


Inside Market Data

May 2010

www.insidemarketdata.com

LATENCY

SPECIAL REPORT



Sponsored by:





COUNTERPARTY RISK.
LIQUIDITY RISK. CREDIT RISK.
CAN YOU RISK
NOT USING RAP?

SYBASE RAP – THE TRADING EDITION ENTERPRISE GETS YOU MORE OF THE BEST DATA, FASTER.

With the different kinds of risk in your business, shouldn't you have software that can manage them all simultaneously? Then get the solution that leading firms rely on for proactive risk management: Sybase® RAP – The Trading Edition™ Enterprise. This powerful market analytics platform now includes time series functionality and integration with Sybase CEP, so you can capture, store and analyze vast amounts of market and trade data in real time. That means faster models and smarter decisions. So why risk your success? Get Sybase RAP today.

Get the answers your business demands at sybase.com/RAP

Copyright © 2010 Sybase, Inc. Unpublished rights reserved. Sybase, RAP – The Trading Edition and the Sybase logo are trademarks of Sybase, Inc. or its subsidiaries. All other trademarks are properties of their respective owners. ® indicates registration in the United States of America.

SYBASE®



Downturn Doesn't Delay Latency Efforts


Fast food, fast cars, fast data—speed makes a difference to us all, and one way or another provides an advantage. In the financial markets, when human traders weren't fast enough, firms began developing algorithms that can instantly analyze market data to execute hundreds of trades faster than a trader can execute one. And when they needed faster data, firms shifted their algorithms into co-location centers that host exchange matching engines, to be as close as possible to the source of the data.

Like fast cars, fast data has become an expensive proposition, where hefty price tags don't deter high rollers. Unlike some areas of market data budgets, spend on latency-related issues has continued through the economic downturn. In fact, some—though not all—believe that the financial crisis and added volatility have hastened firms' pursuit of low latency, to make the most of market opportunities.

But as millisecond latencies become measured in microseconds and nanoseconds, the cost and complexity of further reductions becomes greater, while the potential for returns grows smaller. "We may be approaching the point where costs outweigh return. This has led market leaders to look in other directions," says Sinan Baskan, senior director of global financial services industry solutions at Sybase. These could include smarter pre-trade analytics, or unstructured datasets such as news, where depth of analysis becomes a firm's competitive differentiator, rather than wringing ever-decreasing amounts of latency out of application code and network infrastructure or implementing hardware-based solutions.

While these kinds of inputs may not yet be prevalent across the markets, they are gaining traction among firms with very specific needs who believe they have mastered the art of sentiment analysis and who can capture and process different types of information at comparable rates to regular price feeds.

But even now, regulators are examining whether certain aspects of low latency help or hamper the markets overall. Though regulators are expected to focus on risk-related issues, they are also investigating co-location, now taken for granted in the low-latency arena, and any moves to curtail this would force firms to look to other areas of competitive advantage. ■



Max Bowie

Editor, *Inside Market Data*

Inside Market Data

Max Bowie, **Editor**
Tel: +1 212 457 7768
max.bowie@incisivemedia.com

Jean-Paul Carbonnier, **Deputy Editor**
Tel: +44 (0)20 7968 4588
jp.carbonnier@incisivemedia.com

Vicki Chan, **US Reporter**
Tel: +1 212 457 7758
vicki.chan@incisivemedia.com

Lee Hartt, **Publishing Director**
Tel: +44 (0)20 7484 9907
lee.hartt@incisivemedia.com

Jo Garvey, **Commercial Director**
Tel: +1 212 457 7745
jo.garvey@incisivemedia.com

Elina Patler, **Chief Subeditor (New York)**
Brett Gamston, **Subeditor (London)**
Matthew Crabbe, **Managing Director**

Claire Light, **Marketing Manager**
Tel: +44 (0)20 7004 7450
claire.light@incisivemedia.com

Lillian Lopez, **Senior Production Manager**
Tel: +1 212 457 7761
lillian.lopez@incisivemedia.com

Incisive Media
120 Broadway, 5th Floor
New York, NY 10271
Tel: +1 212 457 9400
Fax: +1 646 417 7705

Incisive Media
Haymarket House
28-29 Haymarket
London SW1Y 4RX
Tel: +44 (0)870 240 8859
Fax: +44 (0)20 7484 9932
E-mail: customerservices@incisivemedia.com

Incisive Media
20th Floor, Admiralty Centre, Tower 2
18 Harcourt Road, Admiralty
Hong Kong
Tel: +852 3411 4888
Fax: +852 3411 4811



When you have finished with
this magazine please recycle it.



An
incisive
media
publication

© 2010 Incisive Media Investments Limited
Unauthorized photocopying or facsimile
distribution of this copyrighted
newsletter is prohibited.
All rights reserved. ISSN 1047-2908.

NEWS ROUNDUP

BT Boosts Ultra Net

Network and hosting provider BT is rolling out updated technology to reduce the latency of capturing and distributing data from markets over its low-latency Radianz Ultra Access network to trading firms.

BT began testing the update, a revised router with faster clock speeds and streamlined code, at the New York Stock Exchange in March, where it reduced processing latency to between 15 and 20 microseconds for passing NYSE data through the router onto Ultra Access, from between 70 and 100 microseconds previously, says Michael Masiello, head of industry solutions for Radianz services at BT.

The router provides 10-Gigabit connectivity, for higher bandwidth and throughput, halving latency between some clients' local datacenters and NYSE to 500 microseconds.

Masiello says BT will upgrade Nasdaq by July, and will then upgrade exchanges based on client demand and where the vendor believes there is most benefit to be gained. Next year, BT will install the new router at exchanges in London—although it already uses the router in its London-based proximity hosting centers—before rolling it out at other exchanges in New York. ■

STAC Opens Benchmarking Tools

The Securities Technology Analysis Center is modifying the test environment it uses to benchmark the performance of feed handlers, to enable clients to choose off-the-shelf latency measurement and time synchronization devices from multiple vendors for performing latency benchmarking.

STAC originally developed its test environment using capture cards and time-synchronization equipment from one particular unnamed vendor, but is now working to open up the choice of vendor solutions it can use for measuring latency.

These could include solution providers currently on STAC's benchmark council, such as latency monitoring systems vendors Correlix, Corvil, NetScout, Niksun and TS-Associates, and pro-

viders of time-sync systems such as IBM, Oregon Systems and Symmetricom.

“If a customer has already selected one of these products for latency monitoring or time synchronization, they will be able to use those same tools, along with our tools in a repeatable process,” to conduct benchmark performance tests, says STAC founder and director Peter Lankford.

Re-tooling the test environment may result in it becoming sufficiently open to be compatible with solutions from vendors that are not members of the STAC benchmark council, Lankford says, though he adds that vendors will need to understand the council's benchmark specifications to ensure their products measure latency in a consistent manner. ■

Fixnetix Launches FX Data

UK-based low-latency data vendor Fixnetix has rolled out data and connectivity services for foreign exchange markets, aimed at traders and market makers, and high-frequency firms trading FX as part of cross-asset, cross-border arbitrage strategies.

Banks can use Fixnetix's infrastructure to provide connectivity between their single-dealer trading portals and hedge fund clients, and to distribute price feeds to clients. Fixnetix can also host banks' trading systems in the same datacenters as multi-dealer ECNs, to provide access to multiple FX markets—which banks can in turn offer to their clients.

“There has been quite an upsurge in the number of banks that are streaming to all of these ECNs,” says Paul Harris, FX specialist at Fixnetix. “If you offer price streaming it's very important that you

have the lowest latency, because you don't want to get caught offering stale prices into the market. So clearly latency reduction is very important for them.”

To support its FX data and trading services, Fixnetix has set up co-located connectivity to major electronic FX trading platforms—including EBS, Hotspot, Currenex, FXCM, Lava and FXall—using its existing datacenter hosting facilities across Europe and the US, which are connected via its proprietary 10-Gigabit-per-second network.

At present, Fixnetix uses ECNs' native interfaces to ensure the lowest latency by eliminating any processing overheads from normalizing data. However, the vendor is talking with a number of multi-dealer ECNs to normalize their feeds into a low-latency consolidated feed of best-bid-and-offer prices across multiple markets in the future. ■

Algo Tech Preps VoR Feeds

Algo Technologies, the data and trading technology vendor headed by former Chi-X Europe chief operating officer Hirander Misra, has applied for vendor-of-record status with exchanges and multilateral trading facilities, while the first clients of its low-latency direct feed service prepare to go live.

Exchanges and banks are now benchmarking Algo Tech's feed service—which aims to offer better performance at a lower cost than incumbent low-latency feed vendors—against their existing suppliers, Misra says.

Algo Tech's first phase of connectivity will include direct feeds and co-location to data from US exchanges such as NYSE Arca, Nasdaq OMX, the BATS Exchange and Direct Edge, as well as European exchanges including the London Stock Exchange and Borsa Italiana, NYSE Euronext's cash markets in Amsterdam, Brussels, Lisbon and Paris, Deutsche Börse, Wiener Börse, the Nasdaq OMX Nordic exchanges, and pan-European MTFs Chi-X Europe, BATS Europe, Equiduct, Turquoise, Nasdaq OMX Europe and NYSE Arca Europe.

Algo Tech plans to connect to the SIX Swiss Exchange and Bolsa de Madrid before the end of the Q2 2010, followed by Oslo Børs, the Budapest Stock Exchange, the Prague and Warsaw stock exchanges and Quote MTF, according to client demand. ■



Savvis Eyes Latency with Corvil

Network and hosting service provider Savvis is building out a global service to help clients monitor latencies for trading venues using technology from Dublin-based latency monitoring technology vendor Corvil.

Savvis will use CorvilNet appliances to monitor data traffic across its global network, and to measure end-to-end latencies, since Savvis hosts or provides cross-connected access to up to 20 different equities, derivatives and foreign exchange markets worldwide.

Savvis began working with Corvil last September, and in the first phase of the rollout deployed the vendor's CorvilNet appliances to monitor the latency of data transfer within its financial network, says Corvil

chief executive Donal Byrne.

The second phase of the rollout, which began recently, will see Savvis double the total number of CorvilNet appliances deployed, while also installing the vendor's CorvilClear peer-to-peer latency sharing software, which will allow Savvis to share latency information with clients and provide them with a fully managed latency monitoring service, while also using it to share latency information with firms that already use Corvil to monitor their own infrastructures, giving them a more complete picture of market data latency, all the way from co-location datacenters where trading venues host their matching engines to data-consuming applications on clients' own premises. ■

XtremeData Preps Data Hardware

Database analytics appliance vendor XtremeData is building a PCI Express (PCIe) card with Field-Programmable Gate Array accelerators and network interface card ports for low-latency data distribution between ticker plants and other applications.

The vendor is testing a beta version of the card with a handful of clients, including several large banks and hedge funds, and plans to release a production card in July, says Geno Valente, vice president of sales and marketing at XtremeData.

Connecting the card to a firm's ticker plant server enables clients to capture and distribute data from the ticker plant to other applications within client datacenters, including publish-subscribe messaging

systems, and algorithmic trading or analytics engines.

Algorithmic traders can use the FPGAs on the card to precompute calculations using data from the ticker plant and add those calculated results to the data stream before it reaches the algorithmic engine, so that—for example—a trading system does not have to spend time performing the computation prior to executing a trade.

The PCIe card can also communicate directly with, and deliver data to, XtremeData's in-socket accelerators—which include an FPGA that connects to system resources and two FPGAs to run user applications—to support memory-intensive tasks, such as order-book building or complex calculations. ■

Exegy Boosts Data Distribution with Solace, 29West

St. Louis, Mo.-based Exegy has integrated its Total Connect ticker plant appliance with hardware-based middleware provider Solace Systems' Message Router and messaging provider 29West's Latency Busters Messaging, to broaden its low-latency data distribution capabilities beyond just high-frequency applications across an enterprise.

Exegy and 29West previously collaborated at several mutual clients, but have now optimized the low-latency data delivery from Exegy's ticker plant to 29West's middleware, while Exegy and Solace are building an integrated version of their products and testing the performance of messages from Exegy's ticker plant over Solace's middleware, and expect the integrated versions to be available by the end of this summer, says Jeff Wells, vice president of product management at Exegy.

Although Exegy's latest ticker plant can connect to up to 40 downstream systems such as algorithmic trading engines, the

integrations with Solace and 29West allow market data from the ticker plant to be distributed to hundreds of applications. This would particularly appeal to large banks and broker-dealers that already have an Exegy appliance co-located at exchanges to provide ultra-low-latency data to their high-frequency trading desks, but who now also want to distribute this data to other areas within their firms—for example, for risk management or portfolio valuation, Wells says.

Although connecting applications directly to the ticker plant is suitable for a relatively small number of consumers and applications, Wells says Exegy has seen demand to extend its distribution and enable clients to connect more algorithmic trading engines and trading applications to the ticker plant, allowing the vendor to more deeply penetrate existing clients' infrastructures, while also reaching new markets.

Since 29West's technology does not require hardware in addition to Exegy's ap-

pliance, users can support the data needs of hundreds of end-users or applications with a low server footprint in their datacenters, Meinel says. "End-users get to leverage the super-high-speed market data [for high-frequency apps], and within the same API can do whatever other processing and messaging they need to do in their [other] apps," such as risk reporting and post-trade processing, using 29West's persistent messaging and queuing functions, says Matt Meinel, global director of business development and marketing at 29West.

Exegy also recently upgraded its ticker plant to reduce data processing latency to within two microseconds—depending on specific datafeeds and market activity—from between 25 and 40 microseconds before, and adding features such as User-Defined BBO and Top of Book Quotes features, which strip out specific datasets from Level 1 and Level 2 datafeeds, enabling downstream applications to shift data processing onto the ticker plant's capacity. ■

SPONSOR'S STATEMENT

Latency Through the Looking Glass

With latency now a critical factor in many trading strategies, the ability to measure that speed and assess the amount of investment required to maintain it have taken on new importance. Donal Byrne, chief executive of Corvil, examines the other side of the coin in the race to zero latency—cost.

The US Securities and Exchange Commission's recent concept release on equity market structure has thrown a spotlight on the topic of latency transparency. To many this is not a well-understood concept in terms of what it means or why it might be a concern for the markets.

The SEC definition of latency transparency is "the disclosure of information that would enable market participants to make informed decisions about their speed of access to an exchange or other trading center." This statement in isolation would bring few arguments from traders whose strategies are sensitive to latency. They are already screaming at their IT staff and pleading with their market centers to provide better latency transparency and information. Why? The answer is cost.

The price of zero latency is infinite spend. Even if one is willing to commit to infinite spend (which nobody is), their advantage may be short lived because of the Law of Latency Relativity, which states that "I don't really care how fast I am, I just need to be faster than you."

Therefore, as soon as one or more competing players invest the same amount, the relative speed advantage is reduced or removed. In practice, the decision process for investment in low-latency trading infrastructures is more complex, but the fundamental issue of latency reduction versus cost is an important question. The key to answering this question is latency transparency. In a world of finite ability to invest in low-latency technology combined with the Law of Latency Relativity, information about speed is likely to become more important than speed itself. This is why we say "Speed is good, but transparency is better."

Latency is not a constant: it changes all the time. Latency information published by exchanges and service providers is typically reported at end-of-day, and lacks critical pieces of information, such as load conditions and specifics of how measurements were taken. This approach focuses on demonstrating that the exchange or service provider is performing at low latency, instead of providing latency information that helps market participants to achieve lower latency at the right cost. The latter is the true test of latency transparency and is the model used by CorvilClear, where microsecond-granular latency information is exchanged between trading parties continuously, in real-time.

If one cannot measure or access latency information continuously and in real-time to support fast capture of market prices and execution, then the only choice is to spend increasing amounts of money on low-latency technology and services, and pray it pays off.

Transparency Test

Before you invest in low-latency technology and services, consider the following:

1. **What you see is not necessarily what you get.** Advertised latency performance metrics or benchmarks often do not represent what is achieved in practice. This is largely because your specific operating conditions are different to those that existed at the time of the advertised latency performance. The second reason is that the measurement method and/or technology could be different, e.g. different measurement precision, accuracy or period. The best way to tackle this is to make your own latency measurements for your specific operating conditions.



Donal Byrne
Corvil

2. **Identify where the real latency bottlenecks are.** We often see people make large investments in certain parts of their trading loop, which fail to improve end-to-end latency in a meaningful way. Therefore, no real improvement in fill rate is achieved. An example of this is co-location. Co-location is not a guarantee that trading performance improves. It is a guarantee that propagation delay will be reduced. Therefore, if propagation delay is a major component of overall latency then co-location should provide a good return on investment. If, however, it is not, then don't be surprised if co-location disappoints. The real latency benefit of co-location often extends from the fact that much higher capacity is achieved at co-lo centers compared to extranet access.

3. **Encourage the "help me help you" model for latency transparency.** Most providers of low-latency services use an approach we call "let me show you that I'm fast." Most consumers of low-latency services don't like this approach and want a model we call "help me understand when you are both fast and not so fast, so I can select the best service and use it better." We believe that a model of "help me help you," adopted by providers and consumers of low-latency data is ultimately the most efficient, where complete latency transparency between parties is required. ■

waterstechnology

Premium pickings

Do you want to cherry-pick the most valuable content and tools for your company in minimal time?

Our brands have long been the choice for professionals within financial-market technologies. From May 2010 our premium content from Waters, Inside Market Data, Inside Reference Data, Buy-Side Technology and Sell-Side Technology (formerly Dealing with Technology) will be available to you via a multi-channel business intelligence platform with daily analysis and news to enable businesses to deliver better strategies with more efficiency than before.


By integrating our five market-leading brands you will get:

- the industry's most respected editorial teams under one web interface
- up-to-the-minute analysis and news – uploaded directly from our editorial desks
- easy navigation – articles linked by topic across all five brands
- free content from our extensive selection of special reports and webinars
- comprehensive news alerts delivering premium business content
- consolidated events and training calendar
- one community to network and interact with
- a multi-level premium subscription service to serve your company's needs – individually, departmentally or globally

Coming soon **waterstechnology**

Premium content for financial-market technologies

If you want to know in advance what you can gain from a corporate site licence to the entire waterstechnology platform contact pat.courtney@incisivemedia.com



The Long Road to Zero Latency

Despite the financial crisis, trading firms have continued to invest heavily in low-latency data solutions to support algorithmic and high-frequency trading, reducing latency levels even further, but requiring firms to invest in new technologies to get the most out of the data. But how are firms handling low-latency feeds of other types of content, and what impact will regulators' investigation of market structure have on latency?

IMD: How have the levels of latency changed over the past year, and what impact have the financial crisis and volatile markets had on firms' ability—and desire—to invest in low-latency solutions?

Lawrence Hansen, director of product management for market data, Lime Brokerage: Overall latency for many market participants has continued to decrease due to technological enhancements across the various segments of the client-venue communication path—for both data and trading. The combination of the financial crisis and volatile markets have brought with them an even greater need for speed across all types of strategies, not just in traditional high-frequency trading strategies such as virtual market making. In today's environment, statistical arbi-

trage, exchange-traded fund arbitrage, algorithmic trading and options trading among others, require faster market signaling, more robust processing speed and near-immediate dispatch to venues to stay ahead of the queue and ensure the profitability of these strategies. Volatility and increased competition bring with them ever-more fleeting opportunities, so strategies with longer time horizons now require the same low latency as microsecond-level strategy horizons. Every firm that trades in the market owes it to themselves and their investors to reevaluate the mechanics and latency of their trading activity. The financial crisis has had the direct impact of compelling firms to more closely evaluate where both macro and micro levels of latency exist in their infrastructure, and the impact on how they wish to trade, rather



“The combination of the financial crisis and volatile markets have brought with them an even greater need for speed across all types of strategies, not just in traditional high-frequency trading strategies such as virtual market making.”

Lawrence Hansen, director of product management for market data, Lime Brokerage

than simply rushing into an expensive infrastructure overhaul. If building the low latency to compete is too costly, especially in light of the current fiscal constraints, there are providers in the marketplace such as Lime Brokerage who have made the necessary investments in infrastructure which these firms can quickly and easily tap without adding too much to the bottom line.

Donal Byrne, chief executive, Corvil: In general, we have seen a significant overall reduction in latency levels within the trading environment. One to two years ago, everyone was talking about milliseconds. Today it is microseconds. Tomorrow, it will be nanoseconds. Investment in low-latency technologies seems to be one of the few areas in banks where growth was experienced [over the past year]. In many respects, the existence of highly volatile markets was positive for many of the trading strategies used in high-frequency trading.

 <p>Donal Byrne CEO Corvil Tel: +353 1 859 1001 db@corvil.com www.corvil.com</p>	
--	---

Mark Skalabrin, chief executive, Redline Trading Solutions: Only a few years ago, low-latency market data systems were considered state of the art if they were processing data in less than a few hundred microseconds. Trading technology has rapidly improved over the past couple of years to where consistent single-digit microsecond performance is required for latency sensitive applications such as smart order-routing and high-frequency trading. In addition, the number of other applications where latency improvements increase profitability has expanded as market infrastructure has improved. These infrastructure changes have come from exchanges and technology providers aggressively competing to deliver better latency performance. What we have seen is that after the shock of the financial crisis wore off,

many firms aggressively moved to improving latency with a clear objective to look outside their firms for leading-edge solutions.

Adam Honoré, senior analyst, Aite Group:

Latency has been an arms race that has not really slowed down through the crisis. Aite Group has seen consistent technology spend for firms that require the lowest latency, despite budget cuts in other areas. Actually, because of the financial crisis and market conditions, I would argue that latency has become a two-tier battle. First, to get faster. Second, to do more processing without getting slower. That second tier is driven largely by increasing sophistication of trading styles. Lastly, now that markets have returned, firms are very focused on expanding the battle beyond local markets and traditional asset classes. Expect Asia as a geography and foreign as an asset class to grow in the next year.



Adam Honoré
Aite Group

Sinan Baskan, senior director, global financial services industry solutions, Sybase:

In the wake of the financial crisis, financial firms’ love affair with low-latency technology solutions continues to grow even stronger. The economic turmoil has given firms reason to reflect, fine-tune and accelerate plans to deliver low-latency solutions to the algorithmic/high-frequency trading community. Firms rely on technology for capabilities ranging from high-frequency trading to risk management to derivatives pricing. If anything, that dependence will only grow deeper in the years to come as companies and regulators seek technological answers to portfolio and systemic risk.

As complex event processing technology continues to advance, levels of latency continue to decrease. Some algorithmic trading infrastructures boast 400 or 500 microsecond response times, resulting in near-zero latency. That’s about as close to “real-time” data as you can get.

John Heflin, senior vice president, global operations, Interactive Data 7Ticks:

We have noticed that latency levels have dropped in almost every part of the technology stack: the telecom circuits, hardware, and software layers. The financial crisis hasn’t necessarily led to a desire to invest in low-latency solutions, but we have noticed that it has fueled a desire to control costs, and we’ve seen an uptick in demand as a result of that. We also believe that volatility equates to trading opportunity. If there’s a lot of opportunity, firms want to get from price to order quickly, and that can drive the need for low-latency solutions.

Steven Sadoff, chief information officer, Knight Capital:

Market data volumes have continued to follow a Moore’s Law-type curve, and most marketplaces are continuing to invest in

ROUNDTABLE

the arms race to lower latencies in a similar fashion. However, I don't believe the financial crisis had any significant impact with respect to this race to zero.

IMD: How can firms take advantage of the current environment to gain a latency advantage over their rivals?

Heflin: Firms can take advantage of the current environment across the stack. At the hardware layer, network hardware vendors seem to always be coming out with new or improved products that can improve performance. In addition, hardware accelerators such as FPGAs and kernel bypass cards are continuing to reduce latency. At the software layer, there have been improvements in areas such as messaging and ISVs that seem to be constantly improving the latency profiles of their core software processing. For firms not using co-location, the telecommunications industry also seems to be constantly improving the latency of its circuit routes. Regardless of where you are in the technology stack, everyone looks to be racing to get as close as possible to zero latency.



Hansen: The current environment has given rise to a whole host of vendor utility solutions that focus on specific pieces or segments in addressing the drive towards lower latency among buy-side and sell-side participants, execution venues and related service providers. Many of these vendors are new entrants taking advantage of current demand but without applying the appropriate scientific rigor or having an established track record to address these issues effectively. Others have been focusing on the issue for a number of years preceding the recent buzz on the topic. These firms typically offer monitoring solutions or other critical functionality to only a portion of the end-to-end low-latency stream for both market data and trading. At Lime Brokerage, we constantly evaluate the market versus what we have engineered internally, but as an agency broker and high-throughput, low-latency technology provider we have implemented an end-to-end solution which we provide to our clients. In recognizing the current financial and regulatory environment, we have brought to market products such as our Citrius market data solutions, and our LimeInside co-located sponsored access solution to drive our own and thus our clients' latency to the lowest levels attainable.

“The challenge today is that many competing rivals are all doing the same thing, and therefore the latency advantage that one may have had 12 months ago has eroded significantly.... As costs exalate, we believe that information about speed is where the next level of competitive advantage lies.”

Donal Byrne, CEO, Corvil

Skalabrin: Today the fastest solutions exploit co-location, hardware acceleration, low-latency switching, and dedicated network connectivity. For most firms, it is not possible to make evolutionary changes to their existing solutions to get the desired benefits from these technologies. Achieving competitive advantage from latency requires addressing each stage of the trading pipeline with solutions that are designed from the ground up to perform at market-leading levels. Fortunately, there is a growing base of leading-edge solution vendors such as Redline that enable firms to quickly replace legacy approaches to achieve best in class performance.

Byrne: The challenge today is that many competing rivals are all doing the same thing, and therefore the latency advantage that one may have had 12 months ago has eroded significantly. Most serious players employ co-location, direct market feeds, and new generation algorithmic trading platforms to maximize their execution speed. However, as costs escalate, we believe that information about speed is where that next level of competitive advantage lies. Those with the best quality and most accurate information on latency will be able to use it to optimize trading decisions while avoiding the infinite cost trap inherent in the race to zero.

Honoré: Technology focused on speed only gets you so far. You also have to build good models, have a deep understanding of the market microstructure, and keep your eye on the future so you don't box yourself in with an inflexible platform. For instance, if you want to add unstructured data to your model or add new asset classes, will what you build to support today's trading opportunities be able to provide for tomorrow's opportunities?

Sadoff: I'm not sure if the current environment changes the classic formula for success. It comes down to getting the best and brightest people, investing in the latest and greatest hardware, and creating a productive team environment.

Baskan: In today's markets, there is an exponential cost associated with every microsecond saved. We may be approaching



the point where costs outweigh return. This has led market leaders to look in other directions. One direction is to be smarter, and so we see more very high-performance pre-trade analytics. Another direction is to news analytics and analysis of a wider range of data, some structured and some unstructured. It is in these areas that the next latency wars may take place—how much deep analysis can you do to identify trading opportunities that today’s ultra-low latency applications can’t detect?

IMD: Where do the highest levels of latency still exist, and what will the industry focus on this year to reduce latency in these and other areas?

Byrne: In our experience there is no single element within the trading loop, which is the standout latency bottleneck point for all environments. Each customer environment has its own challenges. In one deployment it might be a problem with a feed handler, and in another it might be the FIX gateway or any other myriad of problems that might arise. The focus on latency reduction is a function of the type of environment and business at play. If one is co-located, the network will hopefully be less of a problem. However, as people continue to expand their trading strategies to multiple locations then latency between co-location centers becomes a key issue to consider.

Heflin: It varies greatly, depending on the overall trading solution. In some instances, the trading application is a good place to look. For a firm running a third-party ISV in a co-location situation, the highest latency would more likely exist within the trading application, since the firm’s server is sitting right next to the exchange’s matching engine. However, firms that have optimized their application may look toward hardware acceleration. In addition, the telecommunications circuit is the area of highest opportunity in many instances where firms are not using co-location services.



Mark Skalabrin
CEO
Redline Trading
Tel: +1 781-995-3403
sales@redlinetrading.com
www.redlinetrading.com



Skalabrin: The highest latency events for most co-lo deployed systems are the direct result of traffic bursts. Most trading applications do alright on average but fall far behind when one or more exchanges saturate the network with market data. While these bursts are typically short, they can result

in delays of tens to hundreds of milliseconds if not handled properly. Redline has attacked this issue through hardware-accelerated solutions that deliver the bandwidth to consume bursts without delay. Better end-to-end and point-to-point latency measurements with time-accurate traffic patterns are a key industry focus, and are needed to understand and improve this and other similar sources of latency.

Honoré: From what we see, improperly configured networks and bad code seem to be the latency frontrunners. That said, expect to see more hardware acceleration, more solid-state disk drives, and more in-memory database usage. Also look for commodity hardware upgrades based on progress made by Intel and AMD.

Hansen: The war on latency is being fought on two fronts—internally and externally. On the internal front, firms are realizing that their own network infrastructure—along with the actual physical location of their trading and market data infrastructure—is typically a significant handicap in lowering their latency. Advanced hardware, messaging infrastructure and co-location will continue to be a primary focus for the industry this year. Firms that optimize these will put themselves



Lawrence Hansen
Lime Brokerage

in pole position to use and connect to advanced low-latency market data and trading applications offered by the market. On the external front, there is a significant latency disparity among and between the execution venues, the market data disseminated from each, as well as the ways that data is consolidated. Trading venues are making a competitive effort in their quest to demonstrate their value with respect to execution and latency by providing step-by-step latency measurements on flows into and out of the venues, and providing them as a value-added service. This is a clear sign that firms are becoming more focused on latency across all fronts. Specifically you cannot simply pawn off the issue of latency on the exchanges, as they are all better prepared to show their numbers to you and to your clients.

Sadoff: Typically, the longer an asset class has been trading electronically, the lower the latency. That ends up being a function of the volume, the number of participants, and the cumulative investments applied toward the ecosystem. For instance, when we acquired our foreign exchange ECN in 2006, unlike most equity or futures marketplaces, it didn’t have an event-based feed. So in January 2008 we introduced Hotspot FX ITCH, and now all of our latency-sensitive customers use this feed. We did the same for our fixed income ECN, Knight BondPoint. At the end of last year we finished

ROUNDTABLE

the re-architecture of our market data feeds in a similar fashion. Although all of our fixed income clients are on the new feed, at the request of a few clients, we have turned on throttling for their specific feeds because their infrastructure can't handle the transmission rates of our new ticker plant.

IMD: How do the challenges around sourcing, delivering and processing low-latency, machine-readable news differ from other kinds of data?

Honoré: First, the consistency of sentiment is not nearly as high as the consistency of something like an economic indicator from a trading behavior perspective. Second, some high-performance databases used for testing trading strategies don't support unstructured data. Third, you have to be careful about trusted sources. One bad piece of information can cost quite a bit of money or affect the market—ask Apple and United Airlines.

Heflin: Although the content is different than real-time market price and trade data, the core challenges generally remain the same—getting information from point A to point B. However, an important difference is that there are a lot fewer off-the-shelf solutions that support machine-readable news, compared to traditional electronic trading and market data.

“There is an incremental cost associated with every microsecond saved. We may be approaching the point where costs outweigh return. This has led market leaders to look in other directions.”

Sinan Baskan, senior director, global financial services industry solutions, Sybase

Sadoff: Although similar in many regards on the delivery front, news still lags in terms of models to accurately interpret the data. Natural language processing was one of the first goals when the field of artificial intelligence research was founded in the late 1950s, and although it has advanced, the consensus is still that a computer won't be able to pass the Turing Test anytime in this decade. Hence, only a small percentage of firms seem to be using this data in sophisticated fashions above and beyond simply throttling back models when news is published that is relevant to the securities being traded.

Byrne: One of the biggest challenges is that machine-readable news has to be distributed over large geographies and therefore must solve the problem of precision time management. Measuring and understanding precisely the time at which the news is sourced and when it is received is critically important. Millisecond precision on a global scale is required today, and this is likely to become sub-millisecond in the not-so-distant future.

Baskan: With traditional relational database technology, queries run against static data to find information or to summarize or

SYBASE®

Sinan Baskan

Senior Director, Global Financial Services
Industry Solutions
Sybase
Tel: +1 212 596 1150
www.sybase.com/capitalmarkets



analyze data. While those queries may seem fast in some environments, it doesn't provide the continuous and immediate insight needed within capital markets. Complex event processing sends incoming data through a set of pre-defined continuous queries to immediately process it in real time.

The process of sentiment analysis—and other analytics applied to machine-readable news—is an order of magnitude more complicated than dealing with streaming market data. Sentiment engines, such as RavenPack, Need to Know News, and others, can provide low-latency results, but unfortunately, these are available to all subscribers, so it is incumbent on any trading firm that wants to develop a strategic edge to take in multiple streams of raw and analyzed data, apply proprietary logic across a large set of structured and unstructured data, and make decisions. In our opinion, a combination of complex event processing technology (data cleansing, aggregation, keyword identification, etc.)—coupled with a large in-memory analytics capability—is key.

Hansen: Machine-readable news has not yet gained major inroads within the high-frequency trading segment in a significant way, but is gaining traction with certain types of tailored strategies. Recent products are now enabling real-time analysis of the relevance of news events while providing a variety of related metrics and heatmap-style displays in consistent formats. As with traditional data, quality, consistency and speed are critical, but given the subjective nature of machine-readable news, additional refinements will be necessary to make this a viable source for high-speed trading strategies. Vendors apply their own specific techniques to filter and ensure a consistent quality of what comes through, but it is still just subjective evidence of a trend. While Lime can accommodate its rise as firms tailor their algorithms to use these new sources in the search for alpha, any trading decision from machine-readable news requires the ability to execute as expected, and this will continue to be a function of high-quality, low-latency market data and market access.

IMD: What impact will any potential regulation of high-frequency trading and co-location have on the data industry?

Sadoff: It depends on how onerous the regulation. For co-location specifically, I haven't heard anything on the regulatory front that will have a significant impact. On the other hand, if

Inside Market Data | 25
1985 - 2010

insidemarketdata.com

< 25 >

**YEARS SERVING YOUR
MARKET DATA COMMUNITY**

For 25 years, *Inside Market Data* has been the publication of choice for senior level management. With objective, accurate and timely reporting it provides a concise overview of developments in the fast-moving market data business. *Inside Market Data* was the first, and remains the only newsletter dedicated to these markets. Its unrivalled news and analysis gives our readers an insider's edge.

Inside Market Data subscribers form a community of the most well informed professionals in the industry. Join them today with your own subscription.

**For a FREE 4-issue trial register
at insidemarketdata.com/trial**


Leaders

Incisive Media's market data portfolio incorporates the market-leading industry brands serving financial institutions in print, online and in person – through its extensive series of publications, conferences, research, training, briefings and reports.

Senior level management working in market data at financial institutions rely on our journalists to provide them with objective, accurate and timely reporting of developments in these fast-moving markets – allowing them to understand the changing landscape of their industry and develop and deliver solutions to the challenges their businesses face.

We continue to lead the market in providing business critical information to the financial industry through key brands including *Inside Market Data* and *Inside Reference Data*.

insidemarketdata.com/leaders
irdonline.com/leaders





Steven Sadoff
Knight Capital

they outlaw high-frequency trading, impose a transaction tax or take some similar approach, it could be a significant setback. It has never been a better time to be an investor than right now, especially in markets with significant technology advances like equities. Turning the clock back will not just hurt the high-frequency community, but it will also harm other segments of the marketplace, including the retail investor.

Baskan: There are different possible regulation scenarios, and more than one scenario may apply. Firstly, there may be limits on speed of execution, to slow the markets down so regulators can monitor them in real-time. We think this type of regulation is unworkable and unlikely—trading will simply move to another jurisdiction where low-latency trading is allowed. Another option is to make trading strategies undergo approval before they are implemented. Again, it seems laudable to want to understand the “thinking,” but it is probably not feasible. In the end, we think the issue is more about the viability of the products than the buying and selling mechanism—the crisis of 2008 was due to faulty products and ignorant buyers rather than electronic trading. Even the latest 1,000 point drop in the markets was quickly corrected, and (as of writing) was caused initially by human error.

Skalabrin: It is a little difficult to predict, given the spotlight on high-frequency trading and the wide-ranging potential regulations being discussed. What seems most likely at this point is that there will be increased requirements on risk management, whether accomplished through sponsored access or otherwise. While these types of changes will have some affect on the industry’s structure, they should not materially change how trading applications perform.

Byrne: While it is still very early days with respect to regulation, it does seem that improved latency transparency is an area gaining increased traction with regulatory authorities. The publication and reporting of latency performance by market centers in a transparent and universally accessible manner seems to be preferred over and above ideas that might look to handicap the speed of the underlying trading systems. This will likely create further debate and discussion concerning the area of latency compliance reporting.

Heflin: It’s inevitable that technology will play a key role in helping firms comply with regulation and support increasing market data volumes regardless of what regulations are created. As long as electronic trading exists, it is logical to conclude that technology will be involved in supporting and extending firms’ capabilities.

Hansen: Regardless of the specific aspects of new regulations forthcoming on trading and co-location, the overall impact on the market data industry has been very positive thus far. The current attention in the industry to disparities in the mechanics of market data, its consistency and latency measurement standards, have brought forward an effort among the best providers to self-regulate prior to any forthcoming mandates to do so. Industry efforts such as STAC have been well embraced and have caught the attention of providers and consumers in regards to the standardization of latency measurements. We welcome the industry-wide education that such efforts have produced in helping us promote the perspective of the entire trading cycle value chain, and not just the value of a utility or appliance in a specific segment.

“[Regulators] turning the clock back will not just hurt the high-frequency community, but it will also harm other segments of the marketplace, including the retail investor.”

Steven Sadoff, chief information officer, Knight Capital

Honoré: We hear a lot of noise from regulators and politicians about being concerned for the long-term investor. If they were really concerned about long-term investors, they probably wouldn’t worry too much about low-latency trading. Unfortunately, low-latency trading and high-frequency trading have blended together in our vocabulary. High-frequency traders are generally proprietary firms or market makers. Low-latency trading can be done by anybody looking for better execution with access to technology. Both leverage the same basic technology stack. Senator Ted Kaufman issued some guidelines in March urging regulators to take a more active stance when it comes to technical advantages like co-location.

I think the net impact of the execution technology arms race to the average retail investor is positive, even if they lose best execution to an electronic trading firm every time they trade. According to the Investment Company Institute, the median amount invested by investing households in mutual funds is \$100,000. My research into online brokerage habits indicates the average retail investor trades equities once per month. The average number of shares traded per execution across the market is about 240 shares. Instead of costing investors, innovation in low-latency trading saves them \$471 per year if the retail investor loses a penny per share. Even if investors lose 10 cents per share on every trade, they still make \$212 more per year than they would in a theoretical level playing field. We don’t need to legislate a cap-and-trade-style competitive disadvantage in our local markets based largely on misunderstanding and little on fact. Personally, I hope my fund manager beats me in the market every time I trade. If I ever find out they are not winning best execution against me, I’ll start investing in different funds. ■

SPONSOR'S STATEMENT

Latency and the New Competitive Advantage

As high-frequency trading has become more prevalent, low-latency data has become more important, and the increments by which latency is measured have become smaller, requiring sophisticated tools such as complex event processing software to derive value from an ever-increasing flow of ever-faster data. In today's markets, content is king, but CEP is queen. By Sinan Baskan, senior director of global financial services industry solutions, Sybase

In the capital markets, things happen fast. The markets move fast, positions change fast. Reacting quickly is the key to increasing profit and/or managing risk. And when it comes to reacting quickly, seconds just don't cut it anymore. Milliseconds, and even microseconds, are now the competitive differentiator.

On Wall Street, high-frequency trading by investment banks, hedge funds and other players is one of the hottest areas in the securities industry, reportedly accounting for about half of trading volume on the New York Stock Exchange. For example, high-frequency trading enabled one leading capital markets firm to achieve \$100 million trading days on almost 50 days during a 70-day stretch in 2009.

High-frequency trading relies on low latency, which is the time taken by a package of data to move from computer to computer across a network. Attaining "zero latency" rests on the speed advantage of a machine versus a human. The machines are programmed with algorithms to exploit information microseconds (millionths of a second) faster than competitors' computers, and to complete trades seconds faster than humans. Some algorithmic trading infrastructures boast response times of between 400 and 500 microseconds.

The Real Value

While having access to ultra-low-latency raw data is vital, it is only one part of the equation. The key lies in taking it a step further and being able to effectively

manage this real-time data to provide real value. With vast quantities of data on tap, the ability to capture it, make sense of it, and act on it, without getting overwhelmed or delayed, is critical. After all, the point of low-latency data is being able to quickly analyze it and react to the markets faster than the competition.

The total volume and complexity of time-critical, location-specific market data from outside an organization's boundaries—some of it from third-party providers—can overwhelm even the largest financial services firms' ability to analyze and act speedily enough to maximize profits and returns.

This calls for a new approach to cope with the increased data flow and provide the necessary reaction times. Developing applications the traditional way by storing data first and then analyzing it is a slow process that produces significant latencies that lead to missed opportunities. Given the generally accepted metric that a one-millisecond advantage in trading applications can be worth \$100 million a year to major brokerage firms, there is a lot at stake. So how do firms gain the control and continuous, immediate insight they need to stay competitive?

It all comes down to the ability to analyze data arriving from multiple sources, at very high rates, in real-time. Enter complex event processing (CEP)—a technology that monitors datafeeds from multiple sources and creates a relationship within the data to detect patterns that represent opportunities and problems. Using the raw material




Sinan Baskan
Sybase

of large data streams, the software takes less than milliseconds to filter, correlate, analyze, aggregate and compute incoming data.

In short, CEP handles what is required to act on the data. Once it has identified a pattern, CEP technology instructs other systems to trigger immediate and appropriate responses.

Complex event processing technology implements logic applied to fast-moving data, without sacrificing performance. It can be used by anyone from law enforcers looking to "connect the dots" and identify terrorists, or by traders to make high-speed trades on the basis of breaking news. The common denominator among any CEP-enabled applications is the shared need to continuously collect, process, and analyze data in real-time, producing results without delay, even when the data arrives at very high rates.

Data is the critical pillar in understanding operational risk and exposure. For those companies that have the tools and the systems that enable them to gain insight to support informed decisions and confident actions, the opportunities can be significant. In this time of volatile markets and heightened risk, financial firms are learning to understand, track and quantify risk factors across all portfolios and assets. By mastering the capture and flow of data through fast and powerful time-series analysis, these firms turn information into the new competitive advantage. ■



Speed is good, transparency better

Get clear about latency management: Get CorvilClear™

SPEED – the holy grail of electronic trading. But searching endlessly for those elusive extra microseconds in the dark may not be providing you with the competitive advantage you need. In the complex maze of market data and order execution, a latency bump in the road could surface anytime, any place. Latency bumps can cost you bundles.

CorvilClear™ lets you see, and turn latency into a competitive advantage. For the first time ever, we're opening the doorway to inter-party latency transparency. Imagine a unique peer-to-peer technology for exchanging latency data with microsecond precision in a safe and scalable manner. Imagine minimizing risk and uncertainty

without exposing proprietary information. Imagine precision latency transparency between trading partners, market centers, and market service providers — across all the parties at the core of your electronic trading operation.

Isn't it about time you were able to look beyond your own four walls? Now you can with CorvilClear™ from Corvil, the leading innovator in latency management technology.

For more information, please email us at: corvilclear@corvil.com.



An
incisivemedia
publication