

**Newsletter  
Summer 2008**



**Intelligence  
in Finance**

## A Year in the Life of Trading Technology – and SIFMA 2008

Where we will be hosting events and talking with partners, clients and press all week!



Welcome back to Intelligence in Finance – our regular newsletter on technology related activities we see in global financial services. Launched in early 2007, we are revisiting

the inaugural area of focus – trading.

### An interesting year - let's pick out some highlights.

Back in Q1 2007 we announced the opening of the Intel fasterLAB in London and the visibility of how the technology infrastructure could contribute to speed, performance and scale in the trade life cycle. Early tests showed in excess of 20 percent throughput improvement from FIX engines through the use of in-silicon features – and the low latency arms race was on. Early innovative focus on proximity hosting and networks soon became a norm, and the market looked at the technology layers which provided the best ROI if given the engineering treatment.

We seemed to be in a perfect storm. Market regulation (Reg NMS, MiFID) on both sides of the Atlantic was creating a backdrop for new market practices which were playing out in multiple new venues, the threat and opportunity of even more, disintermediation plays, market data

volumes going through the roof with no end in sight and demands on information and historic reporting (best execution) putting new strains on data management.

In such situations technology demand and supply moves up a gear – and the choice of routes to new competitive solutions becomes daunting. Out of this the Intel fasterLAB became established as a facility for Intel's many partners to trial software on the new processor platforms as they revved through dual- to quad- and multi-core, with 45nm processors entering the race in Q4. For clients, the same lab provides the ability to demonstrate and test the effects of individual layers in the solution stack by changing one layer and holding the others constant to see where incremental gain can most effectively be achieved.

As technology choices advanced and selection became an issue, STAC came to market with highly effective performance testing services for new technologies. Comprehensive reporting on market data and trading technologies under the "ever lower latency" banner became the norm – with the opening of the current chapter in September 2007 to standardise tests through the auspices of the STAC Benchmark Council. Collaboratively defined standard tests for market data feeds, distribution and analytics are in beta test and now mark a next phase in the assessment of trading technologies.

So where is Intel in trading technology and financial services? I believe strongly that the last 12 months has seen Intel become established and recognised as a facilitator of innovation in technology design and deployments. Our labs have proved that IA (Intel Architecture) tuning brings performance improvement – improvement that translates directly to extra trades, more market data, improved processing efficiency and reduced heat output from a data centre with the throttle on the floor boards.

The 2006 arrival of Core 2 microarchitecture has seen Intel lead the performance-watt race, 45nm introduced in 2007 delivered staggering improvements and in 2008 Nehalem microarchitecture will come to market and bring with it another quantum leap in innovation.

We hope you find Intelligence in Finance – where we try to give commentary on some of the highlights in the market – a helpful guide to our work with the financial services industry. This newsletter combines with our fasterFS events around the world and the [www.intelfasterfs.com](http://www.intelfasterfs.com) website to continue the Intel FSI community journey. Leave us comments, and we look forward to being with you again in September, before SIBOS with more topical commentary.

**NIGEL WOODWARD**  
Global Director, Financial Services  
Intel Corporation

## News in brief

### QuIC Achieves Landmark Speed With Intel Technology and DataSynapse

By running the Quad-Core Intel® Xeon® 5400 series processors on GridServer® software from DataSynapse, QuIC's Unified Market and Credit Risk Solution has achieved extremely fast calculation time.

*"During our test, we ran two Integrated market and credit risk portfolio calculations featuring this new technology engine and the results beat our expectations for speed and reliability,"* said Tony Coppellotti, chief technology officer for QuIC. It found that Intel's latest generation of 45nm Quad-Core Intel® Xeon® 5400 series processors deliver ultra-low latency and significant performance increases in algorithmic trading, direct market access, and market data delivery.

QuIC ran the tests on DataSynapse's GridServer software which dynamically scales enterprise application services based on performance and business requirements.

• [www.quic.com](http://www.quic.com)

### Making Intel fasterLAB Work Faster

CohesiveFT have deployed their Elastic Server product at Intel fasterLAB to identify the most cost effective way to deliver OPRA feeds with acceptable performance. Elastic Server automates the process of delivering custom application stacks to order, enabling Intel fasterLAB to set up and benchmark multiple test environments in less than an hour.

This testing follows from the breakthrough of 2007 where millions of OPRA mps was achieved using only one server, RabbitMQ AMQP, and Pantor FAST. This stack is now available in virtual formats accelerated by Intel VMDq enabling performance parity to be achieved at a lower cost.

### Red Hat Talks Business Value at SIFMA

Red Hat will be at SIFMA in New York City on June 10-12 providing updates on how offerings like Red Hat Enterprise MRG (Messaging, Realtime, Grid), Red Hat Enterprise IPA (Identity, Policy, Audit), and Red Hat Enterprise Linux bring significant value to the financial services industry. For more information on meeting with Red Hat at SIFMA, contact [financial@redhat.com](mailto:financial@redhat.com)

## STAC - Independent Lab Tests for Investment Banking Technology

You don't have to use the latest and greatest tech – you can always just get out of the (trading) business.

That was the message from Peter Lankford, co-founder of STAC, the New York-based Securities Technology Analysis Center.

*"Trading firms can continue to invest or they can choose to exit markets that require high performance technology. It's that binary nowadays. There is innovation all over the place. I don't remember when so much was focused on industry performance."*

Made up of seasoned trading market specialists, STAC set up to help firms by testing new combinations of hardware and software and measure the performance. Lankford said the market has seen a huge proliferation of viable solution stacks, market data vendor, event processing engines, fast operating systems and new processors.

*"Nobody has time to test all this stuff"*

Nor do they have the equipment and skills to measure latency in the sub-microsecond range. STAC has the expertise, and it can set up tests where all but one element remains constant. Then it can swap out processors, for example, to see what difference a new chip will make.

2007 saw the trading low latency arms race gain pace, a pace in which STAC played an important part as route to achieve public quality independent test results. The emerging issue however was that while test results made impressive reading in their own right – and drew the markets' attention, it was difficult to draw comparisons and identify objective differentiation between solution stacks under test.

Responding to this – in September 2007 STAC established under its auspices the benchmark council – a council of market firms and significant vendors – with the objective of defining de facto standard tests for the major workloads in the trading cycle.

STAC, in full view of the council, is now developing a spec for direct exchange feed or ticker plant solutions, known as test STAC M1, distribution functions will be measured via STAC M2 and analytics and data management (CEP and its derivatives) are on the drawing board.

That gives us high fidelity replays of live feeds, so we can get the spikiness of the markets with all their bursts."

Users who don't need the fastest speed, such as pension funds and insurers, can use the STAC labs to see the tradeoff between very fast and very expensive and reasonably fast and less expensive.

Intel and STAC have been collaborative throughout recent months – with the common view of establishing testing initiatives which are useful to the market. This is slated to grow through sharing lab facilities and working in virtual skills teams which best leverage STAC's trading domain expertise and Intel's infrastructure engineering, tuning and optimization capabilities via the fasterLAB centre.



## Building the Solution: Legacy to Contemporary

Working with the Thomson Reuters RDMS in the Intel fasterLAB is akin to adjusting the spoiler on a Formula 1 race car by one degree to improve performance, says Paul Gow of the consultancy, CJC Ltd.

*"The RDMS is very fast and it is a very resilient piece of software. It has a lot of parameters that can be tweaked to enhance the system, and a lot of those parameters are hidden and people don't know they exist."*

At the Intel fasterLab, CGC works with a base system of RDMS on which firms can use to test their desktop applications.

*"So a client will come in with an application that they want to go faster than the data coming out of the RDMS pipe. We will tune the system and make it go faster and hopefully make it break their application with the update rates and then we help them move forward."*

One of the most common sources of latency in trading is in-house applications that can't keep pace with market data. Firms write their APIs to the Reuters data, but if the API in the application can't handle the data refresh rates, the machine can come to a grinding halt.

As a veteran in the field, Gow is leery of latency claims from vendors and how relevant they are to a production environment.

*"Anyone can put one box at the end of a piece of wire and push data down it and say we have zero latency between those two boxes. But when you put it into client network with switches and firewalls and hops, are those useful measures?"*

As market data experts, CJC examines client systems, designs, and builds and offers support, training, and administrative services.

To improve latency, the firm would

look at the overall design of a system, its configuration, the amount of network connectivity, network speed, and performance of the computers. At the Intel fasterLab it will have access to equipment from, for example, Sun and Cisco, choices of other hardware and infrastructure switches - and be able to replicate client production systems.

*"They can connect their application and start tuning to system to make it go faster."*

Gow also expects to work with virtualisation, leveraging Intel's many close partners in this area such as VMware, to see if it can meet the demands of high-speed trading with the benefit of virtualisation and not lose performance. Although some financial services technologists think virtualisation will be too slow for trading, Gow wants to test it and look at results. Early results in the fasterLAB from Cohesive using AMQP sw and Elastic Server showed effective distribution of trading scale workloads (using Pantor's OPRA simulated feeds) - and subsequent performance runs with other AMQP implementations from RedHat and iMatx have demonstrated that this open source software is potentially ready for prime time in the testing front office environment.

*"New technology arrives every day. Until you can benchmark it against the system then we don't know how fast it can be. You hear stories that virtualisation slows down processing, but does it? With the Sun and Intel guys at the lab - they are serious technical people - so we will see if we can overcome these problems - we can see significant benefits from access to today's leading technologies via the fasterLAB - we can contrast and compare."*

## News in brief

### STAC: Harpertown Cuts RDMS Power Consumption 16%

Reducing power consumption in the data center is a high priority for securities firms today. Intel's recently released quad-core Xeon 45nm Harpertown processors decreased the server's overall power consumption by 16% (333 watts vs 395 watts) and increased the server's market data efficiency (updates per second per watt) by 19 percent over the previous generation Clovertown processor, STAC found. It ran a market data load of 2 million updates per second across a stacked RDMS configuration that utilized as many of the 8 cores in this system as possible.

### SPMC, Swift & Red Hat Present at Intel's fasterCITY event in London 23 June

John O'Hara SPMC will provide context, Carl Triefoff will explain the value of integrating Red Hat's AMQP Messaging, Realtime Linux, and Grid Computing into a high performance distributed computing platform. STAC has recently published a report which discusses record-breaking RDMS 6 performance on Red Hat Enterprise Linux 5.1 and Intel Xeon 5160.

## CALENDAR

### fasterMESSAGING AUTOMATION and STP - Mon 23rd June 2008, 6.00pm | BAFTA Theatre, 195 Piccadilly, London W1

The life blood of financial services and global commerce is built around the exchange of messages between counterparties to trades and transactions. FasterMESSAGING AUTOMATION AND STP taps Intel's sight of market innovation to paint a cross section across this change.

### Sail Intel '08

July 10, 2008 | RAF Yacht Club, Hamble, England  
Capital Markets Technologies Ltd. defend the 2007 trophy.  
Info: [www.intel.com/uk/sailintel08](http://www.intel.com/uk/sailintel08)

### Intel Developer Forum

August 19 - 21, 2008 | San Francisco, CA | Moscone Center West  
Intel Developer Forum (IDF) is where Intel and the ecosystem come together to share their latest innovations and vision for the future of technology. Come hear industry news, learn about trends, and discover opportunities to integrate what is discussed into your projects. Hear breaking news from senior Intel leaders in the seven Forum Keynotes.  
• [http://www.intel.com/idf/index.htm?id=idf08+em620\\_1](http://www.intel.com/idf/index.htm?id=idf08+em620_1)

### At SIFMA

Tuesday 10 June 09.00-11.00 | "Rev your Morning" Intel and STAC seminar on measurement, benchmarks and latency  
East Suite, 4 th floor, Hilton Hotel, 1335 Avenue Americas New York  
Contact [Erika.rieger@metia.com](mailto:Erika.rieger@metia.com)

Tuesday 10 June 18.00-20.00 | Intel - SUN Microsystems - partner reception "Innovation through collaboration"  
East Suite, 4 th floor, Hilton Hotel, 1335 Avenue Americas New York  
Contact [Erika.rieger@metia.com](mailto:Erika.rieger@metia.com)

Wednesday 11 June 18.00-20.00 | Intel and STAC reception "Rev your Evening"  
East Suite, 4 th floor, Hilton Hotel, 1335 Avenue Americas New York  
Contact [Erika.rieger@metia.com](mailto:Erika.rieger@metia.com)

## Acceleration - Applying Niche Innovation

Celoxica, a company which was spun out of Oxford University in 1996, is applying its accelerator technology, which combines hardware, firmware, APIs and services to reducing the latency in market data. Lee Staines, the company's CEO, claims the solution can capture market data faster than any alternative available today.

*"That means that our clients can react to changing market conditions more quickly and trade faster than their competitors. Early performance testing has shown that our solution can deliver under 2 microsecond latency and extremely high bandwidth -- over 10 million messages per second."*

The company uses field programmable gate arrays (FPGAs) that can fit onto a board with Intel processors without requiring more space or power in a data center.

*"We use the hardware to accelerate specific functionality that lends itself to these techniques. We also offer the distinct*

*advantage that our boards can be re-purposed so that they are adaptable and extensible to meet changing requirements."*

That's important because one key complaint has been that FPGAs are difficult to program and relatively inflexible. If the gains are high enough, there will be no barriers in this area, which was born out in April by several participants at the Intel fasterMARKET DATA meeting Intel organized in London, who said they expect hardware solutions such as FPGAs to play a role in the mix going forward.

Commenting on this, Nigel Woodward who chaired the London event, "increasingly we are seeing high speed infrastructures for trading being designed and engineered for optimized performance. Given a clean sheet of paper – rarely available – a mix of technologies could be deployed. Accelerators for particularly high workloads acting as offload engines to mainstream CPUs supporting the firms key trading, risk and analytics applications. We

have to work closely with our partners such as Celoxica, Xilinx, Nallatech Xtreme Data for example to achieve this. Until recently it was accelerators or CPUs, looking forward its now an architected combination.

Offload of specific functions to accelerators can also help the power consumption issue.

Lee Staines of Celoxica is placing some new bets, targeting the company at the market data area – leveraging experience on tap from exchanges and market service providers.

Adding to Woodward's point regarding architecture – using tools such as Intel's V-Tune, mainstream applications are now being tuned and optimized, which closes the gap massively on some early distorted comparisons between the performance of respective emerging and mainstream technologies.

## HP Hardware Collaboration at High Speed

An early supporter of Intel's initiative to establish the fasterLAB was HP – providing hardware and partner products to build out the environment.

You can't get far away from HP in trading. Buy-side firms, major investment banks, and the London Stock Exchange (LSE) run on HP servers and they are all looking for more speed, more power in a smaller footprint with reduced energy consumption.

TABBGroup's Bob Iata was quoted in Advanced Trading saying that HP is the vendor of choice going forward on Wall Street.

*"HP is the only major server vendor we see growing in this market, HP has won the war,"* said Iata.

To get to that position, and stay there, HP works closely with the Intel fasterLAB in London, while engineering staffs from the two companies maintain close contact in the U.S. In fasterLAB, HP and Intel collaborate to tune and optimize the LSE's TradElect application to

meet the demands on the marketplace. This Collaboration extends to their mutual clients & ISV partnerships.

*"A lot of the business HP does in financial service in the UK is directed to the investment banking community,"* says Paul Kember, sales director at HP for the UK financial services industry. Once a rack market, financial services is increasingly looking to the features and flexibility of blades, and HP has met the demand for high power in tight quarters, where limited opportunity to bring in more electricity means its BladeSystem c-Class architecture which can run Intel's latest power-efficient processors has highly relevant and attractive features.

The City of London has some real power issues. The c-Class blade servers reduce power consumption by 15-20 percent, said Jon Buxton, HP Pre-Sales Consulting, while allowing increased density. *"We actively manage the power workload on all the blades we deploy,"* said Buxton.

*"We are supportive with Intel in top to bottom initiatives with other important partners such as Mellanox, Voltaire and Cisco in the high speed interconnect environment,"* adds Kember. *"We work very closely with the Intel account teams to identify customers who would benefit from time with Intel engineers in the lab to fine tune their stack. We invest in equipment for the Intel lab on an ongoing basis, ensuring that this program backs up the client and partner dialogue of our respective teams. Today, with the emphasis on optimisation of infrastructure – where Intel and HP play - our working relationship with Intel is closer (and more effective) than it ever has been before."*

HP is also working with RMDS at Reuters and Wombat at NYSE Euronext. No surprises – both are looking for faster market data which requires faster connections, tuning and integration of the surrounding services. The hardware vendors are catalysing the trading ecosystem.

## Kx - Software Which Uses Every Available Core



Multi-core processors require multi-core capable applications, explains Simon Garland, chief strategist at Kx Systems, an extremely fast in-memory and on-disk database that is a favorite of high-performance trading firms.

*"A lot of database, application and CEP vendors talk about how they are multi-threaded, but they aren't really," he says. "They say they are multi-threaded, but they have a database thread and a user interface thread. So if they are running say 32 cores, they are just using two and the rest are sitting idle."*

Looking for a software company which could make use of a growing number of cores, Intel approached Kx.

*"We thought this would be a brief connection until a bigger company came along," says the unassuming Garland. But software suppliers have been slower to make use of multi-core than Intel had expected, leaving Kx to grow its relationship with the processor provider. David O'Shea, Intel's relationship manager for Kx in New York commented "Intel sees the value in the expertise of specialist software houses – these companies provide the relevance of our technology to the market and we have to find ways of leveraging our respective capabilities and expertise – small is (often) beautiful, and we recognize this at Intel."*

*"If you come to us with 64 cores, we can use all of them for proper parallel execution," adds Garland.*

*"In the standard benchmark for comparing us to someone else on the same machine, we*

*can look unbelievably fast because we are using the full machine and someone else is using just a few of the cores."*

One reason that many applications don't speed up with multiple cores is that they scatter cache around the cores as the application moves the processing from one core to another.

*"The OS internal book-keeping gets so complicated because the data is spread over many different caches, which all have to be checked to make sure they are still current and safe to use."*

Kx's application, Kdb+, benefits from running very tightly written code – about 100 KB -- which has been fully rewritten to run multi-core.

*"We had already produced versions of our code for Windows, OSX, Linux and Solaris on Sparc and Intel. We have rewritten it to make use of multiple cores. By contrast, a big application might take hundreds of man-years to rewrite all the code, find the unexpected interactions, fix the bugs and do QA."*

To achieve speed, Kx's original product was written for distributed processing on multiple machines back in the 90s. Now with 64-bit computers, Kdb+ can run large applications on a single machine. But because it was written for distributed processing, its design matches the requirement of multiple processors and multiple machines. For now, Kdb+ makes use of multiple threads, but Garland expects the marketplace will focus more on multiple processors on multiple machines soon because some of the problems haven't gotten so big.

Subprime, for example.

*"Of the 100 firms that have been hit by subprime, probably only five or six are set up to properly value what they have. Most of them can't do it, but those who can are in a much more interesting situation."* They can run what-if scenarios against portions of their portfolios to see the impact of a certain interest rate change and then determine which five percent they really need to focus on in the face of changing market conditions.

*"It's a big competitive advantage to be able to crunch through your portfolio and value it. You either value your whole portfolio at the individual holdings or you do general reporting based on a guesstimate. There is nothing really in between. With Kdb+, senior managers can ask for a report on what happens to our Florida exposure if interest rates move by X percent and get the results a few hours later or even within minutes. Once you have a really powerful way to view the number you are in a totally different place from your competitors."*

Kx engineers work with Intel engineers and provide input into future designs and suggest approaches that would speed up processing.

*"They share part of their roadmap with us so we can focus our development on software to take advantage of upcoming features. We are way, way faster than anyone else, and as long as we keep doing it on the latest hardware, we will continue to attract customers who have demanding applications."*

## Measuring Latency is the First Step to Reducing it



With the demands of algorithmic trading, the drive to find and eliminate all possible points of latency has become something of an obsession. When you are hunting for milliseconds and microseconds, a stopwatch simply won't do.

So when the Intel fasterLAB was setting up, a top priority – right after getting the hardware and networks installed – was to bring in the measurement experts to generate numbers that would be credible and useful.

Intel leveraged its relationship with Endace, which develops hardware that can drive latency measurement applications, and Trading Metrics, which has built software for the task.

*"For measuring packet delay in digital networks, you need very accurate time stamping, and frankly, NIC (network interface cards) give awful timestamps,"* says Steve Gleave, vice president at Endace.

Endace developed a solution to this problem with their data acquisition and generation (DAG) cards, that work on any interface (circuit, packet, IP and InfiniBand), at any speed, and for any packet size. Packets captured on a DAG card are immediately time-stamped, in hardware, using a clock that is synchronized to an external source

with atomic accuracy. Time-stamps are 64-bit records, offering nanosecond granularity.

*"The DAG card is fundamental to the entire Endace product portfolio. Once you guarantee to capture every packet, time-stamp it immediately, in hardware, and then write it directly into CPU RAM, applications are free to process the data as they please,"* explains Gleave. *"Data can be load-balanced, duplicated, filtered or statistically summarized. For latency measurement, Endace DAG cards can time-stamp the same packet at different points within a trading architecture, thereby enabling partners like Trading Metrics to turn this data into trending information."*

With accurate views of latency, firms can assess where they will get the best ROI from technology investments in their infrastructure to improve the speed of operations. Nigel Woodward of Intel elaborates that *"this is the fundamental raison d'etre of the fasterLAB – to provide a sophisticated sandbox in which to simulate the trading environment and the increments to be gained from engineering at this level."*

Endace and Trading Metrics have worked

together on measuring latency in the Reuters Market Data System (RMDS), and helped launch the Reuters Latency Monitor (RLM) service for measuring latency figures between a number of end points in the trading life cycle - from exchange to customer premise, in the on-site feed handler, the market data hubs, and the point-to-point servers. In addition to conducting diagnostics to uncover latency and help determine what, if any, action to take, the monitors provide alerts during the day if latencies grow beyond a defined range and allow important historical comparisons to be made.

*"Our NinjaBox-LM was designed with Trading Metrics for the RLM service,"* adds Gleave. *"It's a self-contained, passive monitoring appliance for accurate latency measurement. We aren't being used to rip out every nanosecond but we are helping firms to know where the delays are in their trading cycle. In the case of RMDS sites we can measure all points from exchange to trading platform. Once delay issues are identified, diagnosis and remedy can be applied."*

Getting the facts moves discussions past





a blame game between traders and data network operators. By pinpointing issues, the infrastructure teams can work significantly more effectively to meet the demands of the business.

*"It is interesting to see how much money is thrown at latency before people measure it and look to see where it is and where it isn't. We are all about helping people measure the delays very accurately. As official time-keepers of the low latency race, we're helping to pinpoint the difference between winning and losing."*

#### • Trading Metrics

Banks will see real value in using the Intel fasterLAB because they can obtain reliable measures in an environment that simulates all the complexity, and potential latency of real production, says Marie Giangrande, director at Trading Metrics.

*"Clients are using our software to meet needs in trading, operations and compliance. The most common call is one where the customer's algo engine is missing the market. The second most common need is with customers who*

*are obsessed with reducing latency end to end. Finally we also hear quite often that performance improvements in infrastructure upgrades did not meet the expectations the vendors have set."*

She recalls a client in London which tested new hardware and an upgraded network in a sterile environment and achieved a 50 to 70 percent performance boost. Back in the real world, however, the new investments led to a mere 10 percent improvement. Now the firm is using Trading Metrics to go through its production systems, benchmarking performance, and changing network or CPU configurations and then testing again to check for performance changes.

In this case the bank believes that it faces a combination of issues, Giangrande adds.

*"Latency can be introduced at so many different points – by the application, by the throughput into a server and by external events like a huge spike in message traffic during a hot market. Looking at a single instance of performance is not sufficient, especially in this marketplace with extremely variable traffic volumes combined with the capacity limitation*

*of exchanges."*

The key to improvement is a holistic approach to the problem and precise measurements at each point in the process. Trading Metrics measures data flows across the enterprise to catch latency in networks, applications, and CPUs.

*"Our correlations and comparisons help users diagnose the points causing latency. Over the years, we have found the biggest source of latency is applications that have been developed in-house. Code is old, has been grown organically and has not taken advantage of the tuning features available today from most vendors in the stack."*

Together, Endace and Trading Metrics enable Intel fasterLAB users to see definitive numbers in a large-scale environment that accurately simulates real-world production. Exacting comparison can be gained into the respective contributions of infrastructure technology and software tuning to the latency arms race.

## Working close to the silicon

# Pantor's High Performance Gets 20 Percent Boost at the Intel fasterLAB

"Pantor Inside" could become a new tagline for the Stockholm-based company which is achieving very high performance transaction processing on the latest Intel generations of processors. It uses the Intel fasterLAB to test and modify large configurations of the Pantor Framework and Pantor Fast Virtual Machine.

Because Pantor's technology is used by several investment firms and is on the leading edge of differentiating their trading activity – the only public reference it can use is from SunGard Front Arena which has announced it uses Pantor's FAST Virtual Machine in its trading solution.

Designed to handle extreme spikes of data without inducing latency or dropping packets, Pantor Engineering is working with the Intel fasterLAB on continuous improvement of its already extremely high performance products. Pantor Presto, for example, encodes FAST more than ten times faster than the original 2005FAST proof of concept using benchmark

data from OPRA.

*"The lab gives us tremendous benefits because we can test new hardware models before they are launched,"* says Rolf Andersson, CEO of Pantor. *"Whenever we have a client with demanding requirements on speed, we can test the latest and greatest from Intel in engineering white boxes or from production OEMs which is very useful. And we are also able to test large configurations and combinations of infrastructure components to assess optimum contributions to the latency race. With 45nm quad core Harpertown we used six fully equipped boxes with 10Gb and 1Gb networking and got very, very good results. Obviously we wouldn't have seen that at home because we don't normally have access to that many boxes in our lab."*

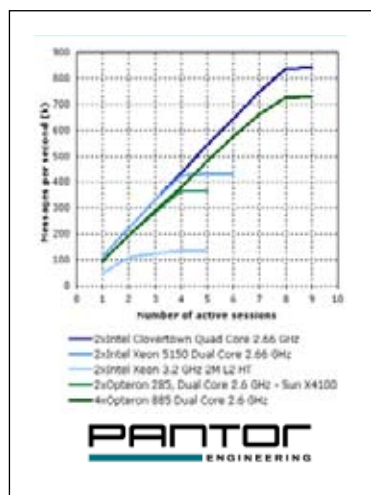
In addition to providing the latest processors, fasterLAB gives users the chance to work with Intel's latest and fastest hardware,

such as its 10Gb NIC, code-named Oplin (Intel 82598 10 Gigabit Ethernet Controller), and its extensive range of software tuning tools – many of which are not widely known in the market.

Without access to several computers, a firm can't determine if its solution will work in realworld production conditions. Users have often discovered that solutions which worked well in a small lab break down in large-scale production.

*"You need this to test horizontal scaling for market data or market feed systems,"* explains Andersson. *"You want to know that it actually scales to 30 parallel pipes and that it will remain fast if you cascade data to hundreds of users. Testing that requires a lot of machines, and ISVs rarely have the equipment to do that in their own lab."*

*"With Harpertown we found linear scaling in places where we didn't think we would. Software performed better than we expected"*



*"Pantor ORDO was able to process and generate best bid offer messages for 300,000 two-sided quotes per second with an average client to client latency of 500 microseconds."*

*and then when we tried it on a large box with multiple CPUs we found it was linear all the way."*

The demands on systems are growing constantly. This summer, options data is projected to reach 900,000 messages per second; in three years the markets have seen a 9x increase in volume, more than 100 percent per year, he says. Messages per second is a measure, but it doesn't tell you how data is flowing within your system.

*"Things are happening within the second, so you should look at your latency period and have that as your measurement. Millisecond – you did that a year ago. Now aim for a tenth of a millisecond. In that time you have 90 messages, six frames. We can do 8 million messages per second, or 8 per microsecond. You have 125 nanoseconds to code your message. Is this relevant? Can we expect to handle several messages in 100 microseconds? Yes. One stock price change can generate hundreds of*

*changes in options prices. As soon as you want to trade with the large market players you may have to decode a message in one-eighth of a microsecond."*

Anders Furuhed, Pantor's chief technology officer, explains the performance improvements the company achieved at fasterLAB.

"Pantor ORDO, a high-performance, fault tolerant application server and gateway configured for a derivatives exchange mass-quoting benchmark, was able to process and generate best bid offer messages for 300,000 two-sided quotes per second with an average client to client latency of 500 microseconds. The low latency of Intel's recently launched 82598 (10Gb) Ethernet controllers with I/OAT technology, combined with the X5440 processors, is key to achieving this level of performance as the fault tolerance built into Pantor ORDO is based on networked servers. Intel's I/OAT technology has helped us to decrease the mean latency by 20 percent while

simultaneously improving transaction rates by 15 percent in our ORDO framework."

*"FAST is purpose-built for market data and for high speed trading and stat-arb trading", adds Andersson. "It is also well-suited to hardware acceleration."*

Using a new Intel Core 2 machine, Pantor did 8 Gbit/s of market data feed with eight parallel feeds from the same machine. It nearly saturated one 10 Gbit/s pipe, he adds. Previous tests were not capable of enough throughput to meet the capacity of a 10 Gbit/s connector but the lower overhead delivered by Oplin enables unprecedented I/O rates.

*"Sixty-six million messages per second was the aggregate peak rate on all parallel feeds with the same network interface. It was to show how quickly you can get data in and out of a modern machine – it was more I/O testing than CPU testing."*

*"Intel's I/OAT technology has helped us to decrease the mean latency by 20% while simultaneously improving transaction rates by 15% in our ORDO framework."*



## European Union Regulation - Global Indicators

### The race is on ... mind the gaps

An industry-wide gap has opened up between new regulatory requirements and current trading infrastructure capabilities in a fragmented market. Today, there are 260 European securities execution venues and 35 of them trade the 1,000 most liquid stocks. In practice, connecting to all of them may not be an option, but getting the best deals and prices and monitoring that deals are done in the clients' best interests will be ever more important to retain or gain customers and keep regulators at bay.

At JWG-IT's recent (15 April 2008) London FORUM, representatives from 30 financial institutions acknowledged that they will be challenged to prove their best execution policies were followed. The majority do not believe the board is aware of the imminent risk of a challenge from the corporate customer (the primary threat) or regulator.

At the FORUM, our Technical Special Interest Group (TechSIG) demonstrated its capabilities with collaborative solutions to liquidity search, record keeping and trade reporting exhibited by combinations of its 30+ member base including Aleri, BEA Systems, Business Objects, Cisco,

Copan, EMC, FRSGlobal, GemaTech, Intel, SGI and Sybase.

The FSA has made its point by imposing the first jumbo fine of £900,000 for careless record keeping, citing it as a "serious failing ... even where the FSA has not determined that the firm mis-sold subprime mortgages and there have been few complaints." Clearly the regulators do not see their investments over the last few years as meaningless and they will now be entering a period of selective and prudent investigation. Maybe the financial penalties will not hurt, but the damage to reputational risk of a trading business in a fragile market undergoing structural change is a clear danger.

EU regulators, though their "principles based approach", are calling on financial institutions to define industry best-practices and avoid 'di minimis' solutions which do not provide transparency to protect European customers. To get it right, a balanced view of the risks and the appropriate investments must be agreed by the businesses and their project managers, the Chief Risk Officer and the technology / operations teams. Priorities must be set and plans established which leverage the "run the

bank" team's ability to use new platforms and operating models to change the bank. Smart order routing, rules-based technology, enterprise data management, monitoring systems and balanced scorecards are just some of the tools required to free up the capacity required.

The board's challenge is straightforward. Deploy the right operating model and you maintain your edge; get it wrong and your costs and operational risk profile explode and more regulatory hurdles will be established. Cooperation across Europe is required to get an agreed league table of priorities. Under the auspices of JWG-IT, over 200 sell-side, buy-side and market infrastructure practitioners started this process in London and the debate will continue at the Frankfurt FORUM on 16 June, to be followed by Milan and Paris later in 2008. New research will form the basis of the discussion, allowing delegates to benchmark their views via anonymous voting. Don't miss the German market's views on 16 June.

• [www.regonline.com/jwg-it-FORUM-FRANKFURT](http://www.regonline.com/jwg-it-FORUM-FRANKFURT)

### JWG-IT Sponsors Technology Special Interest Group (TechSIG) for MiFID

In working with several leading Tier1 firms, JWG-IT has found that most believe that 80 percent of the processes which need to be changed for MiFID compliance are non-competitive and that solutions can be most effectively generated through collaboration.

TechSig provides neutral safe platforms for like-minded firms, vendors and professional service firms to collaboratively develop practical IT infrastructure solutions to resolve MiFID compliance problems. It is based on the belief that MiFID, like most EU regulation, could be implemented quicker, cheaper and better through collaboration across the technology value chain. SIGs take a deeper dive than think-tank forums on specific issues with trusted technical experts from their membership to define implementable IT solutions.

We have together identified:

- A comprehensive framework of over 1,000 requirements that need to change
- What member banks are going to have to do, and their priorities
- How this impacts reference architecture. There is now the opportunity to participate in the special interest groups to:
- Save time and money by applying collaborative insight to your specific needs
- Jointly develop reference solutions which will inform and contribute to your own regulatory programmes
- Be seen to represent the broader financial services community.

For more information and current members see:

• [www.jwg-it.eu/TechSIG-Main](http://www.jwg-it.eu/TechSIG-Main)

If you wish to join, please contact [info@jwg-it.eu](mailto:info@jwg-it.eu)

# NYSE Euronext is...

- 1 An exchange
- 2 A secure financial network
- 3 A software provider
- 4 An ultra high speed market data provider
- 5 Hosted software as a service
- 6 All of the above

**Answer: 6**



*"For the big guys who are concerned about infinite configurability and are able to create the world in their own image, we sell software,"* explains Ken Barnes, VP Hosted Solutions at Wombat, a member of NYSE Euronext Advanced Trading Solutions. *"For everyone else who wants SaaS, we can roll these offerings up to provide a service that resides on a safe network, and the entire environment is hosted and managed. Users leverage it as a service and the whole operation becomes a lot less expensive to run. And we constantly invest in R&D to make sure the platform is the fastest and the best in the world."*

Fastest and best explain why NYSE Euronext Advanced Trading Solutions is working closely with Intel to obtain the highest possible performance from Intel processors.

Barnes sees the exchange world evolving into a group of mega powers including NYSE Euronext, NASDAQ OMX, Deutsche Borse, CME and the LSE and a second tier of regional exchanges, such as Malaysia and Brazil.

*"We could partner with them and leverage our technologies, know-how and network and that would allow them to focus on customers."*

A key competitive advantage in the global

market place is the speed of the Wombat Market Data Platform which NYSE Euronext acquired in January. The NYSE Euronext Advanced Trading Solutions group works closely with Intel to get the best possible performance from Intel processors to maintain its leadership in high-speed, high-volume data handling.

*"We have a deep R&D relationship which is helping us develop the next generation Wombat Data Fabric middleware,"* says Barnes. *"We are working closely with Intel whose architecture we find valuable because it is increasing its front-side bus in line with compute capacity. As the Nehalem processor becomes available, we expect even better performance because we do so much I/O between the CPU and memory, and the new chip will have a great impact on our ability to deliver more messages per second."*

Wombat Data Fabric has shared memory for remote direct memory access from application to application across servers without any use of the CPU for network traffic, allowing the applications to share memory efficiently across a network at unprecedented high speeds. And yet it is all accomplished under the hood of Wombat's mature MAMA API, presenting a set of familiar publish/subscribe semantics to

developers.

*"When we rolled out Wombat Data Fabric, we decided to use the same API so application developers don't have to worry about writing to a new interface. They don't have to change anything with their trading applications or their EMS apps already on MAMA. You can deploy those apps across a distributed fabric, or put them all on the same node and allow them to share messages at millions of messages per second in single digit microseconds. A year ago you couldn't find a server that could handle these volumes, but now with 8, 16, and 32-core machines this is actually becoming possible."*

In the market data arms race, Barnes sees NYSE Euronext Advanced Trading Solutions and Intel converging in platform capabilities.

*"We are trying to bring computers closer together through shared memory and remove as many of the layers as possible in between each application and the processor it runs on. That is happening on a large scale with co-location, and on a small scale with the Wombat Data Fabric and its shared memory which streamlines IO, whether that be server to server or core to core."*





## fasterCITY Event Series

### Providing the Essential Industry Pit Stop on Technology & Business Convergence

March and April 2008 saw a packed schedule in London with 2 fasterCITY events in quick succession – fasterTRADING on March 4 and fasterMARKET DATA on 28 April.

Both events attracted packed audiences to this now established regular forum in The City's calendar which puts technology developments in the context of the markets; key issues.

At fasterTRADING George Andreadis of Credit Suisse discussed the European challenges of multiple emerging MTF execution venues, the quest for speed and smart sourcing of liquidity, the service levels of smart order routing and placement strategies – these are all in the mix for a viable sell side business increasingly under pressure to innovate to stay ahead. Peter Randall, characteristically laid down the gauntlet from a venue point of view by reporting record volumes – finer spreads and market share and a new technology optimised infrastructure to keep Chi-X at the forefront of the exchange business. This was matched by the usual innovation of Kevin Houston discussing the new BidROUTE venture, representative of the new trade value chain innovation that is now possible with low barriers to entry from high performance industry standard technologies (encouraged by Intel!) – and the potential this offers for better liquidity and price discovery across the markets diversity of stocks, buyers and sellers.

Tony Bishop, ex Wachovia CTO and now representing Verari with Dimi Ziakis Principle Engineer at Intel for Financial Services provided the technology context to the subject.

Tony commented on the complexities and sensitivities in engineering today's trading infrastructures. This included the extreme project management disciplines needed to manage all the moving parts and their contributions to optimised performance. Dimi discussed Intel's investments in technologies which directly impact the low latency arms race and increasingly underpin many of the market's new trading infrastructures as they gear up for greater trade volumes and speed. Looking at how solutions are built, Dimi described the work conducted in the Intel fasterLAB to simulate and test the effect of advanced technologies on latency reduction.

The usual panel, this time with representatives from Reuters, Verari, BEA, Fujitsu Siemens, and NYSE-Euronext engaged in discussion around new business drivers affecting service providers to the market and the technology industry's ability to provide cost and environmentally effective technologies to support the change underway.

FasterMARKET DATA on 28 April completed the loop on the trading front office by focussing on this (exploding) life blood of the markets. Mark Reece of The London Stock Exchange discussed how the exchange has been chasing the ever increasing trade volumes and has seen how volume directly correlate to the level at which they publish prices to the market. In their quest for speed of both execution and data dissemination Mark discussed the tuning activities they are conducting for InfoLECT and TradELECT, which include close

collaboration with Intel and partners around fasterLAB. Peter Lankford of STAC brought the audience up to speed with developments to establish an industry standard for latency measurement – facilitating both comparison of sell side service and technology investment selection. Rolf Andersson of Pantor Engineering elaborated on the need to for detailed and advanced architectural design to get the very best out of technology at the leading edge of the front office and the advances in take up and performance of FIX's F.A.S.T market data compression standard.

A fascinating panel of Reuters, FixNetix, CJC-IT, Celoxica and Endace drew the session to a conclusion by discussing the range of market data related topics from legacy replacement of old distribution platforms, proximity to venues and how to harness geography, to the use of emerging new technologies in the acceleration space.

Both sessions – in usual format then advanced swapping notes, making new friends and agreeing actions to the backdrop of tasting of fine wines – compered by Chief Wine Office.

**The next fasterCITY event  
fasterMESSAGING AUTOMATION and STP  
will be on Monday 23 June - information  
and registration.**

<http://www.intelfasterfs.com/fastermessaging/>

**NIGEL WOODWARD  
London, May 2008**

Series sponsors: **FUJITSU** COMPUTERS  
**SIEMENS**

Event sponsors: **Verari Systems, Microsoft Corporation, Reuters, 29West, Petra Financial, BEA Systems, Celoxica, m35**

