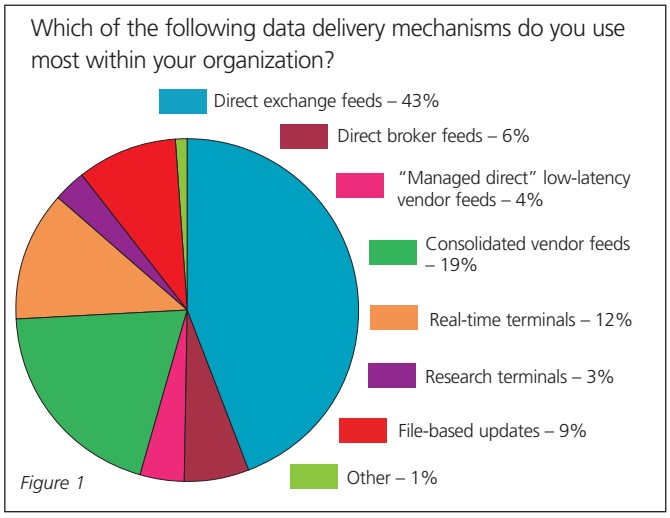

Figure 1

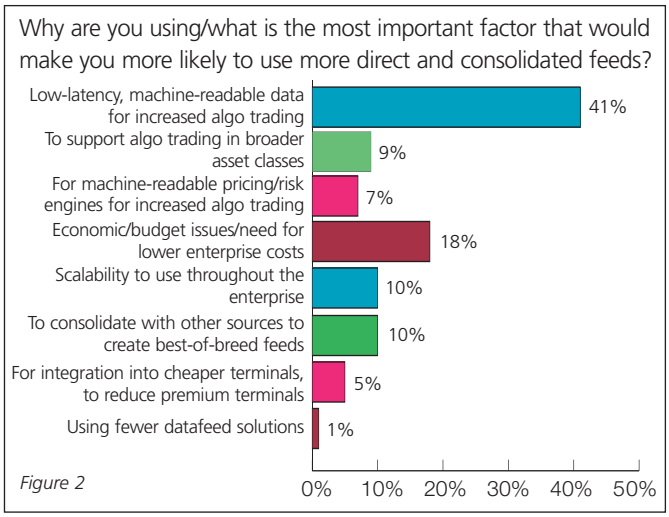

Figure 2

When selecting a consolidated feed, respondents cited comprehensive coverage of exchanges as their second-most important factor overall. And in fact, 70 percent of respondents placed exchange coverage within their top three priorities—only 2 percent lower than the number that placed the same emphasis on cost. Behind these, comprehensive asset class coverage and coverage of over-the-counter sources emerged as mid-level priorities for firms. For example, only 2 percent of respondents cited OTC coverage as their top priority, though 64 percent of respondents selected it as ranging anywhere from their third- through sixth-most important selection criteria.

However, while only 15 percent of respondents cited as a barrier to adoption the ability of feed suppliers to provide the full breadth of data required, a total of 76 percent cited it as one of their top four barriers, behind cost issues.

Yet despite ongoing budget pressures and the often-high cost of direct exchange connectivity, direct exchange feeds remain widely-used through the industry, with 43 percent of survey respondents reporting them as the most-used data sources within

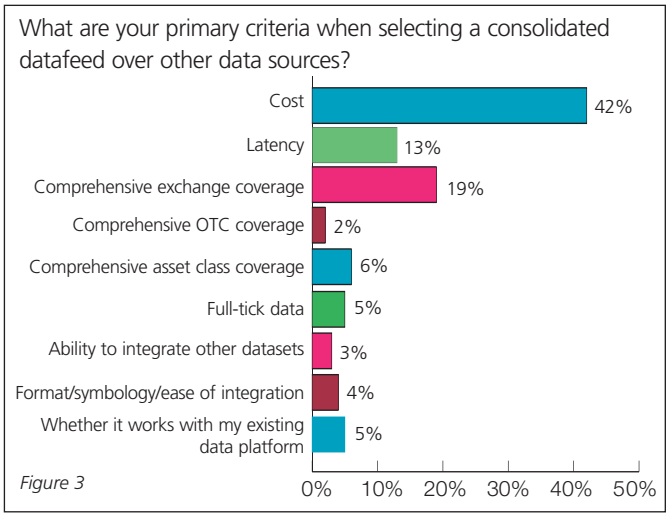

Figure 3

their organizations, followed by consolidated vendor feeds with 19 percent and real-time data terminals with 12 percent. Only 6 percent and 4 percent cited direct broker feeds and the growing breed of "managed direct" low-latency consolidated vendor feeds as their primary sources—though other results of the survey suggest that this last feed type may become more popular in future as firms demand more hybrid direct-consolidated feeds—but these mechanisms both attracted more responses further through the spread of votes, with 57 percent and 58 percent of participants, respectively, citing them as their second- to fourth-most used data source.

Changing Times

The prevalence of exchange feeds and their priority over terminals is in stark contrast to a decade ago and even before the credit crunch and financial crisis, when terminals remained the key source for financial data, especially for over-the-counter asset classes. In the early 2000s, direct exchange feeds were the domain of the early adopters of algorithmic and high-frequency trading, who couldn't stomach the latency of consolidated feeds to power their automated trading models—let alone terminals and the delay of displaying data and relying on a trader's eye-hand coordination to click a trade—while the number of terminals decreased as firms slashed headcount and human traders in response to dual pressures: the emergence of much more efficient trading algorithms, and the onset of the financial crisis.

If anything, these trends have continued—with a brief respite as algorithms pause to figure out their next move, having already largely commoditized equities trading—and direct feeds are increasing their foothold among financial firms.

Though latency ranked only third in terms of top priorities when selecting a consolidated feed—behind cost and exchange coverage—it's a different story when you include direct feeds in the mix. When asked why firms are using datafeeds (both direct and consolidated), 41 percent listed their chief reason as the ability to provide low-latency, machine-readable data to support algorithmic trading, with the secondary pressure being to meet cost pressures by leveraging the economies of scale of enterprise-wide feeds (18 percent), and other latency-related factors, such as to support algo trading in other asset classes and for use in machine-readable pricing and risk engines only the top priority for 9 percent and 7 percent of respondents, respectively—though 52 percent and 60 percent ranked these as their second- through fourth-placed priorities.

Latency continued to be the main driving concern in survey participants' responses when asked what they would want to see made available via real-time consolidated datafeeds in future, with 43 percent saying they want a combination of low-latency, consolidated and delayed data. The second aspect deemed most important was more use of open standards to support integration of feeds from different vendors to reduce vendor lock-in, which 16 percent listed as their top priority—though when it came to picking the least important items on their wish list, 50 percent chose this as being least important. After these, 9 percent cited factors such as volatility data and Greeks, 8 percent each wanted more signals (such as events and sentiment data) and reference data, and 5 percent each wanted index data and evaluated prices delivered via real-time feeds.

Conclusion

Latency continues to be a key driver of datafeed development and adoption. However, with cost issues a more pressing concern—especially as high-frequency traders exhaust the opportunities to be exploited in the markets with the most low-latency data available, and as it becomes harder and more expensive to compete in the "arms race to zero"—budgetary constraints may put a damper on many trading firms' appetite for speed alone, suggesting that future generations of datafeeds will need to diversify their offerings to provide bundled services comprising low-latency, real-time and delayed data via the same infrastructure (suggesting more demand in future for vendor-operated "managed direct" feeds), and carry more content and value-add analytics.

In future, neither speed nor breadth alone will provide enough advantage to keep traders happy: datafeeds will need to provide both—all at an affordable price point to also keep those who control the traders' data budgets happy. n

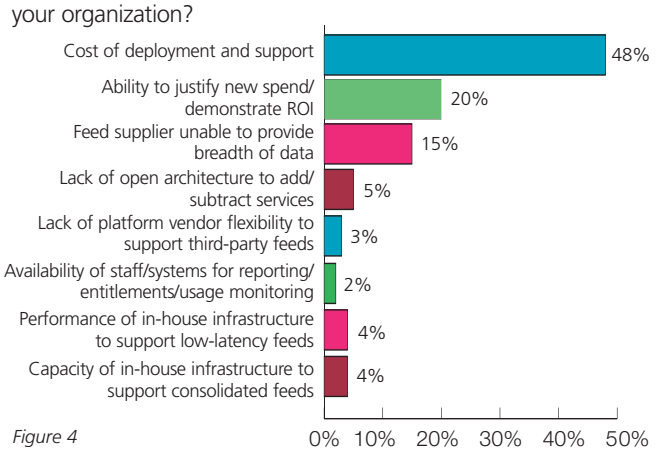

Figure 4

ROUNDTABLE

'Feeding' the Budget Beast: Price vs the Pursuit of Performance

Once the domain of consolidated data vendors, datafeeds have become ubiquitous in the today's world of algorithmic trading, where speed and performance are competitive differentiators. But with so many options to choose from, how can data professionals ensure they are right-sizing their solutions to their needs, and—especially as the highest-performance feeds can come with hefty price tags—are investing in something that will keep them competitive without breaking the bank?

IMD: What are the major drivers of the shift to datafeeds in recent years? How has the recent financial crisis and associated cost concerns contributed to more or less adoption? Has it prompted more enterprise feed projects because of layoffs among traders and the lower cost of a one-to-many feed, or have firms balked at the upfront and ongoing costs of development, maintenance, administration and reporting?

Steven Sadoff, former global head of operations, services and technology at Knight Capital: First and foremost, there is an unprecedented pressure on costs. Both spreads and commissions have shrunk, the cost of capital is becoming more expensive, regulatory demands are only increasing, and trade volumes are down—all of which translate directly toward the drive to lower costs. At the same time, the requirements around market data are increasing on two fronts simultaneously: a larger number of venues, and more data per venue. Specifically, in the US alone there are 13 stock exchanges, 11 options exchanges, and 18 CFTC-Designated Contract Markets. In Europe, there are over 80 exchanges. Even if you ignore the technology aspect of processing different feeds from each venue, it can quickly become daunting on the administration and reporting side—many of these venues perform audits every few years—so even if you are reporting everything correctly, the burden of responding to these audits can become a significant effort by itself. In terms of data per

"The real major driver has been the need for firms to more efficiently consume more content in varying latencies to meet different user profiles, all in the face of real and serious budget constraints. The amount of data that is available and distributed around the globe is astronomical. For years, firms simply took in multiple feeds from multiple vendors to serve all of their internal clients, but what we are seeing today is firms needing to consolidate these resources without losing any of this flexibility."

Brian Cassin, managing director, product and content, and head of Real-Time Solutions, S&P Capital IQ

venue, for the vibrant venues, quote traffic continues to follow a Moore's Law-like curve. To give you a sense of these numbers, for all US stock exchanges, the number of quotes per minute has risen from 1,000 quotes per minute in 1993 to more than 3 million quotes per minute in 2011. On top of all of this, you need to proactively monitor this traffic and deal with the operational aspects of insuring the integrity of all data from every venue. Unless you have a scale business, it's almost impossible to justify the costs of not using consolidated feeds.

If you are a newer entrant to the space, even if cost is not an issue, the time-to-market cost of connecting to many venues will force these players to seriously consider consolidated datafeeds. Just think about how much time it would take to fully set up 100 venues.

Sergei Sinkevich, managing director, direct market access department, Otkritie Capital: The three main drivers in our view are increased reliability of direct data feeds, better latency than that of normalized market data, and lower cost than that of normalized market data (in other words, not having to pay third-party vendors for normalization).

From our vantage point, it appears that the financial crisis itself did not contribute to direct feed adoption. It is more a case of technical progress making it possible for a company to adopt reliable and fast native feeds. Companies are more inclined to use a direct feed than normalized market data, thus saving on payments to normalized solution providers.

Rik Turner, senior analyst, financial services technology, Ovum: The story in the capital markets—starting in the US and then spreading first to Europe post-MiFID 1, and now to other regions—has been one of regulation-driven fragmentation of liquidity, to which the market response is automation in trading systems, which in turn leads to tighter latency requirements and thus a need for direct feeds. A logical conclusion of that process—though by no means a universal trend—is the evolution of high-frequency trading (HFT), which is the ultimate in technical trading, completely devoid of any link to the underlying fundamentals of a stock, and ideally totally market-neutral.

Automation through so-called systematic or algorithmic trading started in—and has largely taken over—cash equities in New York and London, with differing degrees of penetration in other geographies depending on their market infrastructure and sophistication. The global financial crisis may have slowed its advance, but the volatility of the equities markets has itself spurred automation in other asset classes, with foreign exchange an obvious case in point, but increasingly others, too. As Dodd-Frank and EMIR push more swaps through exchanges and central clearing, thus driving standardization at least at the more vanilla end of that market, those instruments too should become susceptible to automation, and consequently Ovum is predicting a gradual spread of lower-latency technologies there, too.

It is certainly true that HFT in cash equities isn't what it was pre-crisis, but very few areas of trading are. HFT was always limited to no more than around 600 companies worldwide, most of them (about 450) in North America, though of course they punch above their weight in terms of the order volumes they generate. I don't see reduced HFT activity resulting in less need for speed in other areas such as investment banks, however, and it is interesting that the real evolution of low-latency techniques—such as wireless line-of-sight links for data transmission—has really evolved post-crisis, or at least after its worst ravages.

Brian Cassin, managing director, product & content, and head of Real-Time Solutions, S&P Capital IQ: The real major driver has been the need for firms to more efficiently consume more content in varying latencies to meet different user profiles, all in the face of real and serious budget constraints. The amount of data that is available and distributed around the globe is astronomical. For years, firms simply took in multiple feeds from multiple vendors to serve all of their internal clients, but what we are seeing today is firms needing to consolidate these resources without losing any of this flexibility.

In addition, the teams that used to manage this content

ROUNDTABLE

"There are a large number of vendors focused on the market data space, so you aren't forced into a one-size-fits-all solution. Depending upon your requirements and appetite for spend, most firms can find a vendor that fits their needs. But, there still seems to be too much hype around low latency, and I get why—all things being equal—a firm would go with the fastest solution, but for a majority of participants, it's overkill and a waste of money."

Steven Sadoff, former global head of operations, services and technology, Knight Capital

and technology are not as large as they were in the 1990s, so they are increasingly looking to their partners and vendors to help them find good solutions to their data and software challenges.

Therefore, it is on us to design smart, tech-savvy solutions that decrease internal costs related to support and infrastructure, making clients better able to focus on meeting the needs of their internal consumers from an implementation and support perspective.

IMD: What different types of feed options are available today (e.g. low-latency and ultra-low-latency direct exchange and "managed direct" feeds, consolidated feeds, conflated feeds) and what are the specific use cases and benefits/disadvantages of each?

Cassin: The industry has really evolved over the years. No longer will one flavor of latency or delivery meet a firm's requirements. Today at every global bank, market data and business leaders are trying to determine the best way to serve their spectrum of internal and external user communities. From deployed dedicated infrastructure to support trading, to conflated and delayed content to support wealth management, they are looking for the best way to optimize limited budgets to meet the needs of these very different requirements.

In this "one-to-many" scalability conundrum, firms are looking for the fewest number of partners that can simplify delivery. An additional level of complexity has arisen over the last couple of years, as sophisticated clients of these banks have started demanding similar informational services for their own needs. As a result, banks must decide between leveraging their existing infrastructure to support these external portals, and upgrading to be able to handle this additional level of reporting and distribution. These trends are driving market data providers to rethink and retool their delivery structures, as missing a piece of clients' "new" requirements can cost business.

Our biggest advantage is technology, infrastructure and personnel, which we acquired through our purchase of QuantHouse last year. QuantHouse's ultra-low-latency, systematic, and algorithmic trading environment was built by traders for traders, and hence has particular expertise in delivering solutions to that market. This has allowed us to look at real-time data delivery in a whole new light. Using the same state-of-the-art infrastructure, we have actually been able to concentrate on the benefits of slowing data down to fit a wider variety of needs. As a result, clients will be able to get the content they need with the latency and structure they require, whether that is real-time, intraday, snap or end-of-day data from a single API.

Turner: There now seems to be an almost infinite gradation of speeds and concomitant price tags, with the HFT brigade obviously requiring the ultimate in low-latency, delivered to systems that are co-located with the exchange's matching engine, and other types of feeds with increasing delays vis-à-vis the event taking place, which feed different trading strategies. It is interesting to see the evolution of multi-asset trading to compensate for the trials and tribulations of equities (not to mention avoiding the table stakes required to compete in the HFT world): while they still like speed, they also require analysis of multiple simultaneous feeds of data on the different asset classes—hence the increasingly mainstream adoption of complex event processing (CEP)—and, since they will usually also be multi-venue in order to trade different asset classes, the kind of co-located race to zero promoted by HFT is less relevant for them.


Steven Sadoff

Sadoff: The good news is that there are a large number of vendors focused on the market data space, so you aren't forced into a one-size-fits-all solution. Depending upon your requirements and appetite for spend, most firms can find a vendor that fits their needs. But, there still seems to be too much hype around low latency, and I get why—all things being equal—a firm would go with the fastest solution, but for a majority of participants, it's overkill and a waste of money.

Sinkevich: The exchanges are making an effort to meet the requirements of all their customers. Small firms that cannot afford to pay much for bespoke development can benefit from TCP FIX Protocol-based feeds. Larger firms are usually offered a faster and more powerful solution named FIX/FAST that is based on UPD multicast technology. And one more option that in most cases outperforms those is raw native feeds. The various exchanges use different terminol-

Inside Market Data | Inside Reference Data | Buy-Side Technology | Sell-Side Technology

waterstechnology

Complete access for your entire organisation to *WatersTechnology*

Our information solutions provide everyone in your organisation with access to the best information resource on the financial market technology industry.

WatersTechnology delivers:

- Breaking news and in-depth analysis on the financial market technology industry
- Detailed features on your market and the deals and people who shape it
- Video features and technical papers
- A fully-searchable archive of content from **WatersTechnology** and from all of the other market-leading titles incorporated under the site (Buy-Side Technology, Sell-Side Technology, Inside Market Data, Inside Reference Data and Waters)
- Full compatibility with any mobile device

To find out more about the benefits an information solutions package would bring, email **solutions@waterstechnology.com** or call **+44 (0) 20 7484 9933 / +1 646 736 1850**

"From our vantage point, it appears that the financial crisis itself did not contribute to direct feed adoption. It is more a case of technical progress making it possible for a company to adopt reliable and fast native feeds. Companies are more inclined to use a direct feed than normalized market data, thus saving on payments to normalized solution providers."

Sergei Sinkevich, managing director, direct market access department, Otkritie Capital

ogy for their native feeds—an option is for large institutional clients who are ready to invest in development in order to achieve maximum results from the technology that each specific exchange can offer. It is also worth mentioning that each participant can develop any technology offered, so even a small client can implement a raw native feed decoder, as all the specifications are openly available to all parties.

IMD: Now that bandwidth is a readily available and affordable commodity, what are some of the other major costs and challenges associated with handling high-volume direct or consolidated feeds, and how can these be addressed?

Turner: Fundamentally, they are related to the normalization, analysis and distribution of the data, though the order of the last two items may be inverted to some degree—i.e. distribute first, then carry out analysis afterwards. In other words, it's about ingesting the data rapidly, carrying out some preliminary analysis once it's normalized, then sending it through the enterprise network to the appropriate place where the trading decisions are made.

That means fast feed handlers in a powerful ticker plant, which—since every data source uses different formats—means buying or developing a feed handler for each new venue you will be trading on, and updating it each time the venue changes something in its protocol.

Cassin: You're right, bandwidth is readily available. But while firms can get bandwidth themselves, often the support and management aspect of that asset becomes cost-prohibitive very quickly. Once you move away from the high-frequency and systematic trading communities that usually require deployed and dedicated infrastructure, hosted APIs and other solutions allow clients to take advantage of providers' scale, support and infrastructure, making the implementation simpler to manage and more cost-effective.

Sadoff: If you need to process every tick across many venues, it's not an inexpensive proposition. Take the requirements for the OPRA feed as an example: 15 million messages per second (1 second interval), 3.9 Gbps (without redundancy), and 25 billion messages per day. Firms that are pushing the envelope have moved to nanosecond-granularity latency measurements. To be able to process this deterministically across the entire day means looking at the entire chain and ensuring there are no weak links. Additionally, you really need to ensure that you have a robust operational performance monitoring capability in place that provides a unified view across all layers, including both the network and application layers.

Sergei Sinkevich
Otkritie Capital

Sinkevich: The main costs are associated with the technology. Yes, bandwidth is available. But what technology is being used to take advantage of this availability? Ethernet via copper wire, or fiber-optic cable? Fiber is far more expensive but has the advantage of fewer hardware delays on the physical level. Next comes the processing unit speeds inside the Network Interface Cards (NICs): Are they 1 Gigabit per second (Gbps) or 10 Gbps? Both options will cope with most of the feeds today. But 10 Gbps does not have a serialization issue that can happen on 1 Gbps, and so is preferable, though slightly more expensive. Next comes the competition within the 10 Gbps product line: One NIC manufacturer claims to decode better than the other. Furthermore, there are manufacturers who take some decoding procedures out of the operating system core down to the NIC processor. That results in fewer operating system interruptions and process management tasks, as well as memory management inside a server. One might say that paying that much to save 50 to 80 microseconds is ridiculous. But it turns out that nowadays, an improvement of 50 to 100 microseconds is too important to pass up.

IMD: Though there are a multitude of feed types and providers to choose from, there are few internal feed distribution platforms for firms to choose from. What technical advances, such as adoption of open architectures, need to occur for firms to be able to exploit the full advantages of a mix of consolidated and low-latency feeds?

Turner: They [platform providers] need to develop the sophisticated strategies, instantiated in software, that can take advantage of a mixture of consolidated (i.e. slower) and direct (i.e. faster) feeds. Once they have done that, everything else is plumbing, which things like the open architectures you mention will of course contribute to improving.

Sinkevich: The advantages of using a mix of consolidated feeds can be achieved by using sophisticated analytical algorithms and a fast reaction to the situation. Say, for instance, that oil prices on IntercontinentalExchange for some reason

"It is interesting to see the evolution of multi-asset trading to compensate for the trials and tribulations of equities: while they still like speed, they also require analysis of multiple simultaneous feeds of data on the different asset classes and, since they will usually also be multi-venue in order to trade different asset classes, the kind of co-located race to zero promoted by high-frequency trading is less relevant for them."

Rik Turner, senior analyst, financial services technology, Ovum

dropped significantly: the algorithm that takes in this information can first benefit in Russia by selling stocks sensitive to the oil price, or even selling the Russian market index before others—who work only with one feed—understand that the environment has changed significantly. Among the technical advances that can help this happen are faster data transmission between different venues, markets, countries and continents; optional availability of universal datafeed standards; proper choice of geographical location; proper choice of hardware components; reliable software implementations; and robust, intelligent and fast algorithms. This is the formula to win.

Cassin: We think it is critical that a solution be capable of publishing to multiple middlewares to simplify the integration. Having to change middleware as part of a solution implementation can lead to significant additional costs, especially if other internal applications are also supported by that infrastructure. We are looking to create more plug-and-play types of implementations, especially with our consolidated feeds. Our goal is to make the API compliant and able to publish to whatever middleware or feed distribution platforms clients might have in-house for downstream distribution to all of their consumers.

Sadoff: I haven't seen a vendor truly embrace an open approach, and I wouldn't hold my breath waiting for a vendor to provide a solution on this front. That being said, I do think there is a tremendous opportunity for a vendor to disrupt the current oligopoly by providing a completely open architecture. I believe the community would truly welcome it, and the vendor would quickly be rewarded with significant market share.

IMD: What will be the next evolution for datafeeds? Will exchanges and consolidated feed providers leverage their direct feed infrastructures to distribute other types of data, and if so, what?

Sadoff: With the client base contracting and the continued client focus on reducing costs, I'm not sure that there is much budget being allocated by vendors on datafeed innovation. Unfortunately, for the foreseeable future, I think it's simply going to be more of the same—incremental improvements.

Sinkevich: I think that as we see consolidation in the exchange industry, they will move toward providing feeds according to a single standard. For instance, the London Stock Exchange and Borsa Italiana have merged, and plan to provide their feed on a Group Ticker Plant technology both in the UK and Italy. As for distributing other types of data via feed infrastructure, my view is that this is not what most professionals are looking for. Specialized infrastructure will always outperform universal solutions.

Rik Turner
Ovum

Turner: Multi-asset would certainly suggest that there is scope for further evolution of datafeeds, perhaps with mixed asset information blended in a single feed. And of course, if the dreaded Big Data hype becomes a reality in the trading world, there will be many additional types of data that could be included for analysis, though many of them will be unstructured and, as such, will require some nifty next-generation analytics platforms to handle them. There is lots of talk of Fast Data in the Big Data world—i.e. stuff that will be big because so much of it will be produced and in such varied, often unstructured formats, but which will still need the kind of speeds that the capital markets are used to for highly structured data.

Cassin: One of the drivers behind the acquisition of QuantHouse was the ability to use its network to distribute non-exchange traded content such as company data, key events and market opinion. There is a lot of proprietary content that S&P Capital IQ publishes in our desktop and enterprise delivery platforms, much of which clients would like to have access to faster or in a machine-readable format. We are now starting to test such content over the low-latency network that would be suitable for systematic traders and others to use in their strategies for alpha discovery. As systematic trading strategies evolve, we believe the focus on speed will be outpaced by the need for differentiated content and other event-driven alerts.

The last thing I'll leave you with: The evolution of exchange market data consumption is still ongoing and unlikely to settle anytime soon. What is important is that suppliers build delivery mechanisms and infrastructures that are open and scalable to support the changing landscape of client needs. The next generation of market data suppliers will need to be nimble from a technology perspective, and robust enough to handle the large volumes of data which the clients will need to power their future applications and strategies. n

SPONSOR'S STATEMENT

How Consolidated Feeds Add Smarts to Speed

As low latency becomes the standard, the ability to deliver "real smarts" through deeper analysis of broader, consolidated datasets becomes the differentiator, says Matthew Spedden, vice president of product and content at S&P Capital IQ.

Consumption patterns for real-time data have significantly changed post-financial crisis, leading firms to rethink their entire data strategies and budget allocations. Different profiles across an organization need access to different levels of market data at different times. Having universal access when not required can add an unnecessary financial burden.

Some users, such as market researchers, may need intra-day, or spot information. Others, such as private wealth managers, may need timely data, but not necessarily to sub-millisecond precision. Yet another group, algorithmic traders, need low- and ultra-low-latency data for executing trades in order to make markets, or to take advantage of price anomalies. We understand that one size certainly does not fit all, though in this article, we will focus on the needs of this last group of consumers.

Regulatory change has further encouraged innovation and technological advancement in the pursuit of a competitive edge. As speed becomes a given, the real advantage comes from being smart. You need to see things your competitors cannot, and better yet, spot opportunities before they do.

As most of today's algorithmic trading is purely a speed game, competitors are pursuing their ideal super-fast data gathering and trade execution systems to take advantage of tiny discrepancies in pricing. With trading volumes down and margins tighter than ever, we are truly in an "arms race" where the winner is the one who is the fastest, pure and simple. The faster the trading system—that is, the lower its latency (the amount of friction or delay in data transmission)—the greater the possibility it can execute a trade.

Only a few years ago, Tabb Group estimated that a five millisecond delay could cost a broker one percent of its deal flow, which can translate to tens of millions of dollars (*The Value of a Millisecond: Finding the Optimal Speed of a Trading Infrastructure*, Tabb Group, April 2008). So, as the markets have become even faster in the five years since that report, who knows how much five milliseconds would cost today?

However, ever-more sophisticated trading systems—including those with co-location offerings, dark fiber infrastructures, and high throughput—will be diminish any competitive advantage brought on by speed. After all, the physical limitations of our universe mean that latency can never reach zero.

But as mentioned before, speed isn't the whole story. Another major component of trade execution is discovery—knowing just what to buy or sell, and when. To do that, you need an incredible level of intelligence on what is happening, and—depending on your strategy—what has happened previously in both correlated and tangential markets. Imagine being able to obtain data lightning-fast and incorporate past performance, historical data, back-testing results, research and other proprietary information to devise new strategies. In this case, technology becomes a catalyst for trading strategy innovation.

To serve this complex trading ecosystem, new systems must deliver a variety of data and content that can be seamlessly integrated into existing third-party and proprietary technologies that are the established backbone for all trading activities. Clearly the daily workflow on the trading floor is more fluid than ever, so it is important that market data vendors adjust to the new realities faced by this particular user profile's specific needs and expectations.

Consolidated datafeeds are not new: data vendors and research firms have been pulling information from various sources and presenting them in one place for decades. But consolidated data, normalized to fit custom user interfaces and delivered instantly, is revolutionary. That said, there are challenges: Consolidated feeds typically aren't "direct" or fast enough, while direct feeds don't have the same levels of normalization, quality checks and value-add data delivered along with the raw price feed. So firms need to choose the right feed for their trading strategy. For example, trading via super-fast algorithms makes sense for discrete changes in equities and currencies, while wealth managers and fixed income traders have little need overall for high-speed trading.

This was one of the drivers behind S&P Capital IQ's recent wave of technology acquisitions. In 2012, we acquired the systematic trading solutions provider QuantHouse to provide traders with a truly fast market data feed. We are already expanding this business into a broader Real-Time Solutions group to leverage this state-of-the-art technology to deliver data at different speeds, as well as for the development of "event driven alerts," utilizing content such as company data, key events and market opinion.

While you can still be the fastest by employing ultra-low latency and high-throughput strategies, there are increasingly streamlined opportunities to be the smartest by utilizing information from diverse sources, and analyzing information using proprietary in-house data and methods. Speed by itself can be an advantage, but with up-to-the-instant market data at your fingertips, you can put the smarts in ultra-high-speed trading. n

Inside Market Data

The new bumper issue

Inside Market Data app now available – FREE for all *Inside Market Data, WatersTechnology Data* and *Premium Package* subscribers.

Download all the content from the latest weekly print issue, build your archive of issues and view them offline.

To download and find out more visit
waterstechnology.com/static/imd-app

An incisivemedia website
incisivemedia.com

Inside Market Data Inside Reference Data Buy-Side Technology Sell-Side Technology
waterstechnology


An
incisivemedia
publication