Inside Market Data

March 2008 www.insidemarketdata.com

LATENCY

SPECIAL REPORT

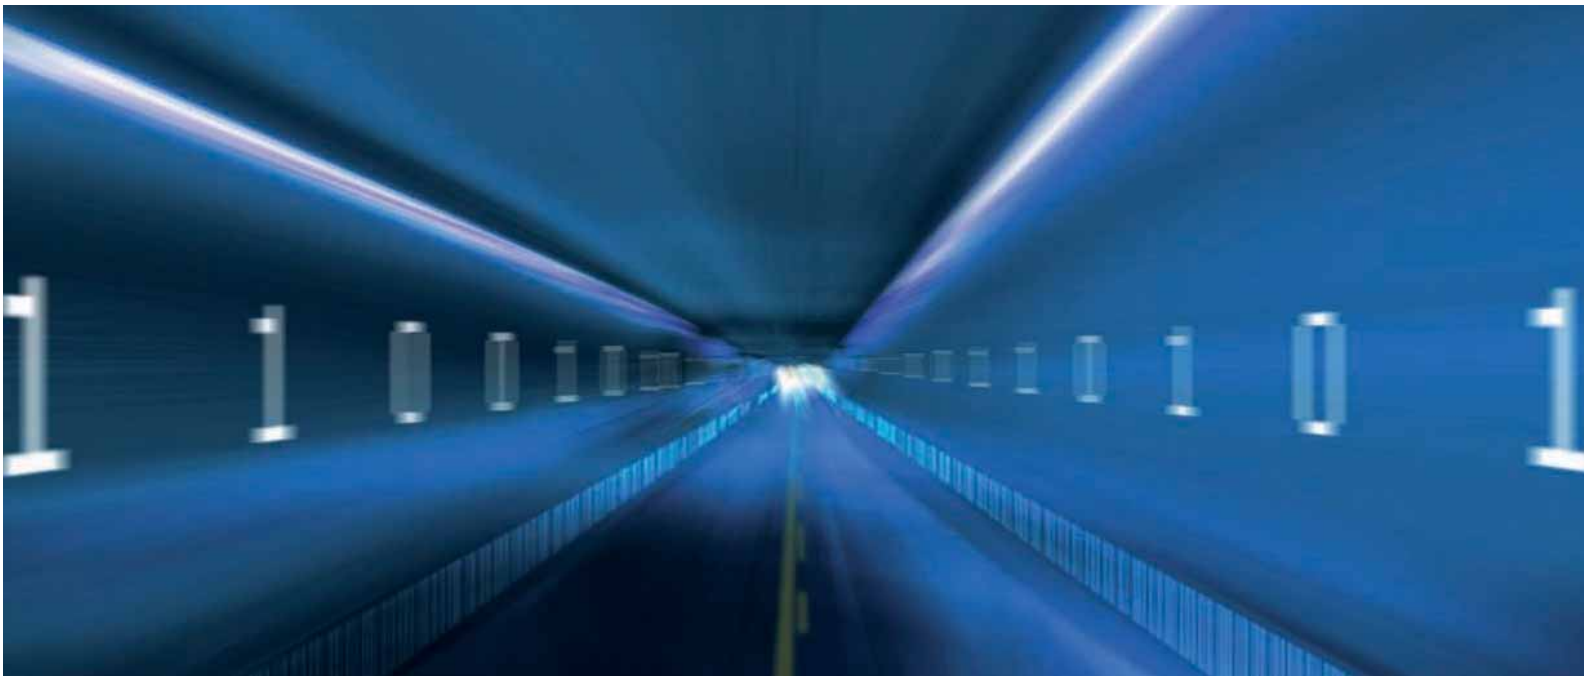

Sponsored by:

Interactive Data | REUTERS | Sun microsystems The Network is the Computer™ | Townsend analytics

LETTER FROM THE EDITOR

Killing Time

Latency—the time taken for a piece of market data to move from one point to another in the trade cycle—has rapidly become the number one data-related issue for firms trading in today's highly competitive markets. But unlike other issues, latency is not something that can ever be solved entirely—it is a constantly shifting environment.

Many projects resemble a 100-yard sprint—or, in reality, more often a marathon—with a defined start and conclusion. Latency, though, has no such "checkered flag." In fact, the finish line is constantly being re-defined as the competitors pick up the baton and set off again in an attempt to beat the original winner and prolong the race—where the "winner" is whoever has the fastest data, the lowest latency and can be first to pick the best prices off execution venues, beating their rivals.

However, in today's markets, a race to the fill such as this occurs thousands of times *every second*. And if the loser shaves a couple of milliseconds—or even microseconds—off their latency figures, they could be the winner next time around... and so on, as each firm inches ahead of each other, only to be overtaken again.

Every link in the trade chain is working on ways to reduce latency—from the exchanges, to the network providers, data vendors and end users, resulting in on-demand exchange feeds, burgeoning co-location facilities and new hardware developments to replace functions traditionally performed in software running across hundreds of servers. And each link is not only pressuring the others to reduce their latency, but is also now looking inward to analyze—and reduce—their own processes that might introduce even the slightest latency.

But part of the problem is that everyone is chasing an impossible dream—to be able to receive and trade on data at light-speed roundtrip times. The real winners in this game will be the firms that not only are first to achieve the lowest-possible latency, but can also figure out what other factors can make a difference on top of that. So, perhaps firms should think of latency as a decathlon, where speed is just one of many competitive assets—and should get in training.

Max Bowie
Editor, *Inside Market Data*

Inside Market Data

Max Bowie, **Editor**
Tel: +1 212 634 4843
max.bowie@incisivemedia.com

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

Elizabeth LeBras, **US Reporter**
Tel: +1 212 634 4809
elizabeth.lebras@incisivemedia.com

Wendy Mock, **Asia-Pacific Reporter**
Tel: +852 2545 2710
wendy.mock@incisivemedia.com

Jo Garvey, **Business Development Manager**
Tel: +1 212 634 4817
jo.garvey@incisivemedia.com

Bernie Simon, **Sales Executive**
Tel: +44 (0)20 7484 9891
bernie.simon@incisivemedia.com

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

Elina Patler, **Chief Subeditor (New York)**
Brett Gamston, **Subeditor (London)**
Vicki Chan, **Editorial/Production Assistant**
Matthew Crabbe, **Managing Director**

Vicky Priest, **Marketing and Subscriptions**
Tel: +44 (0)20 7484 9973
vicky.priest@incisivemedia.com

Lillian Lopez, **Production and Reprints**
Tel: +1 212 925 6990 ext. 134
lillian.lopez@incisivemedia.com

Incisive Media
270 Lafayette Street, Suite 700
New York, NY 10012
Tel: +1 212 925 6990
Fax: +1 212 925 7585
E-mail: customerservices@incisivemedia.com

Incisive Media
Haymarket House
28-29 Haymarket
London SW1Y 4RX
Tel: +44 (0)870 240 8859
Fax: +44 (0)20 7484 9932

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

An incisivemedia publication

NEWS ROUNDUP

Ballista Preps Low-Latency Data Arsenal

Ballista Securities, a new broker-dealer will go live on a hosted ticker plant and multiple North American datafeeds from low-latency data provider Activ Financial when the firm begins operations.

The firm's execution systems are currently in beta-test mode, and have been testing against data from Activ since last October, says Robert Newhouse, chief executive of Ballista in New York.

Ballista has developed a trading platform for trading block options and options volatility, by tying together stock and options trades, so requires feeds from the International Securities Exchange, the Chicago Board Options Exchange and the Options Price Reporting Authority, as well as the CTS and CQS consolidated trade and quote feeds of equities market data, and Nasdaq's NQDS (Nasdaq Quotation Dissemination Service) feed and a feed from its trade reporting system.

"We are very sensitive to latency... because we need to send our crosses [to both the equities and options exchanges] before the market moves-not just [to avoid] monetary loss, but to fulfill our fiduciary responsibility," Newhouse says.

Because of that latency requirement, the firm chose not only to have the ticker and feeds hosted by Activ, but also to co-locate its trading engine within Activ's datacenter in downtown Manhattan. "By co-locating, I save on the extra hop back to my own datacenter, which could be an extra 20 to 40 milliseconds," Newhouse says. Newhouse says that one of the reasons for pursuing this model is to give the firm access to multiple exchange feeds via a single, normalized API, rather than needing to connect to multiple exchange interfaces. ■

Arca Latencies Prompt Multicast Ramp-up

The NYSE-Arca division of the New York Stock Exchange last month rolled out an upgrade to the multicast version of its ArcaBook datafeed to encourage firms still using its legacy unicast feed to migrate to the new version.

Sources say that the unicast feed—which the exchange plans to switch off in July—has experienced "severe" latencies since the start of this year in the face of exponential rises in data volumes.

NYSE-Arca officials blame a sevenfold increase in data volumes since January 2006 for the delays. "We've had delays ranging from a few seconds to 20 or 30 seconds," says Dave Weiss, vice president of connectivity at NYSE-Arca. "We keep throwing hardware at the problem, but the sheer volume of messages is such that this is no longer working."

"Multicast is leaner, has better recovery, and uses less bandwidth," says Paul Adcock, executive vice president at NYSE-Arca. "These unicast setups can't handle these loads... we need to move clients to multicast."

Unicast delivery sends individual data streams to each client on a point-to-point basis, making it more burdensome to manage than multicast, which transmits the same data to all subscribers simultaneously, though multicast can be less reliable than unicast delivery.

Thus, multicast recipients have two lines via which they receive data—a primary and a backup line—so that any data lost by one line should still be delivered by the other. But until last month, the two streams from NYSE-Arca were not synchronized, making it hard to recognize the same data messages on both streams, and causing some firms to delay migrating to the new feed.

According to Weiss, around a dozen firms have already migrated to the multicast feed, with others now looking to migrate. He says there are between 80 and 90 clients of the feed in total. ■

Fuller Enlists Fixnetix Data for New Smart Routing Service

Former Equiduct chief executive Bob Fuller is setting up a new company to provide hosted smart order-routing services for firms that collect low-latency market data from multiple trading venues.

The company, dubbed Exchange Axis, will be an independent utility offering smart order-routing technology to banks and hedge funds trading on European markets.

Fuller has already lined up datafeeds and connectivity to and from exchanges and multilateral trading facilities (MTFs) provided by UK-based low-latency data vendor Fixnetix, of which Fuller is a non-executive director.

Exchange Axis will not redistribute real-time market data. Instead, the service will use low-latency feeds to provide price data to the engine, which will decide where to route each order for execution, based on price and the latency with which each exchange and MTF responds to incoming orders.

"The system will understand the latency of all the venues it's talking to, so if one venue is 10 times slower then another, we'd build that information into the logic of the routing engine," and potentially route orders to venues with faster response times, Fuller says.

Smart order-routing engines will also have to look across other destinations, such as dark pools and systematic internalizers, and will require low-latency data from all venues to feed their decisions.

"When you're using smart order-routing technology, latency is critical," says Alasdair Moore, director at Fixnetix. "If you are 250 milliseconds late, that can make the difference between sending an order to the right destination or the wrong destination. And the problem with sending an order to the wrong destination is that you've shown your hand, even if you cancel it afterwards." ■

Damaged Cables Cause India, Mid-East Latency

Market data in India and the Middle East was disrupted by damaged undersea telecom cables providing Internet connectivity on Jan. 30, though the impact appeared limited to slowing down communications, rather than producing specific outages.

The incident may have been caused by a ship's anchor severing two of the three undersea cables serving the region, reducing the normal 830 gigabits per second of bandwidth to just 205 Gbps.

Trading firms and vendors operating in the region reported only minor service outages, with the disruption mostly taking the form of delays to data transmission.

"We have not received any calls from customers saying that applications are down... [but] there may have been a marginal delay in data delivery," says Pankaj Aher, CEO of Indian data vendor Newswire 18.

Officials at another Indian data vendor also say their services maintained distribution, though the vendor suffered disruptions of several minutes immediately after the incident occurred. However, the vendor adds that communications within the country itself between vendors, users and the exchanges were not affected.

UK-based market data distributor Tenfore Systems, which supplies technology to—and collects data from—India, says its operations were unaffected. "We use a private network to deliver our data to our hubs in India, so we've had no issue," says Richard Barden, head of business development at Tenfore. ■

Xasax Unveils Direct Feed Network

Xasax, a Naples, Fla-based provider of data, proximity hosting and software, is preparing to roll out hosted direct feeds for North American markets, dubbed xsFeed, that will target high-frequency hedge funds and active trader users of trading software from Xasax's sister company, OpenTick.

The vendor's network—points of presence in datacenters operated by Nasdaq and the New York Stock Exchange, as well as third-party hosting centers owned by Savvis and Equinix that host matching engines for various market centers, all joined by dark fiber circuits—has been in production for nearly five months, and will be used to distribute the direct feeds, according to Xasax chief executive Noah Lieske.

The company was created as a result of OpenTick's need to source direct feeds for its trading software. "We had everything below the feed handlers, which we built some time ago for OpenTick," including the vendor's own ticker plant technology and a historical tick database, Lieske says. "We just had to build the feed handlers."

OpenTick previously used a consolidated feed from Nexa Technologies, which Lieske says the vendor will continue to use as a backup feed after it cuts over to the direct feeds from Xasax's network. ■

Reuters Adds Direct Asian Feeds

Reuters has expanded its low-latency Reuters Data Feed Direct service with data from several Asian exchanges, most recently the Tokyo Stock Exchange.

Reuters began providing a direct feed of the TSE's ITA derivatives feed via RDF Direct on Jan. 15, which follows the rollout of feeds from Hong Kong Exchanges and Clearing on Nov. 26, and the addition of the Korea Exchange's KOSPI Futures and KOSPI Options feeds.

Reuters officials say the vendor is currently in negotiations to add 12 more Asian venues to RDF Direct, but decline to comment on a timeline.

The TSE feed provides Level 1 and Level 2 data for bond futures, index futures and options, options on bond futures, equity options and bonds, along with market condition information and data on large orders, large trades and large stop high/low orders, with latency of less than one millisecond. The feed is only available to exchange members.

The ITA feed was originally scheduled to be available via RDF Direct in mid-2006, but was delayed by the exchange, vendor officials say, adding that the Jan. 15 rollout of ITA on RDF Direct—coinciding with TSE's new derivatives trading engine—had been scheduled since Q3 of 2007.

Peter Reilly, regional head of real time enterprise information at Reuters in Tokyo, says that client interest in the TSE feed comes almost entirely from Tokyo-based securities houses for program trading. ■

Trading Metrics Launches Latency Analysis Tool

New York-based software vendor Trading Metrics has unveiled a new product, Market Latency Metrics, for monitoring latency of market data systems and traffic.

The system monitors network traffic to provide latency values for direct and consolidated datafeeds as they pass though points of the firm's infrastructure, such as feed handlers or distributing computers, and records other information that can contribute to latency, such as network load.

"By watching the load... you can see whether a rise in processing time is because there is more information going through a component or over the network," says Trading Metrics founder and CEO Jeff Drew. "You can also look at historical trending to see if there is a pattern of latency increasing at particular times."

Firms can monitor latency performance via a reporting portal that enables clients to run their own reports on the data, or to generate real-time alerts if latency passes beyond client-specified thresholds. These can either be sent by email, for example to a data administrator or system engineer's BlackBerry, or directly to systems management consoles. ■

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

incisivemedia

SPONSOR'S STATEMENT

Adapting to the Changing Low-Latency Landscape

While markets are becoming more consolidated and electronic, the needs of end-user firms are bifurcating, reflecting different levels of demand from each business line, and the need for a broad spectrum of low-latency data options that take other factors into account. By Mark Hepsworth, president of Interactive Data's Institutional Business.


Mark Hepsworth

In recent years, low and ultra-low latency data has become critical for US equity markets and, more recently, markets in the UK and continental Europe. The same drivers of demand in the US—increased competition between markets, and uptake of electronic and algorithmic trading—are spreading to other markets and geographies.

At the same time, market data services are evolving to meet the needs of participants in these changing marketplaces, who are increasingly looking for a spectrum of solutions, governed by particular businesses' needs for low and ultra-low latency data, their desire to manage costs, and the degree to which they prefer to focus on their core business rather than handle the technology themselves.

With the consolidation and "electronification" of the commodities market, we are seeing increasing demand for low-latency commodities data in the US and—to some extent—Europe. Automated trading and very fast, fully electronic execution is increasing, and what we saw two years ago in the US equities market is now happening in commodities. Interactive Data's PlusFeed consolidated global datafeed, which delivers low-latency data from over a dozen US and European commodities markets, as well as Asia-Pacific and Middle East commodities markets, is well positioned to help meet this growing need.

Meanwhile, fueled in part by the roll-out of MiFID and the increasing number of ECNs, low-latency equities data has also become very important in Europe, and is a growing factor in Asia. That's why Interactive Data recently went live with a new European ticker plant, designed to reduce latency by at least 80 milliseconds, significantly decreasing delivery times for European consumers of real-time data from PlusFeed. With the new ticker plant, we are consistently delivering European data to London-based PlusFeed customers at latency of around 60 milliseconds.

Multiple Choices

With the combination of increasing demand for low-latency data and exponential market data volume growth, as well as the cost of managing those rises, firms are increasingly seeking an array of services to satisfy their various needs—for instance, firms may want ultra-low latency data to fuel black-box applications and algorithmic trading, but choose another service for applications that present data visually, or for portfolio pricing.

Clients also have a choice between hosted and deployed datafeed solutions. Many firms see hosted market data services as a way to address both the need for low- and ultra-low latency data and the costs and effort of supporting such datafeeds.

Clients can also choose whether they want a full assortment of data or a more stripped-down dataset. For instance, they can choose a low-latency consolidated feed with value-added information and reference data, or an ultra-low latency direct exchange feed with fewer data points. Or they can decide to receive subsets of certain data, such as our Essential Options Service, which is designed to reduce bandwidth by around 80 percent compared to the full Options Price Reporting Authority feed.

With this range of services, clients can combine ultra-low latency feeds for the most competitive markets, with low-latency data for others, or can choose a consolidated datafeed co-location service to reduce latency and lower communication costs. Today, rather than a one-size-fits-all model, a better solution is often a combination of different technologies.

As the need for an array of ultra-low and low-latency services has grown, Interactive Data has focused on providing the services that firms require:

- **DirectPlus**, a fully managed, ultra-low latency direct exchange feed service with data latency consistently as low as 80 microseconds, allows clients to co-locate their applications alongside DirectPlus at BT Global Financial Services' Nutley, N.J.-based Radianz datacenter.
- **PlusFeed**, a low-latency consolidated global feed providing data on more than five million securities from over 450 sources, includes exchange-based and contributed data, news, corporate actions, and reference and fundamental data.
- The **PlusFeed Co-location Service** enables firms to co-locate applications alongside PlusFeed at the Radianz datacenter. In addition to lowering communication costs, the service is designed to reduce latency between PlusFeed, client applications, and execution venues.
- **PlusFeed Hosted**, Interactive Data's high-speed digital datafeed delivered via leased line or VPN, is designed for financial applications that need real-time and delayed data with reduced bandwidth costs, including applications that require access to streaming watchlists and snap Level 1 data, and applications that need a broad range of exchanges with limited watchlists.

Interactive Data is committed to serving the needs of firms like yours with products for today's highly competitive markets. We welcome your ideas and feedback. ■

ROUNDTABLE

Target: *Light Speed*

Market data latency levels are tumbling as exchanges, vendors and end-user firms all invest fortunes in data delivery and management technologies, in the hope of being the first to trade at the best price. But every valuable microsecond saved means more hard cash spent, and as cutting edge firms inch closer to a light-speed end game, how low can they go?

IMD: What constitutes low and high levels of latency, and how has this definition changed over the past year?

Ambreesh Khanna, chief technologist, Global Financial Services, Sun Microsystems: It is difficult to put into absolute numbers—a lot depends on asset class. But the numbers we are talking about this year are in the single-digit milliseconds to respond to a market condition [whereas] a year ago we were talking about tens of milliseconds. If your algorithmic trading system cannot respond to, say, an arbitrage condition within single-digit milliseconds, you are behind the market this year. I am already having conversations about latencies in the hundreds of microseconds, and technologies that can be used to make software respond in such short time intervals, and tools and techniques that exist to measure and observe causes of these latencies and variances (jitter) from the expected mean.

Kevin McPartland, senior research analyst, Tabb Group: Although it varies greatly depending on the product, latencies have dramatically decreased across the board in the last 10 years. Sending an order from the upstairs broker to a NYSE floor trader's handheld in a matter of seconds was rather amazing in the mid '90s. Only a few short years ago, 50 milliseconds was an achievement, and now we're starting to talk in microseconds. The fact that we are having conversations about hitting physical limits based on the speed of light shows how much the technology has improved. However, in terms of OTC derivatives, we're far from approaching the speed of light, though getting down from days to hours is a major accomplishment in itself.

Interactive Data

Mark Hepsworth
President, Institutional Business
Interactive Data
Tel: +1 781 687 8500
www.interactivedata.com

Mark Hepsworth, president of Institutional business, Interactive Data: We define low latency as less than 100 milliseconds, and ultra-low latency as less than one millisecond. The difference between 100 milliseconds and one millisecond is imperceptible to the human eye or most conventional applications. However, it is perceptible to an algorithm. The most common business metric used by banks to measure the value of ultra-low latency implementations is the fill ratio before and after making the effort. Over the past year, we've seen an increasing number of players who claim to have services in the ultra-low domain, but not many who have been transparent about the tools they use to measure latency. Interactive Data is proactive in measuring both low and ultra-low latency data. We continuously monitor it, and take steps to help keep it within our targeted range by adding more equipment or redesigning software.

"As automated trading has become more and more prevalent during the past several years, the expectations have evolved to the point where traders' algorithms desire market data delivery in a millisecond or less. The emerging frontier for market makers and black-box traders is sub-100 microseconds."

Stuart Breslow, chief information officer, Townsend Analytics

Rick Alm, director, market data strategy and architecture, Credit Suisse: The old conversation was algorithmic-versus-human trading—for humans, 100 milliseconds was acceptable, while for computational trading it was one millisecond. The transition from human-based trading to algo trading changes this dynamic.

Using large numbers for effect, let's look at a scenario of an application that needs time and data to make decisions, where the quality of the decisions are based on the amount of data at hand and how sophisticated the decision-making process can be.
If the data is always delayed by 500 milliseconds from the exchange and the application needs 500 milliseconds to make a decision, then the fastest it can react to an event is 1,000 milliseconds (one second). If you could reduce the latency to one millisecond, then the application can react in 501 milliseconds (half a second). This is where the application (or the developer) can choose to spend another 499 milliseconds making a better decision, or respond to market conditions faster.

These decisions and tradeoffs are now occurring at the microsecond level (one millionth of a second). This makes measuring and quantifying the problem more difficult. But two simple questions can help: The first is whether the data is as close to the exchange as possible, while the second is whether the data is being buffered somewhere. The question isn't so much about how delayed is the data, but what can we do to reduce that value. It's hard to quantify the latency, but it's easy to remove obvious reasons it could be delayed. This is where the focus has been for the last six months.

Robert Newhouse, chief executive, Ballista Securities: The easiest metric is most likely the latency between any single exchange destination and the internal system, and as a second metric—if you use an aggregation system (like Tibco, Activ or Wombat)—what the latency is through the aggregator.

In terms of latency over direct lines to quote venues, while message rates and overall bandwidth needed for quote messaging has increased over the past year, allowable latency has decreased. Whereas the idea of placing a server at an exchange's co-location facility has long been accepted for the purposes of order routing, it is becoming increasingly common to see servers placed at exchange venues to save on the transit time for market data.

We have also seen a general trend in lowering latencies via aggregators. An aggregation solution often classically added up to 40 milliseconds to a quote path. In the last year, we have seen vendors attempt to bring this latency down to more manageable levels with varying degrees of success.

Start Breslow, chief information officer, Townsend Analytics: It depends upon the usage. For example, humans watching screens want the data as quickly as possible, but the reality is that the human eye can't detect changes within 30 milliseconds, and once you factor in human reaction time, market data latencies measured in milliseconds become relatively insignificant. As automated trading has become more and more prevalent during the past several years, the expectations have evolved to the point where traders' algorithms desire market data delivery in a millisecond or less. The emerging frontier for market makers and black-box traders is sub-100 microseconds. Everyone, however, needs to put things into perspective, as a variety of inherent latencies exist that need to be well understood. For example, if a trading algorithm or strategy requires a millisecond to make a trading decision, that will overshadow a predictable 100 microsecond latency relating to market data handling.

Steven Sadoff, senior managing director and chief information officer, Knight Capital Group: Latency continues to be a fast-moving target that is critical to many of Knight's clients. When we acquired ATTAIN ECN's technology in the fall of 2005, order latency was around 75 milliseconds. By last summer when we sold minority equity interests in the ECN, now called Direct Edge, we had lowered it to less than five milliseconds. When Knight acquired foreign exchange ECN Hotspot in the spring of 2006, market data was pushed out every 250 milliseconds. Today, the Hotspot platform pushes market data out in real time, and similarly, order latency on Hotspot has been improved from 100 to around 10 milliseconds.

Steven Sadoff
Knight Capital Group

Kirsti Suutari, global business manager, algorithmic trading, Reuters: Latency has been likened to an arms race, but to be fair, latency requirements are relative rather than absolute. Fastest possible delivery is paramount in highly liquid and fast-moving markets, but where markets are less liquid or slower, latency requirements are not always as stringent. Basically, low latency is as fast as the market, and markets are addressing their own system speeds—so the evolution of low-latency data services has not ended yet.

The concept of latency begins with—though is not limited

ROUNDTABLE

REUTERS

Bill Ruvo
Vice President and Business Manager,
Reuters DataFeed Direct
Reuters
Tel: +1 646 223 4566
www.reuters.com

to—data delivery. Data delivered at the speed of the market will increase the chances that the resultant decision will be timely, though not guarantee it. Speedy inputs will enable a quick decision, but will not ensure an efficient decision processes. So latency is really the sum of all the parts from the inputs to a decision, to the delivery of that decision to the executing market.

Coupled with the concept of latency is capacity. As update rates increase, so must capacity, but without impacting delivery speeds by introducing bottlenecks or queuing, and ideally without increasing hardware requirements to the point of being uneconomical or overly complex.

IMD: As firms look to apply algorithmic trading to new business areas, what effect will this have on latencies and volumes of data for other asset classes? How much of an issue are these topics in other asset classes?

Breslow: As algorithmic trading becomes more widely used for non-equity asset classes, the volume of data for those asset classes will clearly grow. There will be a similar knock-on effect for linked asset classes. We all observed this effect during the past several years with equities and equity options. Lower latency will continue to be sought by those using algorithms, irrespective of asset class, as speed to the market is a clear differentiator for many trading strategies. Data volumes have already begun to show signs of healthy growth. Beyond equities and equity options, which have already demonstrated extraordinary market data volume growth, futures will be next in line for exploding market data volumes as futures trading becomes even more electronic and more mainstream.

Newhouse: It will increase market data exponentially for two reasons. Firstly, trading volumes generally increase with the introduction of algorithmic trading into any asset class, thus increasing quote volumes. Secondly, as algo trading takes hold, loosely correlated trades become increasingly common, such that trading volumes in analogous products may increase in volume as well. A good example of this is derivatives volume increasing as algorithmic trading of equities took hold over the last three to five years, where algos are turned to trade other correlated asset classes in addition to the primary asset class they were built for.

Khanna: Algorithmic trading has a tendency to increase the volume of trading, increasing the ratio of cancellations and quotes to executions. All this generates high levels of market data, which results in more trading opportunities, thus creating a virtuous (or vicious, depending on your perspective!) cycle. Processing requirements are of prime importance. High-speed interconnects, fast CPUs, and tools to observe these trades at very low latencies become critical. As algorithmic trading makes its way into other asset classes, the arms race will continue.

Suutari: The effect of algorithmic trading over time has been to drive fill sizes downward as part of the division of the parent order into many child orders to reduce market impact. Smaller fill sizes have resulted in a commensurate increase in market data update rates. There is no reason to believe that this would not be the case for any moderately liquid asset class subject to algorithmic trading practices. And as machine increasingly competes against machine, the requirement for lower latency in both data delivery and system performance will also escalate in order to ensure competitiveness.

Any asset class with some level of price discovery and liquidity that can be traded via an API has reasonable potential to be traded algorithmically, so while they may not at the moment be subject to latency pressure, capacity issues, and cutthroat competitiveness, these creatures tend to chase their tails.

"The effect of algorithmic trading over time has been to drive fill sizes downward … to reduce market impact. Smaller fill sizes have resulted in a commensurate increase in market data update rates. There is no reason to believe that this would not be the case for any moderately liquid asset class subject to algorithmic trading practices."

Kirsti Suutari, global business manager, algorithmic trading, Reuters

Alm: As the US Equities markets are at an all-time high in terms of volume, and the profitability of execution within these markets are at an all-time low in terms of execution earning per share traded, many banks are looking at improving position and increasing value in other markets. Virtually every other asset class that can be executed electronically is being investigated as a potential new market. The next frontier will be markets that are underserved by the financial community, or where the community goes and creates its own peering arrangements for crossing purposes.

Having learned the lessons from the US options market, newer electronic markets are addressing latency, capacity and throughput from the beginning, which solves some of the existing issues, but opens the door to yet unimagined problems and opportunities.

Solaris™ is open source and free. It runs on IBM, HP (and Sun, too.)

With Solaris™, you can do a lot more. Add reliability and data integrity to your databases. Confidently deploy a secure, scalable Web infrastructure. Plus, you can run Solaris on over 880 x86 Platforms and still benefit from Sun's 24/7 world-class support.

Learn more, download or get your free DVD today at sun.com/getsolaris. And join the Solaris open source project at opensolaris.org.

Call your local Sun Sales Representative, Sun Authorized Partner or (888) 516-9362.

Sadoff: It's a vicious cycle. As asset classes become more electronic, clients are looking for higher performance. When trades happen faster, your typical algorithmic client ends up trading more. When you trade more, you demand even higher performance, and so on. Over time, we expect other higher-volume asset classes to follow this pattern.

McPartland: Latency is an issue across many asset classes, including FX, options, futures and, in some cases, fixed income. As algorithms are used to implement cross-asset strategies, the need for low latency will proliferate beyond equities, as an algorithm is only as fast as its slowest connection. For example, if you are simultaneously trading options alongside the underlying stock, it does no good for your stock order to execute in 65 microseconds if your options order takes 50 milliseconds.

As for data volumes, this is already a major concern and will continue to be so for the foreseeable future. CEP engines will play a huge role in allowing firms to manage and make decisions from an ever-increasing number of high-speed data streams.

"Direct feeds from data sources–mostly exchanges–are no longer a nice-to-have, but a necessity. For any strategy based on low-latency execution, market data from an aggregator and executions via a third-party network just won't suffice anymore."

Kevin McPartland, senior research analyst, Tabb Group

Hepsworth: Low-latency data has become very important in the North American and European equities markets, and is also becoming a significant factor in Asia. We see interest in low-latency data being driven by market consolidation and "electronification." This is happening now in the US and European commodities markets, and firms are applying algorithmic trading strategies to this asset class.

In response to increased demand for low latency data in Europe, we recently rolled out a new ticker plant designed to reduce latency by at least 80 milliseconds for European customers of our PlusFeed consolidated global datafeed. And US customers who trade in Europe also have the option of co-locating their applications in datacenters close to the new London-based ticker plant to take advantage of the low-latency delivery of European data.

IMD: How are low-latency feeds and infrastructures evolving to adapt to faster data at ever-higher volumes?

Khanna: Infrastructures are evolving in contrasting directions. While interconnect technologies are rapidly evolving with lower latency as the primary focus, CPU speeds and memory speeds are not increasing at the traditionally expected rate.

REUTERS

Kirsti Suutari
Global Business Manager,
Algorithmic Trading
Reuters
Tel: +1 416 941 8040
www.reuters.com/algotrading

Moore's Law resulted in almost doubling of CPU frequencies every 24 months in the 1990s and early 2000s. Over the last year—and predicted to hold true in the future—CPU frequencies have remained stagnant. The good news, however, is that more CPU cores are now available in a given package—dual-core and quad-core CPUs are the norm, and Sun even ships a CPU today which has 64 hardware threads on a single CPU. Memory speeds have actually slowed down a bit, though, with the latest memory technology displaying higher latencies than the previous generation.

All this leads to an interesting quandary—how does one crunch through ever-increasing quantities of data if memory and CPU speeds are not keeping pace with the rate of data increases? The answer lies in software—operating systems, compilers and innovative software development techniques.

McPartland: Direct feeds from data sources—mostly exchanges—are no longer a nice-to-have, but a necessity. For any strategy based on low-latency execution, market data from an aggregator and executions via a third-party network just won't suffice anymore. Additionally, the need for hardware acceleration will spread. Although the lowest latency messaging tries to avoid hardware hops where possible, it is nearly impossible to avoid all hardware. Therefore, finding ways to accelerate the needed hardware will increasingly become a point of latency savings.

Newhouse: Infrastructures are evolving in two main ways that I've seen—by throwing hardware at the issue, and transition to hardware-based processing of quote messages. For the former, as hardware becomes commoditized and price per performance continues to fall, one easy way to adapt an infrastructure to higher quite volumes is to simply add more hardware—a paradigm I have seen every shop on the street undertake within the last two years. For the latter, we are starting to see market data aggregators move to a hardware processing model, whereby they are creating specific hardware modules to handle the parsing, with much lower latency and higher throughput than an analogous software solution. A great example of this is the ASIC-based card Activ has released to handle quote parsing on a hardware level instead of via software.

Hepsworth: Many market data firms have become generalists that can provide data for almost every asset class from almost all exchanges and markets. Technical services like middleware and terminals are also geared towards a broad audience of users. But now we're moving towards more focused services that still cover many asset classes and exchanges but have different service levels. Clients are increasingly interested in having a spectrum of solutions, where they can choose a service based on particular appetite for latency, and managing market data volume increases and costs. One option is hosted services. We created our DirectPlus ultra-low latency and fully managed direct exchange feed service in part so clients can off-load the responsibility of taking in feeds directly, keeping ticker plant software up to date, and all the other burdens that go along with direct exchange sourcing. We also offer a co-location service for customers using PlusFeed, our low-latency consolidated global datafeed. This service helps lower communication costs and is designed to reduce latency between PlusFeed and the firm's application, and the execution venues.

Bill Ruvo, vice president and business manager, Reuters DataFeed Direct: Feeds and infrastructures are constantly under scrutiny to ensure they are maintaining performance at the relative pace of the various markets. As these markets upgrade their infrastructures and build new trading engines, this keeps feed handler vendors vigilant. Responsible monitoring and testing means ensuring that performance will withstand the upcoming market data update rates before they are reached—be it due to general trends or momentary spikes. And it means clearing a predictable future level of performance far enough in advance to provide the lead time to make appropriate adjustments. With market data update rates doubling annually, this means that accommodations are generally well ahead of the market. Adaptation can take several forms. Hardware improvements such as InfiniBand have made the same feed handlers more efficient. Where hardware improvements are not an option, optimizing code or applying traffic management are viable alternatives.

"It's a vicious cycle. As asset classes become more electronic, clients are looking for higher performance. When trades happen faster, your typical algorithmic client ends up trading more. When you trade more, you demand even higher performance."

Steven Sadoff, senior managing director and chief information officer, Knight Capital Group

Sadoff: We continue to spend a significant amount of time optimizing software in addition to refreshing and upgrading hardware at a rapid pace. So far, we have been successful at leveraging and customizing distributed commoditized hardware and have been able to avoid dedicated hardware and technologies like InfiniBand.

Breslow: Different solutions are emerging for different usage patterns, and based on the differing capabilities of the users of market data. For example, for the ultra latency-sensitive algorithmic trading users, hardware-based technologies are emerging touting 100-microsecond latency within a 90 percent confidence band for very high data rates. Some users of market data, and some market data distribution companies, have adopted the use of data conflation as data mitigation strategies for usage patterns where not every market data tick is discernible (for example for screen-based users). We're also seeing a much more widespread use of compression as a means of bandwidth conservation, as is evident with the exchanges' evolution to FAST rather than a plain ASCII feed. One of the biggest payoffs that market data users can realize is with a reduction of network hops as market data traverses from server to server. Eliminating routing and context switching latencies results in a big latency save.

Alm: There is a distinct debate in the community as to how to solve this issue. It has become a hardware-versus-software contest of sorts. Many strides have been made in the processing of this data using software methods on commodity hardware. Rather than inject my own opinion into the debate, I would simply ask you to review two other similar situations in the last 10 years.

Cisco Systems was founded as a network routing company that provided a generic hardware platform on which they wrote a software-based routing system. Over time, different hardware-based acceleration methods were introduced into their product line until you have the product that exists today. While an excellent product, and having a software-based "supervisor module," one only has to look at the number of "forwarding engines," "daughter boards," and "content modules" to see that specialty hardware needed to be applied to the problem.

Another quite popular open-source project known as Asterisk is a good example. People realized the ability for a standard

Supercharge your algorithms

This bike may look fast. But without fuel, tuning and skilled control, it's going nowhere. It's the same with your algorithms. They may be the best in the world, but without the **finest data**, **rigorous testing** and the **best technology components**, they won't reach their potential.

Reuters can supercharge your algorithms with ultra-low latency feeds, deep tick histories, machine readable news, analytics engines and testing tools integrated on our market-leading backbone.

To learn more, go to
reuters.com/algotrading

REUTERS

PC to process the voice traffic in a software version of a DSP (Digital Signal Processor) and set about developing a PBX. If you look at the project in its current state, you find many companies that have developed hardware "accessories" for Asterisk. Issues like echo cancellation, compression and conferencing all have hardware "assist" modules available, even though those functions could be done purely in software.

If the financial community is going to keep pace with the volume of traffic, it will have to embrace these ideas. Moving from hubs to switches was industry-changing for network teams, and the same level of change has arrived for the market data community.

"If users require low-latency data delivery, they should focus on reducing the geographic distance from the data source, ensuring appropriate bandwidth to carry the necessary data payload, and then ensure appropriate processing capacity to keep pace with the data rates."

Stuart Breslow, chief information officer, Townsend Analytics

IMD: What part of the market data cycle will the industry address next, and how much potential latency can be eliminated from that part of the process?

Newhouse: I think message normalization and standardization will be addressed in the near term, as evidenced by the introduction of the FAST protocol, etc. I think it's too early to tell exactly how this will pan out, but my hope is certainly that FAST will be normalized to the point where we could expect to see generic FAST parsers implemented in hardware, with the ability to use the devices with quote feeds from any major quote venue.

Robert Newhouse
Ballista Securities

Hepsworth: Along with other vendors, Interactive Data has done a lot of work to help get the data to the client's location expediently. One area where latency can still be reduced is the client's internal market data infrastructure. Addressing this can allow clients to efficiently handle fast-moving flows of market data to support electronic trading. Another area that should be addressed is order execution.

Alm: Poorly performing applications are a significant part of the latency issue. Even the best MDS product suffers when an application makes a huge number of requests in an uncontrolled manner. Developers need to be brought into the design and operational issues of any market data environment so that they can understand the impact of their actions.

The flip side of the coin is that the industry providing the technology needs to embrace the issue. I cannot count the number of times a vendor has told me "It's your network that's dropping the packets," or "This application needs to do this differently." Designers of highway systems do not create roads that a typical car cannot handle at the speed they rated it for. Car manufacturers do not design cars that cannot drive on existing roads.

Why do the providers of the technology insist on providing tools that depend on 100 percent delivery? Adjusting for real-time conditions, allowing for poorly designed routers, switches and congestion events are all items to be expected in the real world. These are the conditions that need to exist in the engineering labs so the products perform in the real world, not just in a clean lab.

The potential latency that can be eliminated is typically discussed as "outliers." Tests are performed and 99.99 percent of the time the product is below a certain latency, with 60 seconds of latency for the remaining 0.01 percent of the time. Why do we applaud this improvement? Next time there is a market event, explain to the trader who lost several million dollars that it was a 0.01 percent-chance event and beyond your control.

Suutari: Eventually there can be no further hardware or application adjustments to improve performance by reducing latency, meaning that physics has become the limitation. When this happens, one of the key areas that companies are looking to address is transport latency—the time taken for data to travel between locations, such as between the exchange and the client's feed handler, loosely calculated at one millisecond for every 100 kilometers of distance. Customers wishing to reduce this form of latency seek to locate their algorithmic trading applications or infrastructures nearer to the market on which they are trading, thus reducing the distance the data travels back and forth. Vendors are accommodating this by offering hosting locations. Local hosting enables firms to trade efficiently on overseas markets without suffering a performance penalty relative to local companies.

McPartland: I don't think there is any part of the cycle currently untouched. Some may be receiving more focus, such as feed handlers and messaging standards, but smart engineers across the financial services sector are looking at the problem in its entirety. Maybe the next issue will be everyone moving their offices back to Wall Street to be near the NYSE.… Additionally, I think we'll begin to see more transparency from the market centers in terms of latency. Speed of execution and how fast market data can be disseminated has already proven a selling point for some ECNs, and will no doubt be a growing differentiator for exchanges going forward.

ROUNDTABLE

Townsend analytics

Stuart Breslow
Chief Information Officer
Townsend Analytics
Tel: +1 312 453 8285
www.townsendanalytics.com

Breslow: There will be pressure on exchanges to disclose and reduce latency within their systems. For example, exchanges timestamp messages differently—some timestamp when a message is created, others when a message is transmitted. The difference discloses some of the exchange's internal latency. Many existing systems are dispersed geographically, resulting in unavoidable latency. Some of these systems can evolve to support co-locating or "close-locating," at least of certain components. An example of this is the fact that at least two of the options exchanges' primary datacenters are several hundred miles away from SIAC, which is responsible for distributing the Opra feed. As a result, there are several milliseconds of latency inherent in some of the options quotes even before they're distributed on the Opra feed.

The data from a feed is not enough to make an automated trading decision. There are many external inputs to trading decisions, including risk management, and reducing latency in these systems will become important also.

Sadoff: Market data continues to follow Moore's Law. Ultimately, the speed of light will be the limiting factor.

IMD: Given the increasing difficulty and expense required to make incremental latency reductions, what other efforts can end users, data sources and data or technology vendors each undertake to improve the flow of data?

Sadoff: Given the significant expertise and investment required to operate successfully in this space, we expect that more firms will outsource all or part of their entire trade execution to companies that have the necessary scale. Knight has been discussing our full-service offering with clients for some time.

Hepsworth: Markets can comply with FISD best practice recommendations and improve their notification policies, particularly regarding changes to infrastructure. If a market decides to go to a new protocol or increase the bandwidth of its feed—which often is necessary as volumes continue to increase—it should give vendors adequate notification so we can prepare our infrastructure and order bigger lines, etc. Unfortunately, some markets don't do this, and it can cause headaches and risks. These systems are complex, expensive and hard to fix, so adherence to good change management practices is vital.

Users can structure their market data requirements based on the needs of particular businesses. For instance, they don't need low-latency data for applications with visual displays or for portfolio pricing, but they probably do want ultra-low latency solutions for their algorithmic trading applications.

For their part, vendors can provide an array of services that offer different levels of latency and content, and that's what we're focusing on at Interactive Data.

Alm: Looking at the US options market for a moment, an overwhelming percentage of traffic is quotes, leaving very little actual traffic related to execution. Methods of improvement around distribution of data, separation of quotes from other classes of traffic at the source, and industry adoption of new standards will be key tools in solving this problem.

The industry needs to embrace standards for distribution. Feed handlers or anything you would like to call "data transformation" are a large component of the latency involved with transporting data. Many firms have dealt with this by foregoing these transformation engines and have brought the raw data into their applications. While providing low latency, the burden is on the application development teams to address each individual change from the source.

FIX FAST goes a long way towards establishing a standard for transporting data, but like other string-based standards, the parsing of the data on arrival adds to the processing time. FIX FAST is good at conserving bandwidth, and is a standard that people can embrace, but does it do anything for routing, speed of processing, etc.? Can you imagine HDTV being sent as comma-separated string values?

Newhouse: From the data source perspective, normalization can be a huge improvement for the industry. If all venues can support FAST or another normalized protocol, it would be quite literally the best thing to happen to market data since TCP/IP came about. From the user perspective, I'm starting to see a bit more done on the mitigation level—that is, having a server parse messages at an exchange co-location facility, then forwarding only applicable quote messages to a lower-level internal core. For instance, if I have an app that needs to see upticks and downticks, I might filter out new trades with no price movement, and not forward those down the WAN.

From the provider perspective, the move to hardware-based processing as mentioned in the second part of the last question seems to be the most focused effort currently.

Breslow: From an end-user standpoint, users can articulate the degree to which low latency is needed—only certain applications or usage patterns require extremely low-latency market data and delivery of every tick update. If they do indeed require low-latency data delivery, they should focus on reducing the geographic distance from the data source, ensuring appropriate bandwidth to carry the necessary data pay-

"The industry needs to embrace standards for distribution. Feed handlers or anything you would like to call 'data transformation' are a large component of the latency involved with transporting data."

Rick Alm, director, market data strategy and architecture, Credit Suisse

load, and then ensure appropriate processing capacity to keep pace with the data rates. If the bandwidth or compute power is insufficient, the end user will experience latencies even if the data is distributed to them flawlessly.

Data sources can help to honor their bandwidth recommendations more closely in several ways. First, data sources could moderate or control their data to ensure that they stay within their bandwidth recommendations. Second, data sources could develop bandwidth recommendations that take into account sub-second microbursts. These microbursts are invisible when rates are measured over five seconds or one second, but they significantly affect latency, either by causing data queuing, or by requiring a re-request of previously transmitted data. Finally, data sources could help to reduce the bandwidth required to process their feeds by helping to suppress redundant data—for example, not sending unchanged values. This is one of the techniques that has helped FAST to achieve its admirable compression ratios.

McPartland: Ensuring the only ones with access to low-latency data are the people and systems truly needing it will often provide quick improvements. Although IT support might enjoy their market data terminal, that terminal does not need a low-latency feed. Doing a full review of usage and other more process-related items such as this is critical. It is not worth spending millions on a low-latency infrastructure without having cleaned up the world around it.

Kevin McPartland
Tabb Group

Khanna: As I had mentioned earlier, vendors are offering systems with more and more CPU cores. The problem is how to take advantage of these additional cores? How does one mitigate the effect of slower memory bandwidth? There have been significant advances in operating system and compiler technologies that allow software developers to scale their applications without any modification. Consolidating application components that communicate with each other over the network using resource containment technologies allows these components to communicate with each at system backplane speeds, while using the additional CPU cores. Recompiling applications with certain flags allows the compilers to insert appropriate directives that auto-parallelize the code, so they can scale across the available CPU cores.

Software developers need to look to operating systems and tool developers to help solve these issues for them. Throwing hardware at the problem does not work any more—it only results in poor utilization, and increased burden on the datacenter in terms of space, power and cooling.

Suutari: Some latency reductions are a matter of common sense, and sometimes the net of a number of small changes can be sufficient to return an operation to competitiveness. Improvements can be the result of:

- Knowing your starting point: Measure and monitor your end-to-end and inter-process latency. Watch for trends, be they inter- or intraday, investigate their cause, and be prepared to address them. Ask about tools that are available to do this and can even become a weapon for arbitraging latency.
- Improved inputs: Swap out legacy data sources for current-version, low-latency direct feeds. As obvious as this may seem, advances in programming and technology of off-the-shelf feed handlers may well have bypassed the performance of older or bespoke versions.
- Co-location: Reduce the distance between you and the exchange by locating your trading application closer to the market.
- Examining your algorithmic trading application: Is it current-generation? Have you exceeded its performance or use case expectations? It may well be serving its intended purpose well, but needs help in offloading some of its more mundane processing responsibilities. Technology is available to do this off the shelf.
- Component interoperability: An algorithmic trading system is a collection of various components. Ensure that they are each optimized, but also that they are integrated as efficiently and effectively as a gold-medal relay team.
- Consistent symbology and interfaces: Data sources that use a common symbol system and API and are already time-synchronized reduce or eliminate translation that can add to latency.
- Configuration optimization: Frequently, your vendor can help you improve the performance of their wares. Talk to the experts to see if your configuration is optimized to the vendor's recommendation and specification, and alternatively, what is required to reach this goal.
- Capacity management: Be certain that capacity is synchronized within and between system components. Bottlenecks and queuing are the enemy of high-performance systems.
- Technology upgrades: Check your hardware and operating system to ensure that you are running the highest-performance version.

Addressing these suggestions is like tuning a Formula One racecar: By upgrading the fuel and adjusting the carburetor, you may find that you had a winning system all along. ■

SPONSOR'S STATEMENT

Outpacing Moore's Law

Ambreesh Khanna, chief technologist in Sun Microsystems' Global Financial Services division, outlines how to reduce latency and jitter by using advances in operating systems.

Market events can create profit opportunities, if one can take advantage of arbitrage opportunities that arise and dissipate within fractions of a second. Responding rapidly to these events is critical. The meteoric rise of electronic trading is testament to this, as computers can respond to events much quicker than humans. Systems that provide information on and respond to external events—such as market data, complex event processing and algorithmic trading systems—must exhibit low latency, defined as the time lag between an event and response.

Typically, increasing the data flow through a system increases the time taken to respond to each data packet. As an example, an eight-socket AMD-based Sun Fire X4600 server with Solaris 10 can handle 282,000 Opra messages per second at 450 microseconds of latency, and 359,000 Opra mps at 511 microseconds of latency. A standard technique for optimizing trading systems is to fix one of the two variables, and push the system to achieve the best result for the other. With this Opra benchmark, the goal was to achieve the highest throughput while keeping latency below one millisecond.

Costs such as power, cooling, datacenter space and network infrastructure contribute a large percentage of the cost of deploying a new application. There is also an executive focus in many companies on their environmental impact, requiring new approaches to support continued volume growth and increased performance, while making better use of power and space by driving up utilization rates.

The Solaris 10 operating system introduced the concept of Containers, which allow individual instances of Solaris 10 to be "installed" and "booted" within a single global instance of Solaris 10. Each of these Containers appear as individual Solaris instances, with their own identity. They are security- and fault-contained, so one non-global Container cannot view data or consume resources from another or propogate faults into another Container.

When communicating between Containers, Solaris determines end-points are on the same physical system and uses the "short-circuited" loopback interface to significantly reduce network latency between processes introduced by physical network hops when processes are deployed on physically-discrete systems.

The chart shows a simple UDP test (TCP numbers follow an identical pattern). The blue line shows latency when communicating processes are on separate systems connected via Gigabit Ethernet; the pink line shows latency over InfiniBand; and the yellow line shows latency when the processes are on two Containers on a single system. The throughput numbers follow a similar pattern—at least an order of magnitude improvement over traditional interconnect technologies.

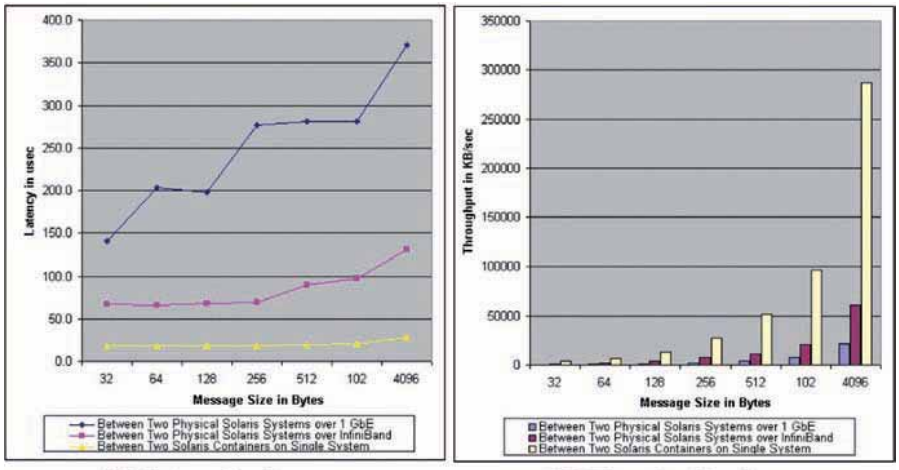

UDP Latency Results

UDP Throughput Results

With the proliferation of electronic trading, minimizing latency across the trading ecosystem is increasingly important. Processes that move large amounts of data should be placed in Containers. For example, implementing a Reuters RMDS system by running RMDS modules such as Source Distributor, Point-To-Point Server (P2PS) and RTIC on a system with Containers would significantly increase server utilization, and increase performance. One can also create Containers that run other applications in the trading portfolio. A data-consuming trading application and the RMDS P2PS module could be run on two Containers on the same system with latencies measured in single-digit microseconds, rather than over the physical network.

Moore's Law almost doubles the number of transistors available to CPU designers every 18 to 24 months. All CPU manufacturers now use these extra transistors to build multi-core CPUs, instead of dramatically increasing CPU frequency. Today, Sun is shipping systems built using a 64-thread CPU. Intel is already shipping a four-core CPU, as will AMD soon. Applications must be built to leverage these large numbers of cores. Sun Studio 12 compilers can multi-thread C or C++ code, if recompiled with the appropriate options. Solaris is a highly scalable OS, with proven scalability across hundreds of CPUs, and technologies like Containers allow applications to be consolidated on multi-core systems if the applications cannot leverage multi-core CPUs.

Many performance assumptions made about trading systems in 2007 are now invalid. No one anticipated today's data rates, which are forcing developers and architects to re-examine their platform choices and invest in technologies that will allow their applications to handle current and future throughput and latency requirements while reducing datacenter footprint. ■

Inside Market Data

insidemarketdata.com

> "*Inside Market Data* provides us with the tools necessary to survive in the market data industry. Its informative content keeps me up-to-date on all the products, mergers, and industry issues. My weekly copy provides me with information on my vendors I didn't know and provides insight into future industry developments."
> *Michael Schwartz,*
> *Lead Associate Market Data Specialist, UBS*

<CONNECT>

WITH THE DATA COMMUNITY

For 22 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.


Inside Market Data

For a FREE 6-issue trial register at insidemarketdata.com/trial

An incisivemedia publication
incisivemedia.com

An
incisivemedia
publication