OVUM

FINANCIAL MARKETS TECHNOLOGY

Monitoring and Measurement for the Capital Markets (Analyst Insight)

Will network monitoring vendors develop application nous?

Reference Code: BFTC2521

Publication Date: December 2009

OVERVIEW

CATALYST

Monitoring and measurement systems for the capital markets are a buoyant sector, with new entrants outnumbering those that have exited the space. Ovum has surveyed the competitive environment to identify the overriding technological trends and discover how the industry is evolving and where the business opportunities are for other tech vendors.

SUMMARY

In this report, the Financial Services Technology team at Ovum looks at the market for monitoring and measurement technology for the capital markets, where the ability to determine the latency with which data is traversing a network and making its way through IT infrastructure is key for trading decisions and troubleshooting, as well as for monitoring how well service level agreements (SLAs) are being met. Its key findings are:

- Buy side, sell side and exchanges all need to monitor latency.
- The FIX protocol is a de facto standard for order flows.
- Proprietary protocols abound in market data delivery.
- Network monitoring tools must address microsecond latencies in the capital markets and
- Measuring application latency means learning new protocols.

Overview

OVUM VIEW

This report shows that there is a firm, ongoing requirement for monitoring and measurement technology in the capital markets. According to estimates from participants around the market, spending on monitoring and measurement in the overall sector is in the $400–500m range, with the high-frequency trading (HFT) segment alone representing around $100m.

As to the future development of this sector, Ovum expects to see vendors that develop greater application expertise, by adding the ability to inspect specific protocols used by the different market data providers and the Financial Information eXchange (FIX) protocol for orders, gain market share as participants in the segment develop increasingly specific ideas about exactly what they need to monitor. Vendors that stick with the more generic network protocols will continue to have a role, however, at least in measuring inter-party latency.

As to which vendors will actually bite the bullet and develop this additional capability beyond the ones that have already committed to doing so, the jury is out. While it is desirable to have such sector-specific knowledge in order to make the product offering a more compelling one, some who do a lot of business in the more generic market for network monitoring may consider it a step too far in the direction of overspecialization to warrant the additional development resources that it will undoubtedly require.

Regarding the overall development of the market, Ovum has noted a marked trend for monitoring vendors to announce trading venues as their customers over the last year, which it attributes to two factors. On the one hand, buy- and sell-side market participants are usually more reluctant to reveal whose technology they are using, making it all but impossible for the vendors to publicize the fact, whereas the venues have a vested interest in trumpeting the fact that they are deploying state-of-the-art monitoring to ensure the best latency through their infrastructure. On the other, competition is increasing among trading venues, particularly, at the moment, in Europe, making it essential that each can demonstrate how low the latency of their respective systems are vis-à-vis their peers.

Ovum expects to see even more latency monitoring systems being deployed over the next couple of years, with the buy side gaining in technical sophistication to reduce its dependence on its brokers, the sell side wanting to show that it is staying ahead to keep buy-siders faithful and the venues using it not only to gauge its own efficiency, but also as a value-add for its customers.

INTRODUCTION

All networks, whether they carry data or voice traffic (a distinction which itself is disappearing with voice over IP technology), need to be monitored to guarantee that they are operating appropriately. If performance has fallen below an acceptable level, network administrators will need to discover what is causing the problem in order to remedy it. If the network is operated by a service provider on an enterprise's behalf, performance criteria will be set in an SLA with financial penalties for failures to meet them, so monitoring becomes even more critical to both parties.

If all this is true in the generic enterprise networking world, it is even more the case in the capital markets. In that sector, the move to electronic trading over the last two decades, followed by the fragmentation of liquidity across multiple trading venues that it made possible, has led to the development of HFT strategies under which companies use automated decisioning systems based on software algorithms to determine the optimum moment to place a buy or sell order, and consequently require the most up-to-date market data.

BUY SIDE, SELL SIDE AND EXCHANGES ALL NEED TO MONITOR LATENCY

This is a market segment in which literally every microsecond counts and, as result, one that invests heavily in monitoring and measurement technology. Traders in the trading rooms need it so they can pressure their IT departments to deliver low latency, the IT departments need it so they can pressure service providers and check that they are meeting SLAs, and of course the service providers need it to prove that they are doing their job.

Moreover, demand is not only coming from sell- and buy-side firms, but also from the trading venues themselves. Theirs is an increasingly competitive world, in which it is necessary for them to demonstrate that an order's passage through their infrastructure is at least as fast if not faster than it is through that of competing liquidity centers (i.e. other venues).

That said, latency monitoring in the capital markets is not necessarily for every vendor. Companies that are making a perfectly good living selling technology to measure latency in seconds or milliseconds on enterprise networks may be tempted by the potential revenue from this most demanding of sectors, but the requirement level will go up alongside the possible profits, so they should think long and hard before venturing into this world of micro- and nanoseconds.

THE FIX PROTOCOL IS A DE FACTO STANDARD FOR ORDER FLOWS

There is also the question of which sector-specific protocols they will need to support (i.e. be able to read and interpret) in order to compete meaningfully in the capital markets. The FIX protocol, which is becoming increasingly prevalent as the industry-standard protocol of choice for transmitting orders, is an obvious first choice, and indeed, a number of companies have already added FIX support into their offerings for this sector (indeed, NetQoS's first entry into the market was with Trade Monitor for FIX).

PROPRIETARY PROTOCOLS ABOUND IN MARKET DATA DELIVERY

Market data, on the other hand, represent a more complex area. There is no prevailing standard protocol for its delivery as, despite the best efforts of FIX Protocol Ltd with FIX Adapted for Streaming (FAST), the majority of providers prefer to stick with their proprietary protocols, not least because they feel this gives them competitive edge. For monitoring vendors, however, this means working with multiple protocols, either through an agreement with the individual providers to get an API for the purpose, or by reverse-engineering to work out how each protocol works.

Corvil is one example of a specialist vendor that has undertaken this work, such that it now supports 80 market data protocols. TS-Associates (TS-A) is similarly focused exclusively on the capital markets and, to this end, its TipOff appliance supports a number of protocols specific to this industry, including RRCP, RRMP, MarketFeed and RWF (all of which are for Reuters' widely used RMDS platform), Tibco's protocols for its Rendezvous product and ActivMiddleware from ACTIV Financial.

One wonders, however, whether the vendors from a generic network monitoring background would be prepared to put in the same kind of highly focused engineering effort for what is, after all, only one of a number of vertical markets in which they operate, even if capital market clients are perceived as having deep pockets to pay for technology to meet their more exacting monitoring requirements.

ANALYSIS

Having established that there is a requirement for monitoring and measurement technology in the capital markets, it should now be noted that, just as latency itself can mean quite different things depending on whether the focus is the network, storage, servers or application, so too there are a number of distinct scenarios in which monitoring and measurement technology is deployed, which have implications for how the vendors in the next section address the market and where they tend to score most of their successes.

NETWORK MONITORING TOOLS MUST ADDRESS MICROSECOND LATENCIES

A lot of the vendors profiled here come from the generic world of enterprise network monitoring, where the requirements are generally less exacting. Latencies of milliseconds, or even of seconds, are often quite tolerable when dealing with the time it takes for a user in a branch office to access an application hosted in the corporate data center, particularly if they are doing so over application virtualization technology such as Citrix's XenApp or Microsoft Terminal Services. This is clearly not the case in the capital markets, and particularly in the growing segment of it that is devoted to HFT.

As a result the first thing such companies have to do in stepping up to the challenges of capital markets clients is endow their technology with the ability to measure down to the micro- and even the nanosecond. An example here is Network Instruments, which in version 13 of its Observer software platform took its latency for reporting and resolution down to 10 nanoseconds.

At this point, the technology can address the needs of monitoring network latency, and indeed it is no coincidence that companies with this background have enjoyed considerable success in meeting this kind of requirement. Corvil, for instance, has made a name for itself in measuring inter-party latency, which is fundamentally a network issue in that what needs to be monitored is the time it is taking for data to travel between one counterparty and another (i.e. across a wire), without any processing or storage going on in the middle. In classic networking terms, it may be thought of as a wide-area network (WAN) rather than a local-area network (LAN) issue, except in the very specific case in which a market participant's algoservers are co-located in the same data center as an exchange's matching engines, in which case it is truly a matter of LAN rather than WAN latency.

MEASURING APPLICATION LATENCY MEANS LEARNING NEW PROTOCOLS

While the network is an important contributor to overall latency, it is only one component, and there is clearly also a requirement to understand what is going on in the application that is processing the data. This in turn means that the monitoring technology needs to be able to decode and interpret the protocol that the application itself is using to carry data. However, this is a capability that generic network monitoring tools lack.

In terms of the Open Systems Interconnection (OSI) reference model widely used to discuss functionality on a network, these tools typically look at layers 2–4, the Data Llayer (Ethernet), Network Layer (IP) and the Transport Layer (TCP). Some of the more versatile ones will also address another layer-4 protocol, the User Datagram Protocol (UDP), which is what underpins the multicast delivery of market data and is increasingly widely used in that sector, particularly in North America. However, the Application Layer of the stack is up at layer 7, and this is not one that network monitoring tools have hitherto been required to inspect.

The situation is analogous, in a sense, to that of network firewalls earlier in this decade, the inability of which to inspect data above layer 4 when security exploits were gaining in sophistication created a market opportunity that the so-called 'application firewalls', also known as 'web application firewalls', emerged to take advantage of, specializing as they did in layer-7 protocols (HTTP, HTTPS, FTP and so on).

Some of the network monitoring vendors have moved to address this additional requirement by adding the ability to inspect higher-level protocols. On the order execution side, at least, their task is made somewhat easier by virtue of the fact that the FIX protocol is fast becoming the norm in such communications, meaning that they only need worry about implementing support for that standard.

Market data, on the other hand, presents a completely different landscape, in that multiple providers are competing to make their feed better, faster and more up-to-date and, as such, they have a vested interest in maintaining an optimized, proprietary protocol to transport their data.

Monitoring vendors thus face the challenge of deciding which market data suppliers to support and carrying out all the technical work to do so. Some have already done so: TS-A has led the way, supporting all the necessary protocols for Reuters, Tibco, ACTIV Financial, 29West and Wombat (now part of NYSE Technologies). Corvil followed suit this year, announcing in its latest version of CorvilNet support for no less than 80 different market data protocols.

It remains to be seen how many other vendors will go down this route. It is certainly not a trivial undertaking and, as such, may be viewed as something of a technological rat-hole by companies with a broad market remit in general enterprise, which see capital markets as a niche.

ITRS, of course, approaches the problem from a different perspective altogether, since it does not come from a networking background: it started life looking at market data feeds and in any case adopts a higher-level view of the monitoring problem, reaching out to an appliance from one of the network monitoring vendors as and when it needs a more granular view of an issue.

COMPETITIVE LANDSCAPE

The majority of companies developing monitoring and measurement technology for sale in the capital markets come from a generic network monitoring background and have determined that the more demanding requirements of this sector, where latency needs to be in microseconds and even nanoseconds rather than seconds or milliseconds, represent a rich vein of business to be mined by vendors with the right product. There are also companies such as Correlix that were actually founded to sell specifically into the capital markets.

For the purposes of this report, Ovum has spoken to the independent software vendors (ISVs) that develop the actual monitoring and measurement products, as well as the hardware developers whose packet capture technology underpins many of the ISVs' offerings when they take them to market as appliances.

It is worth noting here that, while there have been new entrants over the last couple of years such as ClearSight and NetQoS, there have also been casualties during this time. Early in 2008 Reuters announced that it had entered into an agreement with data capture hardware vendor Endace and monitoring software developer Trading Metrics to bundle the latter's product onto the former's and call it the Reuters Latency Monitor appliance, the idea being for Reuters to take it to market as an adjunct to its widely used RMDS market data dissemination platform.

Early this year Trading Metrics went out of business, and its assets were subsequently acquired by Greenline Financial ;Technologies, another start-up in this space, while one of its founders joined Correlix, which around the same time inherited the relationship with Endace as its hardware provider. Meanwhile Reuters is no longer selling the RLM appliance, instead announcing a deal with Corvil to use its CorvilNet product in August.

DEVELOPERS OF MONITORING AND MEASUREMENT SOFTWARE

ClearSight

ClearSight is the most recent generic network monitoring vendor looking to sell its technology in the capital markets, where the demands are considerably more stringent than in conventional enterprise networking. The company has been around since 2001 and its foundation technology is Network Time Machine (NTM), which also underpins its product developed specifically for the financial markets, Cronos, which it launched at the beginning of 2009.

The development of Cronos required three main enhancements of the NTM technology, namely:

- The ability to monitor microsecond latency, since seconds or even milliseconds are far too long for the requirements of the financial markets.
- The ability to detect packet and entire message loss in unicast and multicast environments.
- The ability to detect microbursts, which is when, for a very brief period, network elements such as switches that are sufficient to handle normal levels of traffic are overloaded on account of a sudden dramatic increase in the number of packets traversing the link. Not only do such scenarios need to be detected, but an alert must also be sent to a network admin to execute a script to counter them, halting trading based on the information provided on that link for fear of executing on the wrong data, flushing the buffer and restarting links.

The Cronos product is an appliance that uses proprietary hardware in the shape of specially designed field-programmable gate arrays (FPGAs) to accelerate data capture and ensure that all information is captured, plus a storage subsystem built around redundant array of independent disks (RAID) and serial attached SCSI (SAS) technologies for the longer-term retention of performance data.

In terms of differentiation in what is an increasingly crowded space, the company argues that its ability to see all seven layers of the OSI stack, rather than just the Network Layer and below (i.e. layers 1–3) separates it from some of the competition, although to be fair there are other monitoring players who make a similar claim. ClearSight also points to its ability to capture every single network event, rather than just those taking place after it has detected there is a problem, which is something that Endace also claims, but that several other vendors cannot do.

ClearSight developed the Cronos product in response to a request from “a leading provider of global business news and information services”, according to the announcement accompanying the launch. It is only natural, therefore, that the first instantiation of the product should focus on market data, which are the stock in trade of the customer in question.

However, told Ovum at the time that it intended to move to the other side of the low-latency equation, namely the actual execution of trades, and for this purpose it was developing a range of enhancements, such as the ability to run data models and order flow execution in order to dissect trades and see performance characteristics.

Since 2003, ClearSight has been majority-owned by Toyo Corporation, a Japanese vendor of electromagnetic compatibility (EMC) measurement technology. Given that Asia Pacific is the next big geography for HFT, with the starting pistol being the Tokyo Stock Exchange’s new matching engine coming online in January 2010, it will be interesting to see whether ClearSight can leverage its parent’s position in that region to win business there.

Correlix

New York-based Correlix is one of the vendors that came into existence specifically to develop monitoring and measurement technology for the capital markets, delivering its Latency Intelligence on the NinjaBox appliance from Endace.

It launched the appliance offering in February this year and followed that move in June with the unveiling of a managed service based on the device, entitled Correlix RaceTeam, specifically for measuring inter-party latency.

This is a service whereby Correlix deploys its appliances at trading venues and in market participants’ infrastructure, from where it sends monitoring data to the vendor’s central management server housed in Correlix’s data center. The customer can then access latency information regarding its links to the exchanges and its counterparties on a Web interface.

In essence, RaceTeam can be thought of as a sort of extranet specifically dedicated to monitoring and measurement, whereby the buy and sell sides, as well as the trading venues and service providers such as market data suppliers and hosting companies, gain information on latency and the benefits increase exponentially as the number of venues and counterparties on the network grows.

RaceTeam was launched in June this year, since when Correlix has announced deals with trading services companies Schneider Group and 7ticks (now part of Interactive Data). Most recently, Correlix unveiled a deal with Nasdaq OMX, which will see the exchange group offering the RaceTeam service to its customers and using its data to underscore marketing of its co-location services.

Corvil

Corvil is an Irish software developer that markets appliance-based solutions for high-performance network monitoring.

In recent years the company has focused its marketing efforts exclusively on the electronic trading industry, having identified this segment as a growth area with an immediate need for the type of precision monitoring it provides.

Its customer base includes trading venues such as the London Stock Exchange (LSE), Deutsche Börse and CME, financial service providers like Fixnetix, Savvis, BT Radianz and Thomson Reuters, and buy- and sell-side firms, of which the best-known publicly referenceable names are Credit Suisse and Morgan Stanley, as well as a number of specialist electronic traders and market makers.

Corvil sees as a unique selling point for its CorvilNet technology the fact that it is architected around a distributed data model, whereby packets are captured on a number of probes deployed at key points in a customer's infrastructure and stored locally, with only essential information on each packet being sent over a lightweight, optimized protocol to a central repository.

This contrasts with the approach of centralizing all the data in one place and carrying out analysis on them there, which makes it easier to undertake network-wide analysis and establish a single point of truth, but runs into problems on larger, busier networks on account of the bandwidth required for forwarding the packets to the central database.

Its distributed data model is probably one of the reasons Corvil is popular with companies seeking to address the issue of inter-party latency, since it obviates the need to backhaul vast amounts of performance data from multiple counterparties to a central point for analysis.

Corvil's distributed data model

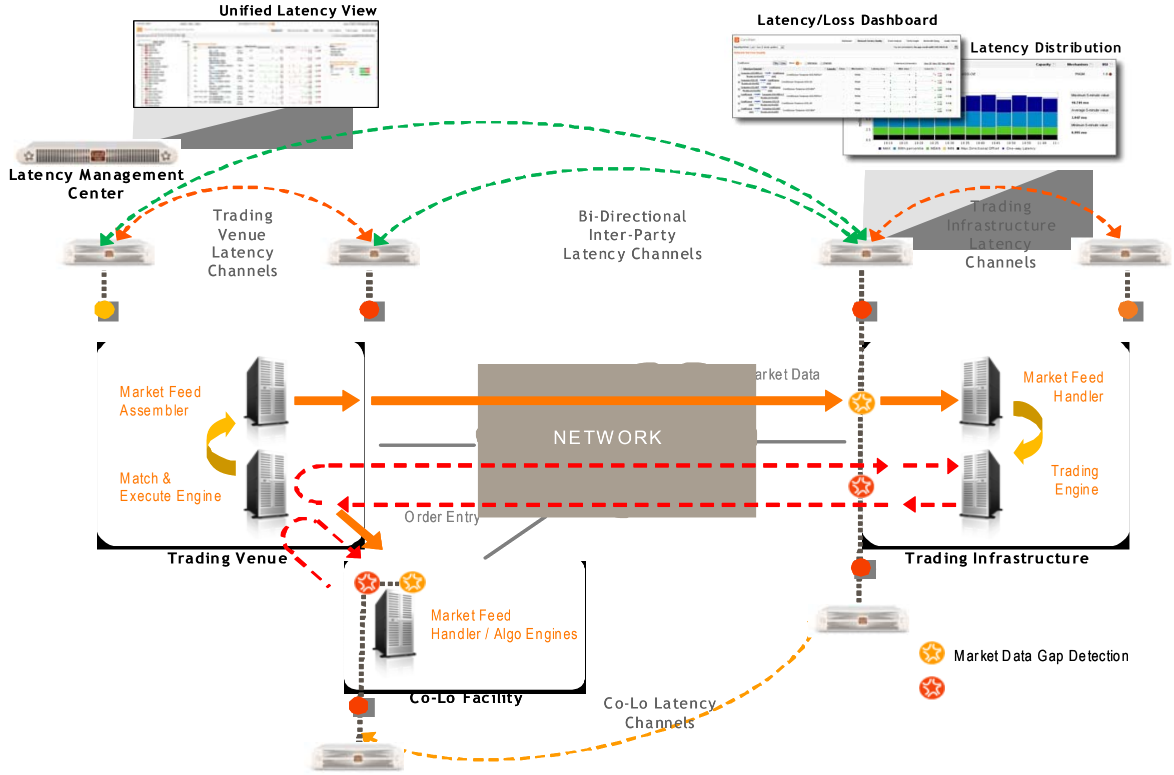

Source: Corvil

There are now efforts underway to come up with an industry standard for inter-party latency of FIX messages, the putative standard being called FIX Inter-Party Latency (FIPL). Corvil views this as a positive development, indicating the strength of interest in inter-party latency management, and intends to participate in standardization efforts as the field matures. The company's strategy in this context is to maintain the competitive lead it has established as the first company to introduce an inter-party solution through ongoing investment in innovation.

Another feature of CorvilNet that the company claims to be unique is the fact that it does not require an external source of synchronization such as Global Positioning System (GPS), Code Division Multiple Access (CDMA) or Precision Time Protocol (PTP). Instead, the software on the appliance includes its own built-in algorithms to solve time synchronization issues (essential for latency measurement in distributed environments). Corvil views this feature as an advantage, because external time-sync solutions are cumbersome to deploy and are not widely used among customers in this space at present.

While many of Corvil's existing customers have purchased CorvilNet for its network-layer latency monitoring capabilities, the company perceives an opportunity to expand from the network-layer to application-layer monitoring. For example, the most recent version (5.2) of CorvilNet includes support for analyzing more than 80 different market data protocols. Corvil intends to continue extending its capabilities in this area in other to broaden its user base and drive more sales to both existing and new customers.

ITRS

ITRS is a UK software developer, with sales offices in London, New York and Hong Kong.

The company started out in the front office in the mid-1990s, monitoring market data platforms such as Reuters Triarch (the predecessors to RMDS) and TIBCO market data services, and expanded to trading systems monitoring in 1999.

On the trading systems side, it began looking at exchange connectivity for systems from Fidessa, GL Trade (now part of SunGard), Patsystems and Trading Technologies. It now has interfaces for over 100 trading systems, including messaging infrastructures such as Rendezvous and MQ, traditional databases like Oracle and Sybase and non-traditional ones such as Coherence, Oracle's in-memory distributed data grid product for clustered applications and application servers. It also monitors exchange connectivity and extends its technology offering into the middle office, monitoring clearing and settlement services.

The core product offering from ITRS is the Geneos product suite, which consists of a central enterprise management framework that controls the distributed data capture components for each of the systems or services to be monitored.

ITRS's Geneos has three conceptual layers

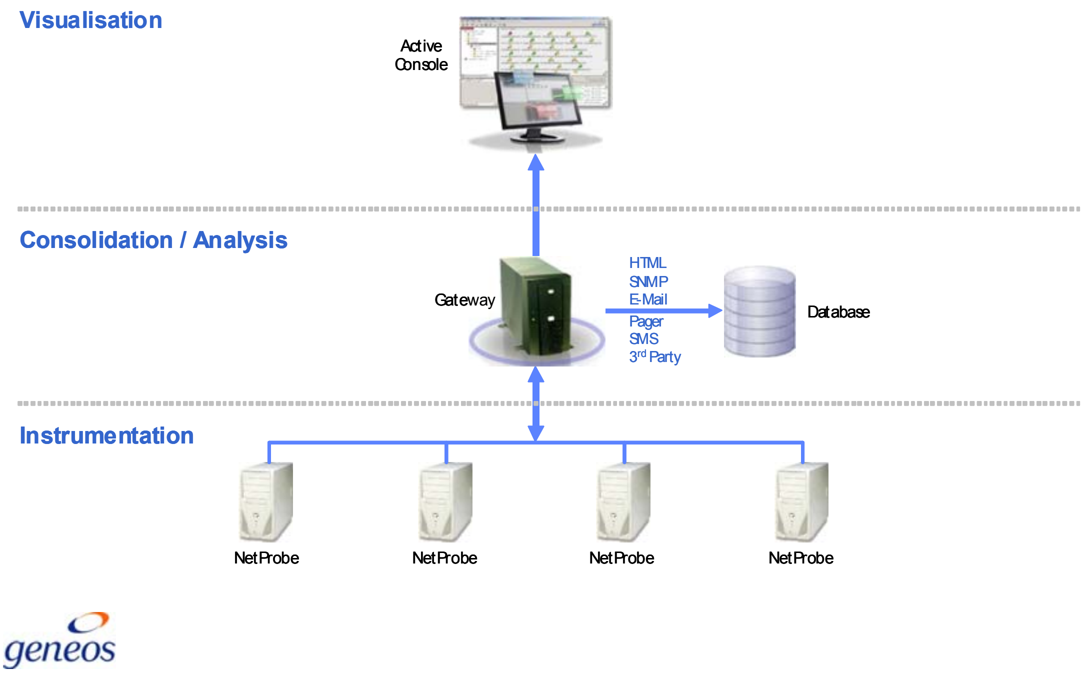

Source: ITRS

The firm takes Geneos to market in two ways: there is a perpetual license option with annual maintenance, or a two or three-year rental contract. In both cases, of course, the product is on-premise software comprising probes on individual servers, which it refers to as its instrumentation layer for data capture, a gateway layer with caching and a rules engine, and the virtualization layer, called ActiveConsole.

The company says it does not need hardware for data capture, doing it all, including timestamping, in software. This can be explained by its overall approach, which is to look only at specific application-level messages rather than every packet, filtering out others that are not relevant for its specific line of inquiry at the time.

For the future, ITRS plans to introduce Advanced Analytics, a system whereby rules on when an alert should be triggered will become self-learning, meaning that they will become more adaptive rather than dealing with traffic in a black-and-white manner. The company also wants to develop Breach Predictor, an ability within Geneos to predict how the order management system will perform later in the day, for instance.

NetQoS

NetQoS made its first foray into low-latency infrastructure monitoring for the financial markets with its Trade Monitor for FIX protocol monitoring in June 2008. It has since expanded the device's capabilities to take in the FAST protocol, representing its first foray into the market data monitoring side of things.

Like several other contenders in this market, NetQoS comes from a generic network monitoring background. It entered the fray in capital markets via its acquisition, in June 2008, of specialist developer Helium Systems, whose technology underpins Trade Monitor for FIX. This is a 1U appliance using packet capture and timestamping hardware from Endace, and is aimed at the trade execution side, where the FIX protocol has the lion's share of the market in terms of the transport layer.

A logical next step would be for NetQoS to extend its portfolio for the financial markets further into the market data side of the equation, adding support for the more proprietary protocols that still prevail in that area. Rather than develop multiple protocol-specific products, the company may opt for an approach that is agnostic.

Of course, the big change in NetQoS's life now is that it was acquired by systems management heavyweight CA in September this year, so there is currently a question mark over the company's strategy, and indeed its ongoing commitment to developing monitoring and measurement technology specifically for the capital markets, given that CA is a broad enterprise player and may thus prioritize that side of NetQoS.

That said, Trade Monitor already has one or two high-profile customers, so CA may feel it is worth devoting resources to further development in this more specialist niche. Furthermore, in a statement to Ovum explaining that Trade Monitor will see its roadmap defined only in the first half of 2010, CA alluded to potential synergies with its Wily product for monitoring business processes, which suggests that Trade Monitor may indeed enjoy a renewed focus under CA's aegis.

Network Instruments

Network Instruments has been in the generic network monitoring market since its foundation in 1994 and increased its focus on the financial sector in version 13 of the Observer software platform that underpins all its products, launched in October last year. Most of the company's portfolio is made up of appliances and, unusually, it also develops its own hardware, rather than relying on third-party providers such as Endace or Napatech.

A screenshot from Network Instruments' Observer

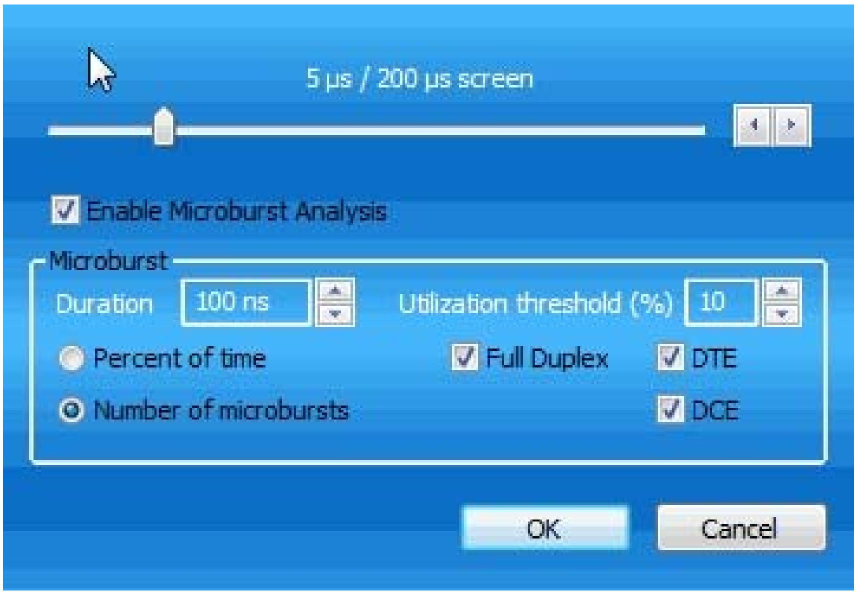

5 µs / 200 µs screen

☑ Enable Microburst Analysis

Microburst

Duration: 100 ns

Utilization threshold (%): 10

○ Percent of time

● Number of microbursts

☑ Full Duplex

☑ DTE

☑ DCE

Source: Network Instruments

Version 13 of Observer added the ability to analyze the FIX protocol for trade execution and FAST for market data, the UTP Quotation Data Feed (UDQF) format from Nasdaq OMX and messaging from capital-markets-specific vendor 29West. It also announced 10-nanosecond latency for reporting and resolution, a microview and macroview graph capability for analyzing microbursts and GPS synchronization, synching to within 80 nanoseconds.

NIKSUN

NIKSUN's approach was quite different from many players in the market. While their idea was to warehouse and index everything on the network, NIKSUN's core technology is more like a combination of Tivo and Google for networks.

Initially, NIKSUN applied this approach to generic network monitoring and found particular market traction in capital markets. As such, it was ahead of its competitors in terms of increasing its focus on the capital markets, and still derives a larger percentage of its business from that vertical than larger players such as NetScout.

Like its larger competitors, NIKSUN delivers appliances with its software on hardware provided by one of the dedicated packet capture vendors. However, some of its appliances have homegrown hardware acceleration for networks such as

10Gbps. It offers licenses for two different types of software on the devices, namely performance monitoring, where it competes with companies such as NetScout and Network Instruments, and network security, where it is up against the likes of NetWitness and Arbor Networks. The performance side of the product is called NetVCR, while the security side is called NetDetector.

In performance monitoring, especially in its modules for capital markets, NIKSUN's patented Camera on the Network Technology captures, timestamps and analyzes all data flowing across a network to identify trade-impacting network anomalies. All occurrences of latency are reported to actionable parties in real-time so as to enable prompt initiation of cause discovery, forensic analysis and remedy processes.

While it started life on the network side, NIKSUN has added application awareness to its products, meaning that it can now understand FIX, Tibco Rendezvous, 29West and other formats. It lists as some of its advantages:

- advanced capabilities to measure one-way delay;
- real-time microburst alarming and cause-analysis;
- ability to discover, diagnose and remedy multicast gaps;
- centralized visibility of the enterprise-wide network;
- common, multi-time repository of trading information
- Dynamic Application Awareness (i.e. fingerprinting applications by content and not just by IP address, port number or URL).

NIKSUN's technology is used by capital market institutions to monitor multicast data feeds, FIX engines and exchange connectivity. Meanwhile the algorithmic trading divisions within its customer base use a NIKSUN API to monitor FIX transaction times by exchange and by order ID number, creating their own portals so that different constituencies within their organization can get performance views in accordance with their specific requirements.

By hooking up the NIKSUN API to an algo engine, moreover, the latter can actually take intelligent order-routing decisions based on latency on point-to-point links, as well as on end-to-end round trip times and per-hop losses, with NIKSUN monitoring across multiple networks.

In this context, it is clearly useful that NIKSUN not only understands TCP/IP and UDP, but also traditional carrier networking protocols such as T1/T3, OC-48 and 192, SONET/SDH and Multiprotocol Label Switching (MPLS), not to mention Stream Control Transmission Protocol (SCTP), an alternative transport layer protocol to TCP that the company says is becoming increasingly popular in Europe.

SeaNet Technologies

SeaNet began life as a consultancy firm working on network monitoring in the general enterprise market in 2001. It brought its flagship product, SeaView, to market in 2004 to deliver performance and SLA monitoring of business applications.

It is this application-centric view that SeaNet highlights as being its differentiator, in that it sets out to instrument the application such that it performance can be understood in full business context from the user's perspective, where many competing products are essentially looking at network performance only. It cites deployments in a variety of sectors, including satellite TV companies, cellular operators, manufacturing firms and retail banks.

It was also in 2004 that SeaNet began using the Data Acquisition and Generation (DAG) cards from Endace to capture data packets off the wire at line speed and timestamp them, and the following year it began work in the capital markets, selling into the investment banking community. The capital markets sector now represents around 95% of SeaNet's business.

SeaView works in two stages. First it decodes the data from every application of interest on the SeaView/RTM probes deployed around a customer's infrastructure, timestamping them using CDMA or GPS technology, which it describes as a more economical alternative to an atomic clock but with the same accuracy. It supports both 1Gb and 10Gb connections as well as InfiniBand for this purpose and claims to be able to decode up to 200,000 messages per second per probe at line speed.

Product architecture for the SeaView platform from SeaNet

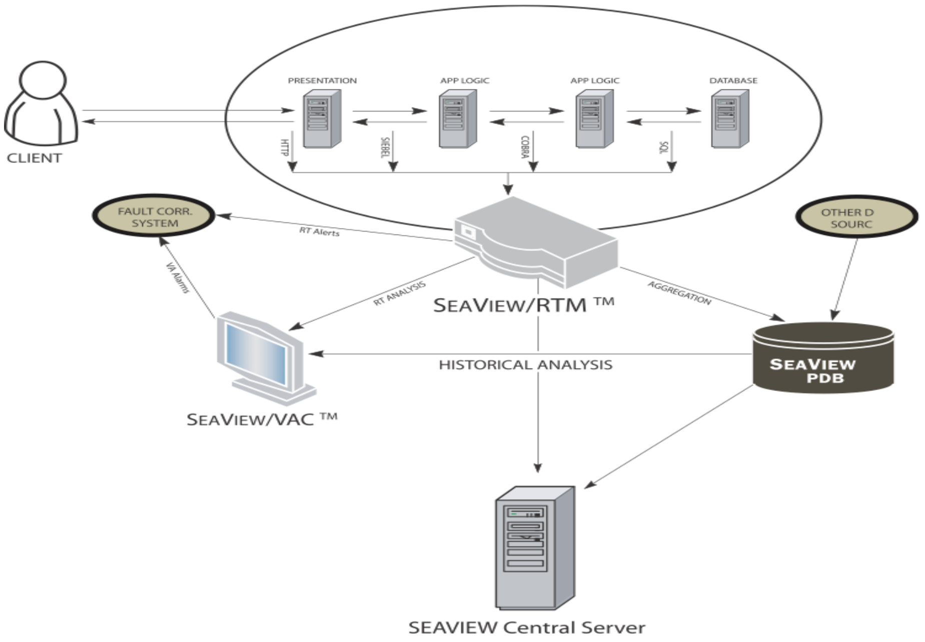

Source: SeaNet

This decoding is where the application intelligence first shows itself, since SeaView can deconstruct each packet and identify individual tags within it such as price, ticker, customer identification and source and destination IP addresses. The decoding takes place at every probe location. The resulting information, which it calls SeaView Events, is temporarily stored locally on the individual probe, prior to being forwarded to a central console engine, called the SeaView/Matcher.

At the matcher engine SeaView has technology to reconstruct market data and/or order flows, mapping every request to every response in what the company describes as a deterministic manner, differentiating it from the probabilistic way used by other systems on the market. The matchers recreate every tier and hop of a message's journey, a correlation capability that again the company feels differentiates it from what else is on offer in the market.

Once these two stages are completed, the information can be displayed in near real-time on a SeaView/VAC dashboard and the Analysis Cube for visualization and analysis purposes. SeaNet cites NYSE/Euronext as a customer, in a deal announced in mid-2009. NYSE is using the technology to monitor both order flow and market data latency.

SeaNet is headquartered in New York with offices in the UK as well as other US locations. The product is sold directly and through SeaNet's partners.

TS-Associates

TS-Associates (TS-A) is a UK-based company that started life in 1999 as a consultancy and began launching products in 2003, when it unveiled software that could detect data feed and system problems in trading systems based on Tibco's Rendezvous messaging technology.

Its current flagship is TipOff, which was launched in 2005. The company describes the product as a real-time network, middleware and transaction analysis appliance. As is the norm in this market, TS-A develops the software for its products and ships it on specialist hardware, namely x86 servers fitted with packet capture cards from Napatech, Accolade and Endace.

TS-A considers one of the principal differentiators of its TipOff vis-à-vis its competitors is the fact that, while most of them measure packet latency, i.e. looking at the TCP/IP packets traversing a network, TipOff is also able to measure message latency, thanks to its ability to understand the protocols used at that higher layer in the stack. That means, for instance, that its latency calculation takes into account retransmission requests and actual retransmissions.

While messaging protocols are standards-based such as the Financial Information eXchange (FIX) protocol for order transmission, the vast majority of protocols on the market data side remain proprietary. This means TS-A has had to go through the process of developing its expertise in each individual protocol as required by its customers, though of course by now it supports the majority of the most widely used ones.

In this context, TSA draws on the expertise gathered from its background as a consulting firm specializing in the financial middleware arena and points out that the ability to decode messages in multiple formats is of value in that latency can be seen throughout a system, even if the wire or message protocol changes; for example through a feed handler, or distribution system.

Another advantage TS-A cites with TipOff is its ability to support market arbitrage strategies by publishing latency data directly into the trading systems its customers are using. While all developers of latency monitoring technology must be able to present their findings, most of them publish the information to a database and then provide their customers a schema through which to extract the data.

TS-A's product, by contrast, can publish a metadata feed straight into a CEP engine or smart order router within a customer's infrastructure, using the same API as the market data itself, thereby keeping the latency data in the same context for purposes of analysis and avoiding the additional latency of a database look-up.

A significant development this year was the announcement by TS-A of Application Tap, a PCIe card that enables the implementation of precisely time-stamped software instrumentation with what it calls negligible performance overhead. The Application Tap was previously an internal component within TipOff, but is now also offered as a product in its own right.

The decision to separate out the Tap reflects the evolution of trading infrastructure from highly distributed systems towards consolidation of multiple applications onto multi-core servers, making less inter-process communication available to capture on the network.

Such consolidation of different functions onto a single box also makes it increasingly important to monitor software events occurring inside the multiple applications, such as order matching engines, and emit data associated with those events.

The Application Tap provides a user space API that enables software events within applications to be instrumented with minimal overhead, with the aid of an FPGA based co-processor and precise hardware clock. Instrumentation code added to applications is able to pass instrumentation metadata to the co-processor, which time stamps the event to 10nS resolution and forwards the instrumentation metadata to an external monitoring appliance such as TipOff, using an on-board Gigabit network interface. Clock synchronization meanwhile is supported using either the Pulse Per Second (PPS) or Precision Time Protocol (PTP) methods.

TS-A is now contributing to the specification work for a new standard within the context of FIX Protocol Ltd (FPL), the not-for-profit organization that defines FIX itself and all related standards. The new one is the so-called FIX Inter-Party Latency (FIPL) and will be designed to enable interoperability between different products monitoring inter-party latency.

This is an important development since, while an exchange and or trading venue may be using one vendor's package for this purpose, brokers and other sell-side firms may be using another, making it very difficult to agree when an SLA has been met.

PROVIDERS OF DATA CAPTURE HARDWARE

Endace

This New-Zealand-based company, specializing in traffic capture for the security and monitoring markets, made its name with its DAG cards, which it supplies on an original equipment manufacturer (OEM) basis to application providers in both the network performance monitoring and security markets.

DAG cards use direct memory access (DMA) and FPGAs to capture data packets running across a wire in a non-intrusive (i.e. out-of-band) manner. It then uses a variety of sources of accurate time (GPS, CDMA and the IEEE's 15888/PTP protocol) to timestamp every packet individually.

Endace sells into governments, service providers and general enterprise markets all over the world. However, the cards' capabilities are clearly of value in the context of the capital markets and their need to know exactly when messages left a given location and arrived at another, as well as how long they take to pass through the different parts of a company's infrastructure. Thus DAG cards are used by application vendor SeaNet to underpin its appliance offering.

Meanwhile Endace introduced an appliance of its own, a device called the NinjaBox, in 2007, and that product is used by Correlix, a monitoring developer for the capital markets,, which offers it as part of its Latency Intelligence (Correlix LI) Suite, with versions for trade execution and market data. Indeed, as Correlix puts it in its marketing, "Endace NinjaBox appliances and Endace Data Acquisition and Generation (DAG) interface cards are the foundation technology for the Correlix Latency Intelligence solution."

While Endace's technology supports 10Gb Ethernet up to 40Gb/s, one of the company's differentiating features is its ability to inspect InfiniBand, a low-latency network connection used in high-performance computing environments. Not surprisingly, it also counts among its partners Voltaire, the company whose name is most widely associated with InfiniBand on account of the switching fabrics it makes for the protocol, even though it also supports 10Gb Ethernet.

Napatech

Napatech provides what it calls "intelligent network adapters for real-time network analysis". This terminology is intended to convey the dual virtues of intelligence and real-time packet capture that distinguishes Napatech network adapters from standard NIC and server adapters, such as those provided by Intel. It also reveals that Napatech is entirely focused on the "network analysis" niche market and has designed its network adapters to address the needs of high-speed packet capture with zero packet loss and time synchronization, both features that are important to network monitoring and latency measurement.

Indeed, not only does Napatech focus on the network analysis market, but it also focuses solely on OEM network appliance vendors, including ISVs that need hardware acceleration. From a technical perspective, it is focused on x86 servers with a Peripheral Component Interconnect Express (PCI-Express) form factor and Ethernet (1Gbps and 10Gbps) today, but says it is open to supporting other form factors and protocols should the business case be attractive.

Napatech's network adapters are designed to address two issues:

- the inability of standard network adapters (NICs) to guarantee zero packet loss at all packet sizes;
- The need to accelerate application throughput via central processing unit (CPU) data-handling off-load and multi-CPU support.

For network monitoring and latency measurement, it is vital that no packets are lost, as the integrity of captured data will be compromised. At full 1Gbps or 10Gbps line-rate, the throughput of standard NICs falls rapidly at low Ethernet frame sizes with significant packet loss. Napatech's claim is that, with its technology, full throughput at full line rate is provided under all conditions.

Napatech's 10Gbps in-line packet processing technology

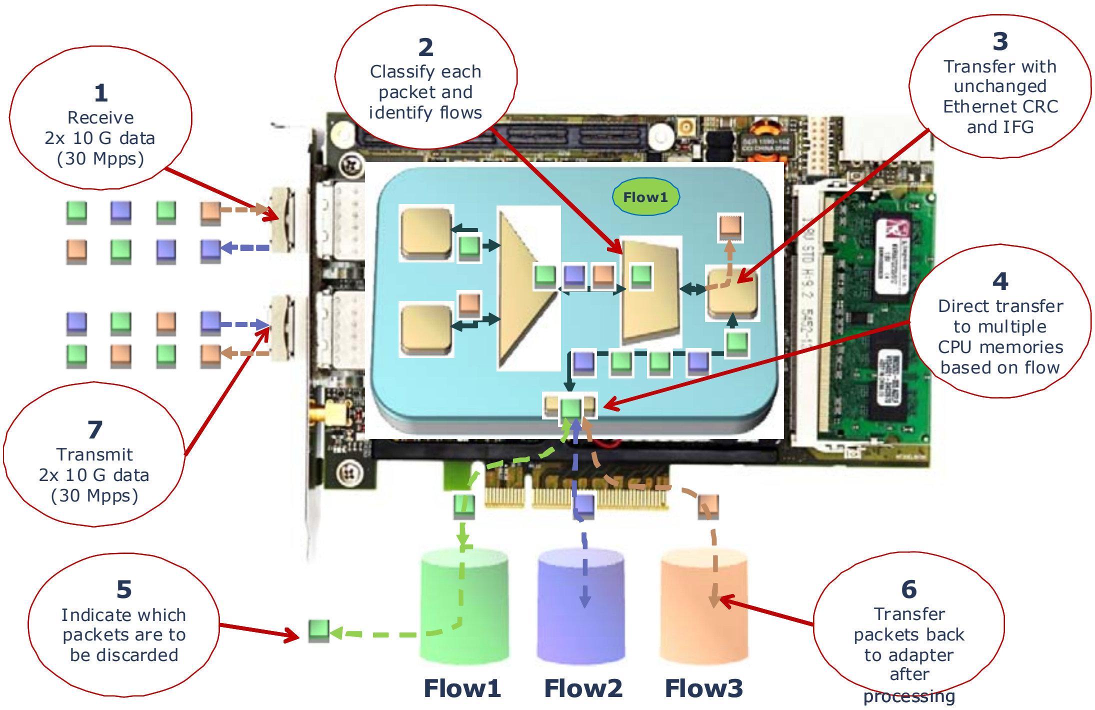

Source: Napatech

At 10Gbps, up to 75% of CPU cycles can be expended on data-handling. Napatech includes a number of features designed to offload this data processing from the CPU to the network adapter, including packet decoding and classification, the merging of packet streams from multiple ports, filtering and de-duplication.

In addition, 17 different hash-key modes are available for defining flows based on Ethernet, IP, TCP, UDP header information and various types of tunneling protocol information such as SCTP, Generic Routing Encapsulation (GRE) and

GPRS Tunnelling Protocol (GTP). This information can then be used to distribute the flow to up to 32 CPUs on a per-flow basis or for balanced load. A zero-copy DMA process is used to transfer data.

A number of offset pointers are also available to indicate where the packet payload and layer-4 header information can be found. This supports multiple virtual LAN (VLAN) tags and MPLS labels, as well as Cisco's Inter-Switch Link technology.

One of the most important features of the Napatech network adapters for monitoring and latency measurement is nanosecond precision time-stamping with hardware time-synchronization. As the market moves towards 10Gbps networks, nanosecond precision becomes vital.

Napatech provides hardware time-stamping with 10ns precision on each network adapter. The free running clock on the network adapter can be synchronized with a number of different Pulse Per Second (PPS) time sources, such as GPS, CDMA and IEEE1588v2 to an accuracy of 50 nanoseconds, although of course it only guarantees this accuracy in relation to the PPS output of whatever third-party device the card is loaded into. Beyond that, accuracy will depend on how well that device controls its connection to an external UTC source, which will typically be around 30 nanoseconds, resulting in an overall accuracy of some 80 nanoseconds.

Distribution of the PPS and UTC time signals can be accomplished via a daisy-chaining mechanism from one network adapter to the next using a proprietary Napatech PPS protocol, or using an external distribution unit known as the Napatech Time Synchronization Unit (NTTSU).

RECOMMENDATIONS

Ovum sees market opportunities in monitoring and measurement technology for the capital markets, in that all the components of the market (buy side, sell side and exchanges/trading venues, not to mention service providers such as network operators and market data suppliers) need to deploy something to track how well their respective infrastructures are working, both for troubleshooting purposes and to monitor whether SLAs are being met. Its recommendations are:

Systems integrators (SIs) and consultancies

There is a role for SIs and/or consultancy firms in this market, particularly for smaller customers on the buy side, who may lack the technical nous to determine which type of monitoring technology they require and thus narrow down their list of potential suppliers, then actually deploy the one they choose, road test it and tweak it to meet the precise requirements of their infrastructure.

Data center operators

With proximity hosting growing as HFT grows as a percentage of overall trading activity around the world, data center operators should deploy technology to demonstrate the latency of incoming market data and outgoing orders from their facilities. They will typically be looking for tools to monitor inter-party latency, though it would also be advisable to deploy something inside their data centers to measure latency between the ingress/egress point and their customers' ticker plants and trading platforms/order management systems. On the inter-party latency side, they would do well to accompany the development of the putative standard in this space, namely FIPL, which should change the dynamics of that market as and when it is ratified.

Service providers and trading venues

Service providers need monitoring technology to demonstrate that they are meeting the SLAs they have signed with their customers, while trading venues need to show that their infrastructure offers round-trip latency which is at least in line with, if not better than, their competitors. An interesting move in this context has been Nasdaq OMX's deal with Correlix in November, whereby it promotes the latter's RaceTeam service to its customers as independent third-party verification of its latency numbers but also uses it to help promote its ancilliary services, such as co-location. Network providers with data centers from which they offer proximity hosting services could well look to emulate Nasdaq OMX's strategy to validate their own latency and help push their offerings.

APPENDIX

DEFINITIONS

- **Algorithmic trading** – in electronic financial markets, algorithmic trading or automated trading, also known as algo trading, black-box trading or robo trading, is the use of computer programs for entering trading orders with the computer algorithm deciding on aspects of the order such as the timing, price, or quantity of the order, or in many cases initiating the order without human intervention. Algorithmic Trading is widely used by pension funds, mutual funds, and other buy-side (investor driven) institutional traders, to divide large trades into several smaller trades in order to manage market impact, and risk. Sell-side traders, such as market makers and some hedge funds, provide liquidity to the market, generating and executing orders automatically. In this HFT computers make the decision to initiate orders based on information that is received electronically, before human traders are even aware of the information. Algorithmic trading may be used in any investment strategy, including market making, inter-market spreading, arbitrage, or pure speculation (including trend following). The investment decision and implementation may be augmented at any stage with algorithmic support or may operate completely automatically ("on auto-pilot"). As of 2009, HFT firms account for 73% of all US equity trading volume.
- **Complex event processing (CEP) engines** – CEP is primarily an event processing concept that addresses the task of processing multiple events with the goal of identifying the meaningful events within an event cloud. CEP employs techniques such as detection of complex patterns of many events, event correlation and abstraction, event hierarchies, and relationships between events such as causality, membership, and timing, and event-driven processes. In his blog, Mark Tsimelzon, president and CTO of Coral8, writes that "a CEP engine is a platform that makes it easy to write and deploy applications that process and analyze real-time data."
- **Electronic communication network (ECN)** – an ECN is the term used in financial circles for a type of computer system that facilitates trading of financial products outside of stock exchanges. The primary products that are traded on ECNs are stocks and currencies. ECNs came into existence in 1998 when the SEC authorized their creation. ECNs increase competition among trading firms by lowering transaction costs, giving clients full access to their order books, and offering order matching outside of traditional exchange hours. In order to trade with an ECN, one must be a subscriber or have an account with a broker that provides direct access trading. ECN subscribers can enter orders into the ECN via a custom computer terminal or network protocols. The ECN will then match contra-side orders (i.e. a sell-order is "contra-side" to a buy-order with the same price and share count) for execution. The ECN will post unmatched orders on the system for other subscribers to view. Generally, the buyer and seller are anonymous, with the trade execution reports listing the ECN as the party. Some ECNs may offer additional features to subscribers such as negotiation, reserve size, and pegging, and may have access to the entire ECN book (as opposed to the "top of the book") that contains important real-time market data regarding depth of trading interest.

- **Ethernet** – Ethernet is a family of frame-based computer networking technologies for LANs. The name originated from the physical concept of the ether. It defines a number of wiring and signaling standards for the physical layer of the OSI networking model, through means of network access at the media access control (MAC) /data link layer, and a common addressing format. The data link layer is layer 2 in the OSI model, such that Ethernet is often referred to as a layer-2 protocol. Ethernet is standardized as IEEE 802.3. The combination of the twisted pair versions of Ethernet for connecting end systems to the network, along with the fiber optic versions for site backbones, is the most widespread wired LAN technology and the efforts of bodies such as the Metro Ethernet Forum seek to popularize its use in Metro and wide area networks (MANs and WANs).
- **Feed handler** – since trading venues began offering data feeds with their prices streamed directly to market participants rather than as part of a consolidated feed covering multiple venues (so-called direct feeds), it has been necessary to install software at the receiving end that normalizes them, i.e. puts them into a standard format, prior to sending them on to the relevant consuming applications. This software is called a feed handler.
- The financial information exchange protocol (**FIX)** – FIX is an open specification intended to streamline electronic communications in the financial securities industry. FIX supports multiple formats and types of communications between financial entities including email, texting, trade allocation, order submissions, order changes, execution reporting and advertisements. FIX is employed by numerous financial vendors and has emerged as the favored specification among trading partners. The concept originated in 1992 when several brokers expressed interest in using the fledgling Internet to improve the speed, volume and efficiency of their trading activities. FIX is vendor-neutral.
- **Internet protocol (IP)** – IP is a protocol used for communicating data across a packet-switched internetwork using the Internet Protocol Suite, also referred to as TCP/IP. IP is the primary protocol in the internet layer of the Internet Protocol Suite and has the task of delivering distinguished protocol datagrams (packets) from the source host to the destination host solely based on their addresses. From the perspective of the OSI networking model, IP operates at the network layer, which is layer 3, such that it is often referred to as a layer-3 protocol.
- **Jitter** – jitter is the time variation of a characteristic of a periodic signal in electronics and telecommunications, often in relation to a reference clock source. Jitter may be observed in characteristics such as the frequency of successive pulses, the signal amplitude, or phase of periodic signals. Put more simply, jitter can be thought of as variation in latency.
- **Latency** – latency is a time delay between the moment something is initiated, and the moment one of its effects begins or becomes detectable. The word derives from the fact that during the period of latency the effects of an action are latent, meaning "potential" or "not yet observed". Latency in a packet-switched network is measured either one-way (the time from the source sending a packet to the destination receiving it), or round-trip (the one-way latency from source to destination plus the one-way latency from the destination back to the source). Round-trip latency is more often quoted, because it can be measured from a single point. Note that round trip latency excludes the amount of time that a destination system spends processing the packet. Many software platforms provide a service called ping that can be used to measure round-trip latency. Ping performs no packet

processing; it merely sends a response back when it receives a packet (i.e. performs a no-op), thus it is a relatively accurate way of measuring latency. Where precision is important, one-way latency for a link can be more strictly defined as the time from the start of packet transmission to the start of packet reception. The time from the start of packet reception to the end of packet reception is measured separately and called "Serialization Delay". This definition of latency is independent of the link's throughput and the size of the packet, and is the absolute minimum delay possible with that link. However, in a non-trivial network, a typical packet will be forwarded over many links via many gateways, each of which will not begin to forward the packet until it has been completely received. In such a network, the minimal latency is the sum of the minimum latency of each link, plus the transmission delay of each link except the final one, plus the forwarding latency of each gateway. In practice, this minimal latency is further augmented by queuing and processing delays. Queuing delay occurs when a gateway receives multiple packets from different sources heading towards the same destination. Since typically only one packet can be transmitted at a time, some of the packets must queue for transmission, incurring additional delay. Processing delays are incurred while a gateway determines what to do with a newly received packet. The combination of propagation, serialization, queuing, and processing delays often produces a complex and variable network latency profile. Significant latency in the capital markets was, until a couple of years ago, measured in milliseconds (ms), but that has changed, with the advance of algorithmic trading, to microseconds (µs).

- **Multicast** – multicast addressing is a network technology for the delivery of information to a group of destinations simultaneously using the most efficient strategy to deliver the messages over each link of the network only once, creating copies only when the links to the multiple destinations split. The word "multicast" is typically used to refer to IP multicast which is often employed for streaming media and Internet television applications. In IP multicast the implementation of the multicast concept occurs at the IP routing level, where routers create optimal distribution paths for datagrams sent to a multicast destination address spanning tree in real-time. At the data link layer, multicast describes one-to-many distribution such as Ethernet multicast addressing, Asynchronous Transfer Mode (ATM) point-to-multipoint virtual circuits or Infiniband multicast.
- **Multilateral trading facility (MTF)** – an MTF is a multilateral system, operated by an investment firm or a market operator, which brings together multiple third-party buying and selling interests in financial instruments - in the system and in accordance with non-discretionary rules - in a way that results in a contract in accordance with the provisions of Title II of MiFID.1
- **Multiprotocol label switching (MPLS)** – MPLS is a mechanism in high-performance telecommunications networks which directs and carries data from one network node to the next. MPLS makes it easy to create "virtual links" between distant nodes. It can encapsulate packets of various network protocols. MPLS is a highly scalable, protocol agnostic, data-carrying mechanism. In an MPLS network, data packets are assigned labels. Packet-forwarding decisions are made solely on the contents of this label, without the need to examine the packet itself. This allows one to create end-to-end circuits across any type of transport medium, using any protocol. The primary benefit is to eliminate dependence on a particular data link layer technology, such as ATM, frame relay, SONET or Ethernet, and eliminate the need for multiple layer 2 networks to satisfy different types of

traffic. MPLS belongs to the family of packet-switched networks. MPLS operates at an OSI model layer that is generally considered to lie between traditional definitions of layer 2 (data link layer) and layer 3 (network layer), and thus is often referred to as a "layer-2.5" protocol. It was designed to provide a unified data-carrying service for both circuit-based clients and packet-switching clients which provide a datagram service model. It can be used to carry many different kinds of traffic, including IP packets, as well as native ATM, SONET, and Ethernet frames.

- **The open system interconnection reference model (OPRA)** – OPRA (OSI reference model or OSI model) is an abstract description for layered communications and computer network protocol design. It was developed as part of the OSI initiative. In its most basic form, it divides network architecture into seven layers which, from top to bottom, are the application, presentation, session, transport, network, data-link and physical layers. It is therefore often referred to as the OSI seven layer model. A layer is a collection of conceptually similar functions that provide services to the layer above it and receives service from the layer below it. On each layer an instance provides services to the instances at the layer above and requests service from the layer below. For example, a layer that provides error-free communications across a network provides the path needed by applications above it, while it calls the next lower layer to send and receive packets that make up the contents of the path. Conceptually two instances at one layer are connected by a horizontal protocol connection on that layer.
- **Point of presence (PoP)** – a POP is an artificial demarcation point or interface point between communications entities. In the USA, this term became important during the court-ordered breakup of the Bell Telephone system. A point of presence was a location where a long-distance carrier could terminate services and provide connections into a local telephone network. An internet point of presence is an access point to the internet. It is a physical location that houses servers, routers, ATM switches and digital/analog call aggregators. It may be either part of the facilities of a telecommunications provider that the Internet service provider (ISP) rents or a location separate from the telecommunications provider. ISPs typically have multiple POPs, sometimes numbering in the thousands. POPs are also located at Internet exchange points and co-location centers. In the context of this report, the term is used to signify a node on an extranet to which customers can connect and, in some cases, where they can also locate their algoservers for lower-latency performance.
- **Transmission control protocol (TCP)** – TCP is one of the core protocols of the Internet Protocol Suite. TCP is one of the two original components, with Internet Protocol (IP), of the suite, so that the entire suite is commonly referred to as *TCP/IP*. Whereas IP handles lower-level transmissions from computer to computer as a message makes its way across the Internet, TCP operates at a higher level, concerned only with the two end systems, for example, a Web browser and a Web server. In particular, TCP provides reliable, ordered delivery of a stream of bytes from a program on one computer to another program on another computer. Besides the Web, other common applications of TCP include e-mail and file transfer. Among its other management tasks, TCP controls message size, the rate at which messages are exchanged, and network traffic congestion.
- **Ticker plant** – a ticker plant is a system for distributing financial market data that receives ticker feed data from many exchanges throughout the world, processes and formats the received data and then distributes or

broadcasts the data to regional customers in the form of securities transactional data denoting the security identity and related transactional data. A ticker plant contains multiple feed handlers, one per direct feed.

- **Unicast** – in computer networking, unicast transmission is the sending of information packets to a single network destination. The term "unicast" is formed in analogy to the word "broadcast" which means transmitting the same data to all destinations. Another multi-mode distribution method, multicasting, is similar to IP broadcasting, but implemented in more efficient manner. Unicast messaging is used for all network processes where a private or unique resource is requested making most networking traffic Unicast in form. Unicast is used where two way connections are needed to complete the network transaction. Certain network applications which are mass-distributed are too costly to implement on Unicast. These include streaming media of many forms. And when multicasting is unavailable, unicasting the exact same content to many users can be costly. Internet radio stations may have high bandwidth costs because of this. These terms are also used by streaming content providers' services. Unicast based media servers open and provide a stream for each unique user. Multicast servers can support a larger audience by serving content simultaneously to multiple users.

METHODOLOGY

- Ongoing vendor briefings – Datamonitor conducts interviews with software, hardware, networking and services vendors serving the financial markets industry on an ongoing basis.
- End-user surveys **–** Datamonitor conducts regular surveys with IT decision makers at financial institutions around the world.
- Secondary research – other secondary sources of information include international organization statistics, national governmental statistics, national and international trade associations, company filings, broker and analyst reports, company annual reports, and business information libraries and databases.

FURTHER READING

- Datamonitor (2008) *Seeking Low Latency In Financial Markets (Strategic Focus)*, October 2008, DMTC2275
- Datamonitor (2008) *QuantHouse Adds Connectivity to its Portfolio (Analyst Opinion)*, November 2008, BFTC2241

ASK THE ANALYST

Rik Turner, Senior Analyst, Financial Services Technology

rturner@ovum.com

OVUM CONSULTING

We hope that the data and analysis in this brief will help you make informed and imaginative business decisions. If you have further requirements, Ovum's consulting team may be able to help you. For more information about Ovum's consulting capabilities, please contact us directly at consulting@ovum.com.

DISCLAIMER

All Rights Reserved.

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission of the publisher, Ovum.

The facts of this report are believed to be correct at the time of publication but cannot be guaranteed. Please note that the findings, conclusions and recommendations that Ovum delivers will be based on information gathered in good faith from both primary and secondary sources, whose accuracy we are not always in a position to guarantee. As such Ovum can accept no liability whatever for actions taken based on any information that may subsequently prove to be incorrect.