

# Innovating across

## Reducing technology costs with standardised platforms

by Nigel Woodward, global director of financial services, Enterprise Solution Sales, Intel Corporation



**Let's put aside the usual rhetoric of solutions, partnerships, ecosystems, and strategy and get practical.** Today, financial services organisations must invest in technologies that work together and provide reliable, sustainable, cost-effective solutions across the trade lifecycle – from buy side to sell side and venue, from front to middle and back office. This means covering all key functions: market data, trade generation and routing, risk management, accounting, settlement, and payment.

The challenge is figuring out how to navigate your way through the many choices, from CPU to business interface, and to set the right architectural direction. This includes chip design, operating system, middleware, and networking and means balancing innovation, performance, strategy, and tactical response.

Looking up from its safe haven at the bottom of the solution stack, Intel uses its expertise to help lead the way. This means keeping a sharp eye out for innovation coming into the marketplace and selectively using its lab-based engineering skills to propagate and accelerate time to market for new solutions.

Let's begin with an assumption – that today's powerful, affordable solutions are based on using an operating system from Microsoft and underlying processors from Intel.

Why Microsoft and Intel? Because these companies produce affordable, standardised technologies that can ensure consistency for any compute task and lay a strong foundation for a next-generation technology strategy. Companies worldwide universally use Microsoft operating systems, although the company must prove itself capable in high-end workloads away from the desktop. Intel is championing industry-standard, low-cost hardware and achieving the marketplace's best performance through throughput and low latency figures with trading workloads.

Everything starts with market data, so that's the right place to

start building a technology story. Market activity is reflected in prices and drives behaviour at the high end. This area has established "Big Brand" players with great technology, but there is still room for innovation.

Consider Neotick from the highly demanding world of gaming. Using a Microsoft-Intel environment, Neotick has built new code to handle market data, from feed handlers to consuming applications.

The Neotick Gen V\* trading station manages data through to the trading engine and displays it to the trader. It handles technical differentiation by combining a fast language, XNA\*, the Neobus\* data transport messaging bus, Intel processors, and selective use of proprietary graphics from NVIDIA for display.

As a result, Neotick's Gen V trading station can process hundreds of thousands of market data updates per second and refresh the screen faster than the eye can see – achieving sub-millisecond execution without disrupting the trading process. It offers a three-dimensional interface that permits a trader to personalise the display with ladders, graph animations, and market characteristics in a virtual world.

"The core problem in financial services is that the sheer volume of data coming in is so high that when you try to render it to your screen, it freezes your computer," explains Belal Faruki, president of Neotick. "That's why we offload the screen refresh to a video card while keeping the CPU to do the serious work."

Low latency has received a lot of attention in recent years, and for good reason, says Faruki. The fastest systems can detect price changes and act on them before their competitors, who wind up paying the price the first movers charge the market laggards. Neotick's solution is designed for arbitrage price opportunities in commodities like oil and gold – opportunities that typically exist for just a fraction of a second across multiple exchanges. The traders with the fastest system win.

Faruki cites a niche hedge fund scenario where traders are working on an 18 millisecond lag. As the hedge fund manager told him, "We get there first,





# the trade life cycle

and by the time the traders' orders come in, they are buying from us." A firm using Neotick's solution with 300 microsecond latency can hit the market with 20 orders in the time it takes a user with 9 milliseconds of latency to get in just one order.

"We have stayed true to our mission," says Faruki. "Neotick will offer one-way transmission time of 300 microseconds. Neotick has streamlined its development process so much that it can rewrite the system every year to take full advantage of new developments in technology innovation across the stack – Intel processors and Microsoft operating systems, development techniques, and so on."

Especially today, as the market restructures and balances, the focus in financial services is across all market participants. While large banks command the headlines, hundreds of firms make up the market's fabric. But not all are technology shops – nor should they be. To facilitate access to this innovation, Neotick will offer a hosted solution connected to some of the most optimal network real estate, ensuring the lowest possible traffic times or providing an in-house system for larger firms. Initially co-located to serve the financial hubs of North America, Neotick's technology can now connect customers globally to exchanges and most derivative volumes.

With the data inside the firm, the next task is to get it into a form where it can be seen, understood, and then driven into the engines that generate trades, and view and manage risk.

Xenomorph is making significant progress in providing a visual and productive front end to market data and the mathematics that sit below it. Xenomorph cleanses and analyses the data before the numbers hit the risk systems with its TimeScope\* product for the centralised management of reference and market data and pricing analytics.

"Poor quality reference and market data and pricing valuations can fatally undermine the success of any risk management system," says Brian Sentance, chief executive at Xenomorph. "Poor quality data leads to GIGO – garbage in, garbage out. While in a rising market that might not matter, during a downturn, with increased regulatory scrutiny, inaccurate incoming data won't be tolerated. It is increasingly raising concern among regulators."

Regulatory concerns are not just something to put in footnotes on financial reports; they have a direct impact on profitability. Sentance adds,

"Some years ago, regulators challenged a US investment bank's numbers on derivatives' valuations. The company improved information management to allay the concerns and freed up \$1 billion in regulatory capital which, at least at that time, seemed like a lot of money.

"Firms want to make sure everyone can see the same data and get access to the same models," explains Sentance. "The more people who can see and understand the data, the greater the chance that any errors or problems will be picked up quickly. Risk management is the function where the terms and conditions, historical data, and pricing have to come together across all asset classes. Risk managers just need a consolidated view so they can improve the data and feed it into the risk models."

While the priority at European banks appears to be risk management, in the US many buy-side fund managers are improving their data management and analytics in a search for trading opportunities, says Sentance. "They want to use larger data sets and analyse them quickly and have the data running on something more robust than a spreadsheet," he explains.

Running on Intel and Microsoft, Xenomorph's solution offers impressive speed. Xenomorph developed its own fast vector database. Using this solution, value at risk (VaR) calculations can shrink from hours to minutes compared to similar mathematical techniques on older systems. Since the data analysis challenge will only grow, this is the right side of the road to be on.

While the vector database is fast, some users want to work with standard tools and languages. Working with Microsoft SQL Server database developers, Xenomorph has now hosted its proprietary database functionality inside SQL Server to provide the best of both worlds – both de facto standards-based functionality and proprietary performance. The company has tested it on a 12-terabyte database with 750 billion price records.

"The performance in data retrieval was comparable with our own object database, but it also had other benefits in improved scalability," explains Sentance. "A number of clients have expressed interest because they see it as another way to leverage their existing knowledge of databases within the organisation. Microsoft is part of the technology fabric of firms, so this helps extend its effectiveness from desktop to server and mission-critical – rather than having someone who

learns about a specialist proprietary system or language.” The company is also working with Microsoft’s High Performance Computing Group to apply that technology to pricing and analytics.

While many firms have centralised their data management to improve quality, only some have applied the same attention to the quality of pricing and analytics. This is especially important for structured products. As firms move toward more multi-asset trading, Xenomorph’s TimeScape\* solution allows them to combine database technology, real-time and historical data management, and end user data analysis through a unified and transparent view to access a firm’s data.

Insight Investments, a top-ten UK asset manager, wanted to develop a comprehensive derivatives capability that included straight-through processing for over-the-counter (OTC) trades along with a golden copy of all static, market, and pricing data. It chose TimeScape. The Xenomorph product provided the combination of flexibility for users and central control of the data in a platform that would support any type of derivative in high volumes while maintaining complete records for any future audits. TimeScape allows the firm to add its own proprietary models. Once registered with TimeScape, they can take full advantage of the speed of the database to model and price new trades intraday and to run overnight batch processing of analytics and valuations. Users at the asset manager can also add new instruments and models, which means they can launch new products faster.

Further down the line, in the complex area of trade generation and management and becoming established on the trading Formula 1 circuit, is QuantHouse. Founded in 2003 out of a hedge fund business, Paris-based QuantHouse has designed its technology to provide clients with low-latency data directly from exchanges over fibre optic lines. QuantHouse is a leader in the new disciplines of systematic trading with out-of-the-box solution capability. While the Markets in Financial Instruments Directive (MiFID) and Regulation National Market System (Reg NMS) have brought best execution and smart routing technologies, systematic trading raises the bar again on processing performance with pro active construction of scenarios reacting to market signals. The need is for a sophisticated development framework with thousands of objects and plug-ins so that firms can easily develop high-speed trading applications and update them in-house as market conditions change.

“We provide leading-edge technology for hedge funds, CTAs, and prop trading desks so they can develop models and run them in the market

automatically,” explains Stephane Leroy, head of global sales and marketing at QuantHouse.

Deployed in-house or co-located, the QuantHouse solution takes feeds directly from exchanges and, with a feed handler built with multi-threaded code, takes direct advantage Intel’s multi-core processors. Rather than develop a custom trading application, the firm has designed framework on Microsoft .NET. Called QuantFactory\*, this framework includes a large library of plug-ins that trading firms can use to quickly develop, test, and execute their strategies.

QuantHouse is a regular visitor to fasterLAB, Intel’s low-latency lab facility in the UK. In March 2009, its feed handler was tested on the 2.93GHz Intel® Xeon® processor 5570 series. The results raised the bar, with 3.3 million messages decoded per second compared with 2.0 million on 2008’s Intel Xeon processor 5440 series.

“What makes this unique,” adds Leroy, “is that QuantFactory is delivered within Microsoft Visual Studio, so clients can use their .NET, C#\*, or G\* to program and develop their models in-house. We are packaging sophisticated technology development features for wider availability to smaller firms that do not want to be tech shops. The technology demands are huge, the learning curve can be steep, and the risks are large when it comes to developing automatic trading systems entirely in-house.”

Leroy says that as the market moves from human traders at screens to fully-electronic, computer-based trading, hedge funds and prop desks must move quickly to stay current with the market leaders or be on the wrong side of an arbitrage deal.

QuantFactory as a framework brings years of technology IP to market in the form of an off-the-shelf product that enables firms to develop new strategies and modify them as markets change – without waiting for a technology provider to act. Firms have the flexibility to develop and execute new, complex strategies at will. QuantHouse has thousands of objects and plug-ins that function almost as easily as Lego blocks, so traders or their IT staffs can create new functionality, test it, and have it functioning fast enough to benefit from new opportunities.

While big is definitely not necessarily beautiful, the future belongs to the hedge funds and other trading operations, insensitive to economic dynamics they simply see the market as a source of multiple signals and underlying trading opportunity.

Says Leroy: “They can read the market data trends and trade on them faster than the competition because their systems are automated. For those firms, 2008 was a fabulous year and

2009 will be as well. Computer programs don't get caught up in market emotion; based on the rules in their programs, they just search the market and buy or sell for quick profits, making money from even small opportunities in pricing. And for companies like QuantHouse, which are selling technology to silicon traders, this is the beginning of a wonderful expansion."

Kee-Meng Tan, managing director of Knight's Electronic Trading Group in Europe, adds: "QuantHouse's end-to-end, low-latency solutions allowed us to reduce time to market and to provide our European clients with a unique, advanced platform for dark liquidity."

The world now has a heightened awareness of risk, and there is an increasingly active prerogative to ensure risk is not only seen, but also controlled and managed. The risk discipline is complex and a challenge to implement effectively.

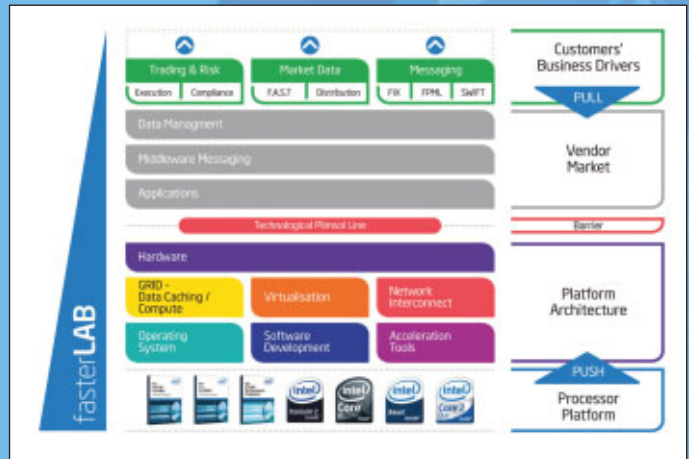
In the same functional area, Progress Apama have become established with leading products and functionality surrounding their CEP (complex event processing) capability. Apama powers capital markets applications like algorithmic trading, real-time risk, market surveillance, market aggregation, and smart order routing in equities, FX, derivatives, and fixed income markets. In recent extensive lab tests across the range of the company's products on the new Xeon 5500 series processors – the Apama smart order routing system achieved a 2.12x improvement in latency over Xeon 5100 series introduced just 18 months ago in 2007.

Today more than ever, managing the risk from trading and credit functions is high on the agenda. Ultimately, we need to get "risk" close to the point of trade, dynamically managing exposure rather than controlling under constricting limit structures.

New entrant Hyper Rig claims some ground-breaking progress. By redesigning code from the ground up and combining it with the horsepower of new, high-performance processor technology, Hyper Rig manages access to transaction data and performs multi-dimensional analysis of counterparty and asset class exposures. Hyper Rig recently carried out a performance benchmark to prove the advanced computational performance of its solution.

"Real-time risk management also requires a change in risk management, from after-the-fact monitoring to collaborating with traders. The change is underway," says Michael Coleman, chief executive and chairman of Hyper Rig. "We maintain the state of a firm's exposure in real time and we have very fast calculators that are able to do Monte Carlo simulations significantly faster than the current industry norm."

Hyper Rig accomplishes its high-speed



processing on standard commodity computers using Intel processors and Microsoft Windows Server 2008 using a test portfolio designed to represent a large trading book from a Tier 1 organisation containing one million trades (70% of them vanilla fixed-income and FX products and 30% exotic instruments such as interest and credit default swaps, options, etc.).

Hyper Rig computed a benchmark of profit and loss, Greeks, and full Monte Carlo value-at-risk with 5,000 iterations after a two-minute burst of 100 price updates and new trades per second, with a 37-second latency on a cluster of 16 four-core machines using the latest Intel Xeon processor 5500 series. This performance means that large institutions can now perform full-risk simulations for their largest books before putting on a trade in a matter of minutes instead of hours. A full stress testing of all volatility and correlation assumptions – which implies an exhaustive recomputation of all scenario and risk results – took less than 20 minutes.

Due to threading technology on an optimised modular proprietary grid and caching technology, Hyper Rig achieves near-linear scalability as processors are added. Alternatively, Hyper Rig can be integrated into third-party grids such as DataSynapse\* where these manage the high-performance computing (HPC) infrastructure. Within Hyper Rig, all of a bank's positions are maintained in the cache so total risk is updated with each incremental change. This simplifies the process and eliminates the need to move large amounts of data across the enterprise to a single location for batch processing. The software can process millions of trades per day and handle sharp bursts in volume, reducing traditional compute times from hours to minutes. As the regulators demand larger data sets to be analysed across more scenarios, the focus of risk will be in data handling.

“By maintaining the state, we are not managing huge volumes of data in one chunk,” says Coleman. “We don’t need to move the data around networks, and we can calculate changes very, very fast.”

Market research by Hyper Rig indicates that financial firms want systems that can integrate and manage data from multiple sources in a way that is both streamlined and cost effective. They also want systems that are flexible and scalable to accommodate both organic growth and rapid expansion through acquisition – plus the consolidation that is now taking place across the markets.

Hyper Rig includes a data management component that operates in real time with proprietary extract, transform, and load (ETL) tools. The system offers a dashboard through which users can register any type of data source – including databases, Web services, XML, queues, or spreadsheets – through a simple drag-and-drop interface. Users can review the data graphically, open the data to examine it, extract schemas, and transform them. From the dashboard, they can also impose data management rules, from simple alerts to more complex rules governing system response to price volatility or very high data bursts. Once loaded into the high speed cache, the data is immediately available to users and subject to the management rules.

Hyper Rig is essentially both a data management system and a risk management program, fully integrated and tuned. Users can implement the platform faster – and likely for 20 to

30% less cost – than with separate modules.

Hyper Rig claims this kind of functionality is attractive to firms that cannot grow their (relatively) plain vanilla derivatives businesses because existing systems can’t provide real-time exposure. With growing resistance to highly complex derivatives, banks are pressed to increase their numbers of trades in more standardised products, but that requires highly scalable systems.

“Banks are shutting down their highly structured proprietary derivatives trading while trying to expand their more straightforward underlying business and operating it more effectively,” says Coleman. “Banks realise they have to trade products in which investors have confidence. The business is about reducing risk and maximising opportunity, and that is why it has to be in real time. You can’t trade into opportunities a day or a week after the event; you have to spot them as they are coming.”

QuIC Financial Technologies, a global risk management solutions provider, is hitting new speed benchmarks with Intel® processors on everything from a single-processor box running Windows XP to enterprise-wide calculations. Headquartered in Vancouver, Canada, the company has a global presence and a strong professional services capability to help clients implement solutions to their risk management problems.

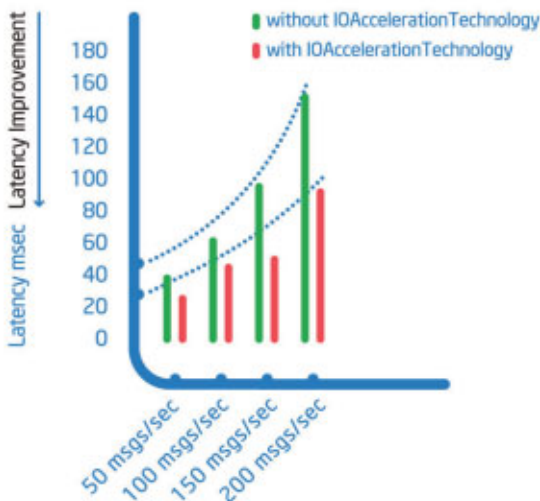
QuIC recently launched an incremental risk charge (IRC) solution to help banks comply with new IRC regulation from the Basel Committee on Banking Supervision, which requires them to measure and hold capital against specific risk that is incremental to those risks captured in the bank’s VaR.

“Since any bank with an internal VaR model will be affected by this pending regulation, we anticipate overwhelming demand from banks looking to implement their IRC solutions now, ahead of the regulatory deadline,” says Nigel Cairns, president and chief executive of QuIC. “Anticipating this demand, we welcomed the opportunity to work on an IRC solution for one of our long-term clients that approached us some time ago. Their solution is currently under regulatory review and, if approved, will be one of the first submissions to gain regulatory approval.”

The QuIC IRC Solution\* enables risk managers to rapidly implement an out-of-the-box initial model that can be easily modified to meet their internal methodology. This gives them complete flexibility to implement their IRC solution quickly and easily, to identify any regulatory gaps earlier, and to understand the potential impact of IRC on their business well in advance of the regulatory deadline.

QuIC, which began as a specialist in credit risk and then expanded into a comprehensive risk management company, has maximised the performance available from Intel processors to

### IO Acceleration Technology Reduces Latency in FIX Messaging





achieve speeds that have impressed both clients and prospects.

One global bank increased its risk analysis performance by more than 400% while dropping its IT costs significantly when it deployed QulC. Since implementing the QulC Credit Risk Solution, the bank's annual trade volume on a single type of credit derivative has increased by 70% and annual trading volume has risen by several million U.S. dollars while the bank has, at the same time, sharply reduced capital requirements for its credit-derivative trading books.

"There is a potential revolution underway in risk technology," says QulC's chief technology officer Tony Coppellotti. "Risk management is high-end data management, traditionally requiring large, proprietary servers – expensive in capital and running costs. Banks are coming to realise they achieve higher pure performance, and far better price/performance, running on the likes of Intel's industry-standard machines. While increasing performance by 400%, the bank was able to cut its portfolio run time from nine hours to two hours on QulC and reduce the number of servers by a factor of 10."

QulC's modular approach to risk management fits with the conundrum in which the market finds itself today – more pressure to manage group risk deep and wide and lower budgets for capital and operating expenses. QulC's solution integrates with and augments existing systems, fixing today's problem while offering a clear future direction.

"Today, no one wants to make big decisions on infrastructure," says Cairns, "so our ability to integrate with in-house systems or work alongside another vendor's technology and accelerate the performance is a real benefit."

Time to implement is an increasingly important requirement as boards come under pressure to provide better risk reporting to regulators and investors. For example, when a large hedge fund needed future earnings projections to be included in the fund's regular reporting, the risk management team was at a loss. Its existing in-house risk system couldn't easily calculate the risks and certainly couldn't do it fast enough for the demands of rapidly-changing markets.

"We always relish the opportunity to prove our technology's flexibility and value," says Cairns. "This was a chance to do so with one of the world's largest and best-managed hedge fund companies. The entire implementation process, from specification to development, testing, and global launch, was completed on time and on budget, within the desired three-month timeframe. Performance improvements are as promised, if not faster."

Cairns expects that the turmoil in financial markets will continue to drive growth for fast and

cost-effective risk management solutions. The pending IRC regulation, which requires banks to measure and capitalise issuer-specific risks on their trading books as a way to contain the risk of a counterparty failure, will add another layer of complexity to existing VaR models. The regulation requires, at a minimum, credit default and migration risks for positions subject to credit risk. It also requires capturing spread risk and equity price risk.

"As we've experienced firsthand, banks developing an IRC solution will quickly discover that speed, performance, and flexibility are critical, and extending an existing market risk solution to include IRC is not a viable strategy," says Coppellotti. "Regulatory guidelines mandate a 99.9% confidence level, which for highly-rated products can imply running ten-million-plus scenarios. Pushing an existing market risk solution to run so many scenarios will cause performance issues and have serious hardware and cost implications."

So, traversing the trade lifecycle, market data has been analysed and trading opportunities are identified and, we trust, managed inside acceptable risk parameters. Now we have to route the trade to the market, quickly and in volume. Exchanges now multiply into new forms, venues compete on depth of book and speed of execution, and the performance pressure conducts back down the value chain to routing – from trading engines or new direct market access (DMA) services.

Rapid Addition, with a heritage at the centre of the FIX Protocol, has long been at the forefront of trade messaging and was one of the earliest beneficiaries of engineering tuning for low-latency, working with Intel in the fasterLAB back in 2006. Recent work in preparation for the launch of the Intel Xeon processor 5500 series shows that its applications run 51% faster on the new processor – and that's straight out of the box, before any modifications to the software.

Toby Corballis, chief executive of Rapid Addition, is pleased with the results and expects that the company will see further benefits from the processor.

"That's a 35% improvement in throughput and a 21% reduction in latency without any specific tweaks," says Corballis. "With the various configuration changes we could make, and doing our multi-threading a little differently to fit the chip's profile, we are confident we could achieve even faster results. We are very excited about this entire new Intel technology. People should pay attention to what Intel has accomplished if they are in a business where performance matters."

Rapid Addition worked with Intel technicians in the lab this spring to make adjustments for even better performance.

Since its users value low latency for algorithmic trading, The London Stock Exchange recently announced it will use Rapid Addition's RA-Professional to develop an optional FIX 5.0 interface to TradElect. RA-Professional supports all versions of FIX, from 4.0 to 5.0, with low latency and high throughput performance that outpaces that of competing FIX engine providers, according to Corballis.

In laboratory tests, Rapid Addition's FIX gateway was able to process transactions more than 35 times faster than the best published figures currently available. RA-Professional is written in C# for the Microsoft .Net environment and can run on Windows, Linux, UNIX, or Mac OS operating systems.

Intel's continuous improvements to its processors both drive and support the low latency race. Just 10 years ago, says Corballis, 5 milliseconds was considered extremely fast. Now the company is measuring performance in microseconds.

Rapid Addition also works closely with Microsoft, especially on two key initiatives – a faster version of Microsoft's BizTalk middleware and applying the power of its HPC platform to trading.

"Some of our clients are already using HPC and other clients should be on it," Corballis says. "They absolutely need the capability to process large volumes of data in a very short time, and they can't do that on a single machine."

For BizTalk, Rapid Addition provides an adaptor called ShortCut for FIX. Developers from London spend time with the programming teams at Microsoft headquarters in Redmond, Washington, reviewing the next generation of Microsoft .NET technology.

"Microsoft has committed to making BizTalk low latency, which it is not now, and we are working very closely with them on that," says Corballis. Rapid Addition also continues to work closely with the Microsoft development teams in Reading, UK. Corballis adds, "Our engineering relationship with these major technology players – Intel and Microsoft – really pays dividends. It gives us a strategic edge. FIX and trade messaging are at the front of the low-latency arms race and there is no substitute for engineering excellence everywhere. Today, the infrastructure layers, previously discarded as commodity, are becoming critical.

"We make sure we are able to achieve outstanding results and we are absolutely committed to that relationship and to increasing the performance of our FIX and FAST engine," says Corballis. "Leading-edge customers ask about the firm's relationship to the technology platform providers and are pleased when they learn how closely RapidAddition works with Intel and Microsoft."

The importance of low latency hasn't gone away in the market downturn, Corballis adds. "I know we

are seeing fewer trades, but if you need to get to a price in the market, you can't afford to say we'll invest a little less in technology this year and run a bit slower. If you are going to take that approach, you might as well go do something else completely."

The trade is complete, analysed, constructed, risk measured and sent to the exchange, and confirmed back to the counterparty or prop book. So is post-trade the poor relation in the equation? It doesn't need to be, since new technology is coming to market to bridge the traditional front- to middle- to back-office divides and consolidate cross-asset class.

TwoFour's trading technology combines knowledge of business processes acquired from the complex systems that dominated the 1980s and 1990s, with the latest technology offering rapid time to market and lower costs. The company's multi-asset trading system is built on Microsoft Windows .NET, running on Intel processors, and supports seamless workflow across the front, middle, and back office, from capture to settlement.

Users can quickly configure the application, which provides speed and flexibility with a low total cost of ownership, and integrate it with their accounting systems. With over 26 installations, TwoFour is the success story in this area of the last 10 years. TwoFour has been replacing legacy in-house systems as well as once-popular, now-aging systems from major traditional vendors – which are proving an overhead in terms of inflexibility to new techniques and, even with capital expenditures written down, loading the cost of operation and data centre capacity.

Chris Davis, director of sales and marketing at TwoFour, says, "The software works out of the box and includes tools which, by minimising manual activity, reduce the total cost of ownership through user-driven configuration and also give clients control of their own destinies."

Rather than sit on a waiting list when they want an extension or modification, the clients' developers can make the change in-house using familiar Microsoft tools. Because the TwoFour System runs on Intel, it delivers top-of-the-line performance on industry-standard-priced hardware.

In the post-trade area, arena transaction and message sizes cannot be compared with front-office latency statistics; however, in a recent benchmark at Intel's fasterLAB, a prospective customer wanted proof that TwoFour could handle 400 trades a second for a sustained 15 minutes.

"We did that and didn't tax anything," says Davis.

So the tester stepped up the transactions: "We stopped at 4,000 transactions per second for 30 minutes," Davis says, "or a total of 7 million new transactions with external sources, blotters



updating in real time, and numbers reconciling back to the totals.”

The test was run on an Intel processor-based server which retails for about \$10,000. Davis adds, “The whole concept of these larger financial institutions needing massive hardware to process significant volumes is obsolete. We have conquered that with TwoFour.”

For TwoFour, the Intel fasterLAB has been a key component in its development and sales program. As a relatively small, early-stage company, it gives access to equipment which proves technology beyond normal affordability and gives customers peace of mind.

“The lab isn’t just Intel providing a facility and telling vendors to go install the software,” explains Les Gosling, head of sales in Europe, the Middle East, and Africa for TwoFour. “The level of support they provide is phenomenal, which is what you want when you’re undertaking a benchmark. Their support saves us time and money. When we hit a bottleneck, we may need to make some tweaks to the system to overcome it. If we were testing on our own, that would require us to buy more kit. At the fasterLAB, we might want to look at the impact of changing variables in the infrastructure (e.g., smaller disks and more spindles to address I/O congestion for the database server). It really works as a service, and it’s the kind of support that engenders loyalty by vendors and customers alike. We see it as being a responsible supplier.”

TwoFour can quickly test its software on different versions of Microsoft operating systems, Microsoft SQL Server, or database server configurations using the technology at the lab.

TwoFour has sold its software – which includes front-office trading, middle office, and back office clearing and settlement – to banks, brokers, asset managers, global wealth, and corporate treasury departments. In today’s financially challenging environment, the company’s initial sale often consists of a few specific modules for one department. Then, as the client learns how powerful and easy-to-use the TwoFour technology is, its footprint expands. The horizontal scalability of Intel® architecture is ideal for this incremental approach to legacy replacement. There is no need for big, up-front hardware costs and, in many cases, the new hardware pays for itself in the first year through data centre power and space savings.

Total cost of ownership is a major driver in TwoFour’s sales. The company estimates its systems provide significant percentage gains over current operations running on proprietary operating systems and RISC\* environments. Intel architecture has become de facto in the front office low-latency space. TwoFour believes that work with Intel has

proved that the vast areas across post-trade are now ready to be addressed.

“We aren’t talking about hundreds of thousands of dollars worth of hardware,” says Gosling, “we are talking about deploying server machines that cost a few thousand dollars.” Because TwoFour is written in C# within the Microsoft .NET framework, it can easily integrate with in-house systems and software from other vendors, meaning that banks or brokerages can start with one department and build out from there as the software proves its value. The fact that it runs on Microsoft, with many management features, means that firms can do their own modifications and interfaces rather than putting their names into a long queue with the vendor.

“Companies like to have control of their own destinies rather than relying on the vendor’s development schedule for tweaks to the system,” adds Gosling.

The TwoFour design team, learning from the way banks and corporate treasuries used previous systems it had built, designed its solution around workflow, says Gosling. Managing operational risk is built into the system and reconciliation is simplified through the use of dashboards to make the process more proactive than traditional systems.

“Workflow is relatively new in the marketplace; not many financial services products are workflow-driven by design.”

To summarise, there are many important functions along the trade lifecycle: inter and intra-party, data, and transaction processing. As financial markets return to a sustainable, normalised state, new demands will be placed on data, information, reporting, messaging, and connectivity. The demands on computer automation will also change, with new investments both viable and inevitable. Even so, the good news is that the technology landscape has changed, with new systems that can save money. Intel has moved to a position from which it is influencing the facilitation of solutions in financial services. From a position of relative independence deep in the infrastructure – as a de facto standard, Intel and its technologies are having a direct impact on performance. A collaborative business model enables support for innovation from many parties. Intel works across the board with independent suppliers using Intel architecture and technologies. We have just taken a quick look at recent activity in conjunction with Microsoft. In today’s world, there are many new ways to reduce the cost of computing through standard technologies and encourage optimisation through infrastructure engineering – it’s worth exploring.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

\*Other names and brands may be claimed as the property of others.