Researched and written by:

A-TEAMGROUP

MARKET DATA PLATFORMS

Infrastructure to Handle Your Data Challenges

June 2009

An industry briefing prepared by **A-Team Group** for

NYSE Technologies℠

Executive Summary

Demands on market data infrastructure within investment firms have increased in number and intensity. With algorithmic trading, pre-trade risk calculation and dynamic portfolio balancing all moving into the mainstream, it's no wonder that processing systems and infrastructure are being reinvented on an almost semi-annual basis. The deluge of immediate, relevant information that serves as input into today's trading decisions is mind-boggling when compared with just five years ago.

Existing market data architectures are stretched beyond original, decades-old design parameters, and with mixed results. Many systems do not adequately handle modern data volumes from the expanding myriad of markets, while others do not deliver data quickly enough. Precious few provide both. Increasing server capacity with traditional horizontal scaling architectures – to service new markets, high throughputs and additional data points - is enormously expensive. As a result, operations staff have been forced to weigh cost against speed and settle for something less than ideal. Clearly, a new solution is needed.

Rising costs of energy aside, recent improvements in multi-core chips and multi-socket computers have paved the way for efficient, super-scale computing in an economical package. Direct Memory Access (DMA) – used locally within a single computer or remotely over 10 gigabit or InfiniBand network I/O interfaces – permits a new approach to making large amounts of data available. DMA can deliver data at the lowest latencies possible – at mere nano-seconds of delay for millions of messages per second - to the most demanding applications while remaining environmentally and fiscally responsible.

Modern market data platforms should not only provide traditional data to brokers but should also serve the processing-intensive, low-latency transactions performed by smart order routers and algorithmic trading engines. This is a tall order, with few technology vendors up to the challenge – and even fewer providing products that can address all concerns. Nevertheless, it's clear that firms need to invest in an infrastructure that gives them agility and power whenever it is needed, while simplifying the acquisition of new data.

Ideally, modern application programming interfaces (APIs) abstract network topology and infrastructure concerns from the developer, thereby speeding time to market and providing a flexible operating environment. Combined with a full suite of feed handlers for the world's markets, new systems can give maximum throughput with minimal end-to-end latency while providing a host of functions for the enterprise's data demands.

Adoption of new technologies like DMA can afford operations managers the chance to have it all – and have it now. NYSE Technologies' Market Data Platform V5 co-opts DMA, Remote DMA and other technologies and techniques to introduce a next-generation infrastructure for today's trading environment, allowing developers to overhaul their high-performance market data systems while meeting a broad array of requirements within the enterprise.

Market Data Processing Has Come A Long Way

An exponential expansion of market data volumes in recent years has been caused by the confluence of a number of technology and business drivers – from the emergence of electronic communication networks (ECNs) to the advent of smart order routing systems and algorithmic trading engines. Coupled with the increasingly divergent data sets needed to support complex decision systems, there is now a requirement for high-performance data management infrastructures throughout the enterprise supporting high-frequency automated trading as well as traditional market data processing functions like securities master file updates, risk analysis, and back-office trade support.

Many of today's market data environments are built on infrastructure designed in the 1980s and 1990s, when new systems were needed to support electronic trading and to provide a backbone for distribution of real-time data to traders' and brokers' desktops. From securities master files to client holdings and fundamentals to sector analysis; from historical performance to the activity in the order book; market data systems became an intrinsic component of every trading desk.

Traditional market data systems matured organically: by integrating applications and data sets as necessary; by reworking components to meet new needs; by adding restrictions on access to limit exposure for fee-liable data; and by distributing data to remote systems for further processing and distribution. These platforms served the industry well until 2000, when quote and trading volumes began to skyrocket. Networking was expanded to distribute the deluge of data to larger and larger collections of servers for simultaneous processing. Adding computers, meanwhile, has meant larger computer facilities and higher energy costs.

At the same time, decimalization squeezed spreads. The adoption of algorithmic trading models – and the slicing and dicing they entail – caused a massive increase in volumes of orders and data. And new alternative trading systems and multilateral trading facilities boosted competition for liquidity, creating a broader array of market signals, as new venues provided new data points. The result: still greater volumes of data.

In the face of growing competition, traditional exchanges were forced to lower their fees, and old-school monopolistic settlement relationships are still being dismantled, further fragmenting liquidity. Smart order routing systems – which seek out that liquidity and execute on it – aggregate information from multiple venues, generating virtual order books that provide a consolidated view of the fragmented liquidity to drive algorithmic trading engines and transaction systems in order to guarantee best execution. And so the cycle starts again.

Anticipated capacity and the underlying messaging of many market data platforms derive from observations of market data 15 to 20 years ago, when sub-second delivery of perhaps as few as 500 updates (messages) per second was considered high-performance. Today, rates of 1.5 million messages per second (mps) are common in the North American markets – where planning for 3 million mps is not considered forward-thinking enough. Overall latency of data – from source to consumer – is now measured in microseconds and is quickly approaching the nanosecond range. Modern systems, especially when feeding algorithmic trading engines, must also minimize latency jitter –

the variance in observed latency – so the end application observes a predictable latency profile even during peak market data periods.

The time is ripe for change in the mechanisms financial institutions use to deal with data, long regarded as the lifeblood of the trading environment, but increasingly seen as an asset that needs to be collected and managed properly – and efficiently – if it is to deliver on its promised benefits.

Beyond a Quick Fix

This onslaught of mainstream electronic trading brings into focus two concerns that resonate today: networking and computer input-output (I/O) bandwidth limitations; and the lag time from a market event happening to the time it is available for processing (latency).

Proximity-location of primary computer centres and direct market access have eliminated the most obvious component of latency – the unnecessary distribution of quote and order data to a geographically distant processing centre and redistribution back to the end user. But the adoption of these techniques has brought the volume and latency problems solidly within the investment firm's walls, necessitating a more sophisticated market data system.

Firms have found their internal resources insufficient to deal with the complexities of market data processing at unprecedented extent and intensity, but they have dealt with the situation as best they can. A disjointed architecture has emerged where data is segregated into manageable chunks. Making use of sunk costs in existing platforms and infrastructure, firms have deployed disparate components to tackle each specific challenge as it came into play. The result is a system that cannot process the right data quickly enough. Many firms are looking to overhaul their systems again – and those that haven't are in desperate need of a quick solution.

With the economic environment forcing scrutiny of budgets everywhere – including those for IT – data managers are being asked to do more and more with less and less. Data volumes are going through the roof, while IT budgets are testing new floors. But competition among struggling financial institutions means that 'no' is not an acceptable response.

An improved solution is needed – one that requires a smaller hardware footprint, enables flexible deployment scenarios, features minimal time to market for critical applications, while at the same time demonstrating superior latency and throughput profiles. It's time for a modern approach to high-performance data management.

Location, Location, Location – And New Hardware

Regardless of the processor or network interface, one fact does not change: when microseconds of latency make a difference and the distance between cities is measured in milliseconds, proximity counts. Co-location or placement of large processing centres near markets and network points-of-presence (POPs) continue to drive many trading infrastructure decisions. The physical proximity of hardware at or close to the execution venue site – and the networking to connect it to the source data – remains the most expedient and economical way to reduce overall latency to sub-millisecond levels. Once that's optimized, the next focal points for latency reduction are the software and, importantly, the hardware it runs on.

Recent computer processor improvements have focused on adding more cores on a chip and more processors in a box. Clock speed has not significantly increased, and hovers around 3 GHz for typical systems, mostly due to localized heat-dissipation issues. As a result, chips are becoming more efficient and better priced for today's ecologically and fiscally sound operating environments.

Dual processor systems with quad-cores are commonplace in blade and rack configurations with octa-core chips promised by the end of 2009. On-chip design enhancements include multiple non-contentious paths allowing from 8 to 32 separate CPUs to process common data without incurring I/O penalties or excessive cost. This improvement, however, creates an I/O bandwidth bottleneck limiting the amount of the data that can be made available to these super-scale processing environments. Systems engineers constantly balance computational power and throughput with data availability and cost.

New 10 gigabit Ethernet, InfiniBand, and even 100 gigabit interfaces now provide data flow rates that can feed these hungry processors. This requires that the operating system and network stack are properly customized and use more direct communication methods such as Remote Direct Memory Access (RDMA) which can add as little as 1 microsecond of delay per interchange to the overall latency picture, while processing millions of messages per second.

Middleware and Applications

Increases in performance through hardware and infrastructure will require interfaces to new technologies, and optimization of algorithms to handle market data efficiently, all at a low level. To boost efficiency, developers use an abstraction layer to hide the vagaries of the network topology, communications methodology, and lower-level protocols (See panel on Page 6).

In the simplest of market data processing systems, a single application collects market data from the feed source, performs basic calculations, massages incoming data for storage and downstream utilization, then passes the information to other systems and processing functions.

More sophisticated, resilient and flexible implementations incorporate separate processes and computers for the collection and redistribution functions. Often, the collection process keeps a coherent view of each instrument and performs calculations on the incoming data, which can be responsible for as much as 30% of observed latency in the feed handler.

Reference Market Data Platform Architecture

NYSE Technologies' Market Data Platform V5

A reference market data system, here based on NYSE Technologies' Market Data Platform V5, includes data feed handlers (FHs) as separate applications that deliver processed market data over a logical bus to distinct applications. These applications – algorithmic trading engines, calculations engines, data stores, or display terminals – use a messaging API to interface to the bus and a market data API to abstract the components of financial instruments and market information.

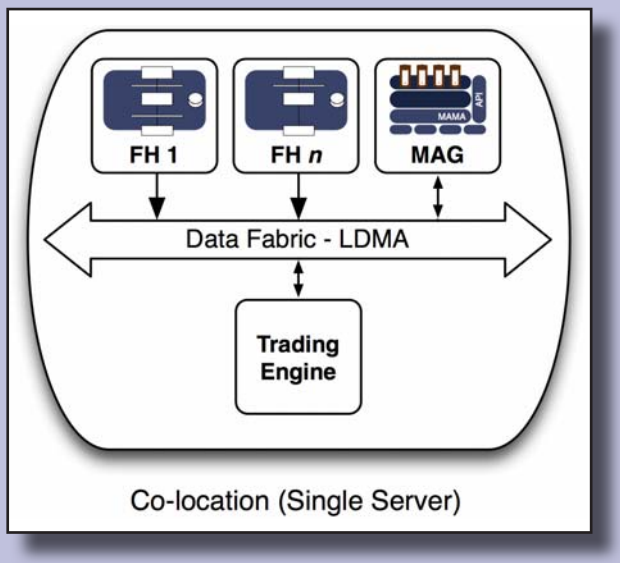

Communication with the logical data bus may take one of four basic forms: a) Local Direct Memory Access (LDMA), b) Remote Direct Memory Access (RDMA), c) TCP/IP fan-out, or d) other messaging middleware including multicast.

To minimize network configuration issues, the programmer uses an abstract messaging API configured to use any of the communications methods available as needed. For lowest latency, the LDMA option may be used where all processes run on the same multi-socket, multi-core computer. For large systems it may be necessary to use TCP/IP fan-out to downstream applications. But the programmer need not worry about the specifics network topography.

The market data platform also provides a view of information – about each instrument, transaction or other data carried on the bus – within the API for retrieving, manipulating and delivering updates. Depending on the implementation, the market data API may interact with the messaging system to retrieve, or publish, specific information from, or to, the bus.

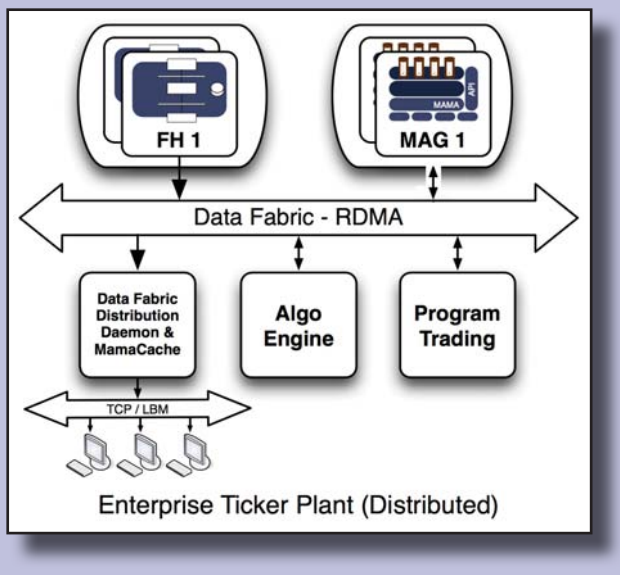

Efficiencies can be obtained by streamlining the primary acquisition function to simply pull raw data, normalize it to internal messaging formats, and pass it on to primary consumers at the same time it is passed to ancillary processes that perform value-add calculations or maintain the current state of each instrument. In this way, algorithmic trading strategies that require a consumable form of the raw data are given access as quickly as possible.

End applications – smart order routers, algorithmic trading engines, market data terminals, and risk calculation engines – require an array of networking and computational facilities in concert with various combinations of the data within the overall system. To aid in development, software libraries – application programming interfaces (APIs) – provide conveniences specific to market data and inter-process communication without encumbering programmers with the specifics of data feeds or network topology.

In essence, the programmer views the computing environment as a distributed source of data without worrying about the speed or size of processors, the number of cores or nodes available, or the specific location of compute nodes in relation to each other. The programmer requests a collection of data and the API aggregates it and makes it available. Alternatively, the programmer dispatches new data to downstream applications without worrying about which components may consume it. Such an abstraction permits flexible deployment of software and systems to multiple environments without intervention by the programmer. Instead, the messaging middleware is configured as demand dictates by operations staff versed in real-world needs.

These software abstractions permit the underlying middleware to be configured to use any transport or subsystem to handle requests and dispatch data. For example, Local Direct Memory Access (LDMA) within a single multi-core computer can reduce inter-process communication latency to sub-microsecond for millions of messages per second. A motherboard with four sockets utilizing quad-core processors with sufficient memory, disk space and I/O channels could easily handle common algorithmic trading functions including data collection and the trading venue interface.

For communication between computers within a single rack or cage, Remote Direct Memory Access (RDMA) over a 10 gigabit Ethernet or InfiniBand with proper networking gear, tuning and application distribution can keep latency in the 5 to 50 microsecond range – obviously slower than LDMA but a significant improvement over current norms. For example, communicating with a daemon dedicated to fan-out via TCP/IP may increase latency to about 100 microseconds due mostly to overhead in the protocol and network interface interrupts to the operating system.

Similarly, traditional multicast solutions or socket-based reliable messaging solutions, even with optimized implementations, have a wider degree of variance in performance and latency 'jitter' – the difference between the fastest and slowest updates. This typically results in even greater latency for distributed processing pressing into the double-digit millisecond range.

A Next-Generation Solution From NYSE Technologies

New approaches to minimizing latency – involving the use of Local and Remote Direct Memory Access (LDMA/RDMA) technologies – are presenting developers with opportunities to overhaul their high-performance systems while meeting a broad set of requirements for market data distribution within the enterprise. In essence, these techniques are heralding a next generation of market data infrastructures that allows programmers to do more with less.

An Accelerated Market Data Platform

NYSE Technologies, the innovative commercial technology unit of NYSE Euronext, has released its next-generation market data infrastructure, the Market Data Platform V5, to meet the growing and demanding requirement for high performance market data systems. This new platform offers rewritten, streamlined feed handlers, I/O acceleration and optimization, and ultra-high performance utilizing its Data Fabric messaging middleware. Market Data Platform V5 delivers significant improvements over previous generations of its products and establishes NYSE Technologies as a leader in enterprise-wide data infrastructures, high-frequency trading applications, and co-location solutions.

10 Times the Performance with 20% of the Hardware

NYSE Technologies' Market Data Platform V5 provides continuous throughput of 500,000 messages per second on a single CPU core (scaling linearly on additional cores) with less than 40 microsecond end-to-end latency (i.e. from raw market data input on a feed handler machine to a subscribing client application on a separate server), even at peak market data rates on today's commodity hardware, making it ideal for high-frequency electronic trading. Requiring one-fifth the hardware and delivering 10 times the performance of previous generations, Market Data Platform V5 reduces infrastructure thus saving capital expenditures and reducing maintenance costs.

Ultra low-latency and enterprise-wide distribution capabilities are a feature of NYSE Technologies' Data Fabric high-performance middleware. Data Fabric supports Local Direct Memory Access (LDMA), Remote Direct Memory Access (RDMA) and traditional TCP/IP fan-out, all of which can be configured to meet the needs of any application on nearly any hardware platform that supports Linux. Additionally, Data Fabric ensures that minimal latency 'jitter' – the range between the fastest and slowest arriving update – is very narrow, even during high volume bursts such as market open and close. As shown in Table 1 below, Market Data Platform V5, with Data Fabric LDMA and streamlined feed handlers, reduces processing core requirements dramatically leaving plenty of headroom for expansion and still maintaining less than 40 microseconds observed end-to-end latency.

Feed	# Cores	Avg CPU	Avg Latency	% < 40 µsec
CTA	1	11.8%	25 µsec	95.84%
OPRA	4	24.0%	20 µsec	99.00%
ITCH	1	10.0%	11 µsec	99.94%
OBU	1	21.7%	7 µsec	99.38%
ABM	1	25.8%	24 µsec	89.64%

Collected under real-world conditions showing percentage of messages processed and delivered in less than 40 microseconds from input on interface card of feed handler through to receiving client application.

Table 1: Core, CPU and Latency for Streamlined Feed Handlers

Flexible Deployment Options

Market Data Platform V5 lends itself to proximity hosted deployments where the feed handlers and trading applications can be installed on the same server thus netting nanosecond transport latency. NYSE Technologies' Market Access Gateway platform can be easily integrated to provide seamless order routing or distribution throughout the enterprise.

Within NYSE Technologies' Market Data Platform V5, the Middleware Agnostic Messaging API (MAMA) allows programmers to develop to an abstract view of the messaging infrastructure that hides the underlying processor and network configuration allowing flexible deployment scenarios with the same code base (See Figure 1 below).

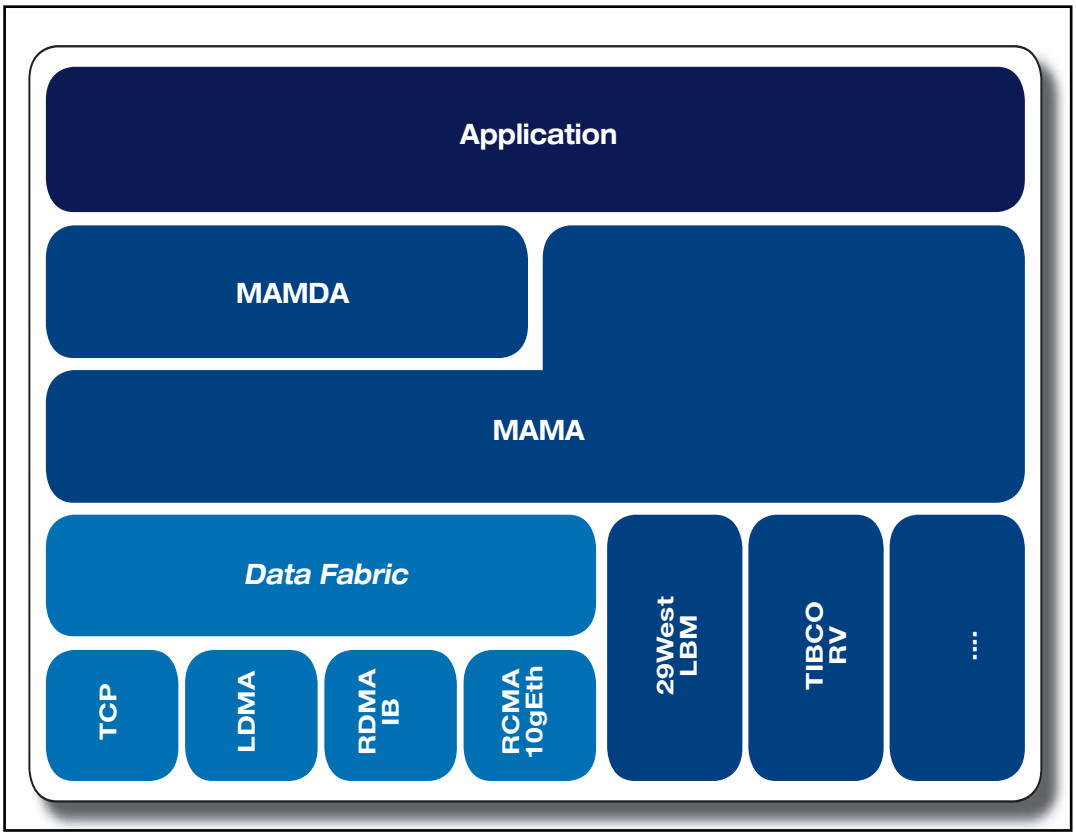

Figure 1: Application Architecture
Showing logical layers between middleware and logical data bus

Thus, Market Data Platform V5 can be configured in high-frequency, low-latency applications using Data Fabric LDMA. The desired latency profile and scalability parameters can dictate the required configuration (See Figure 2 overleaf).

Market data processing and manipulation is available through the Middleware Agnostic Market Data API (MAMDA) including the latest state-of-the-art order book processing algorithm permitting developers to focus on the value-added functions of their applications instead of the underlying infrastructure necessary to collect the data they need. MAMDA interacts with MAMA, when necessary, to access market data on the logical bus – whether it comes from a feed handler on the same processor, or across town in the enterprise data centre.

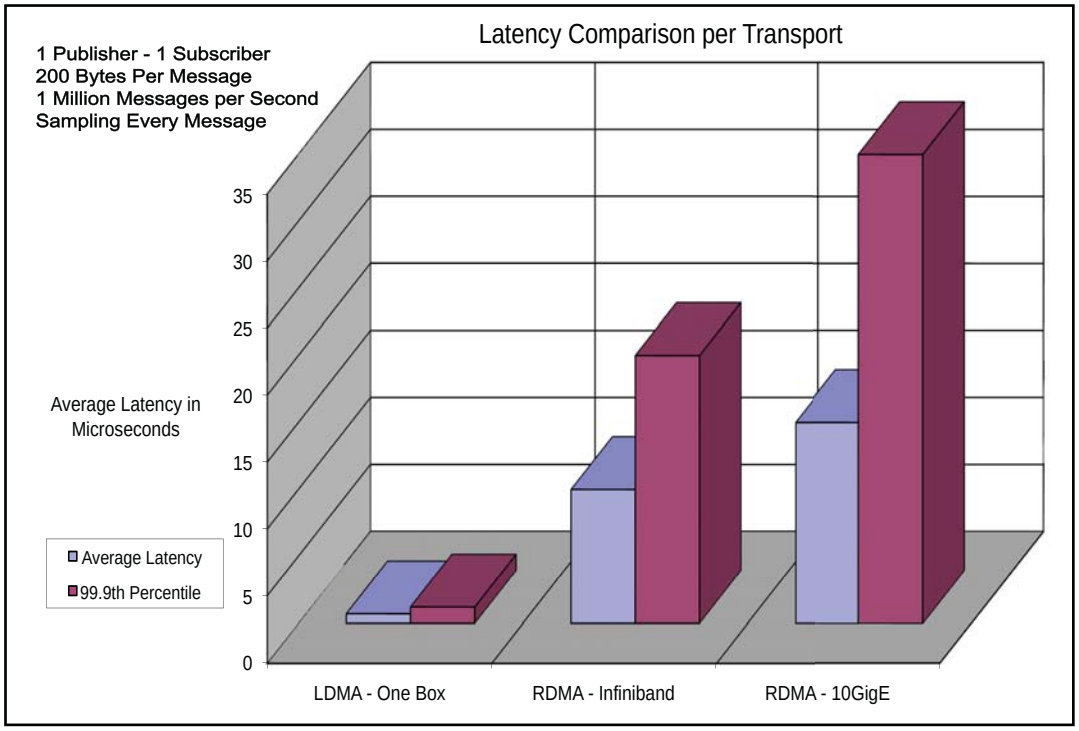

Figure 2: Comparison of LDMA, RDMA Infiniband, and RDMA 10GigE fanout Each transport method addresses a different issue providing different latency and throughput characteristics

Because Market Data Platform V5 is developed to the MAMA and MAMDA APIs, the feed handler improvements and Data Fabric enhancements are cross-compatible with NYSE Technologies' existing global Feed Handler Suite. An incremental migration is possible incorporating Data Fabric and the new feed handlers as they are needed (currently available for the fastest North American markets with coverage for all major markets soon).

NYSE Technologies' radical approach to modern market data throughput issues and low-latency concerns means Market Data Platform V5 leapfrogs current technology from other vendors to provide world-class capabilities in a truly scalable product. From single, ultra-low latency, black box implementations to multi-facility, enterprise-wide data solutions, NYSE Technologies' Market Data Platform V5 is a capable, proven solution from a global-scale technology provider.

For further information, please download the Market Data Platform V5 product sheet at **www.nyse.com/technologies**.

NYSE Technologies℠

www.nyse.com/technologies

A business division of NYSE Euronext (NYX), NYSE Technologies is a leader in providing innovative software, market data products and data management applications, connectivity solutions, exchange technology, and transaction solutions for trading firms, vendors and financial markets around the world. Offering the most comprehensive suite of best-of-breed commercial technology solutions that also include advanced messaging platforms, co-location systems, ultra low-latency order routing capabilities and a one-of-a-kind hosted networked community, NYSE Technologies' products attract and enable the free flow of liquidity for the global marketplace. With offices across the U.S., Europe and Asia, NYSE Technologies' advanced integrated solutions for fully-deployed, hosted, managed or on-demand services has earned the trust to power the trading operations of hundreds of global financial institutions and exchanges.

For additional information visit:
www.nyse.com/technologies

A-TEAMGROUP

www.a-teamgroup.com

A-Team Group, founded in 2001, is a publishing, research and events group that provides the global community of IT and data professionals in financial markets with the business intelligence they need to excel in their roles.

We're widely known for our focused series of publications, research and events across mission-critical functions including high-performance trading infrastructure, low-latency market data and connectivity, enterprise data management, reference data, market structure, risk and regulation and more.

A-Team Group's publishing division publishes a range of online and in print news services including Risk & Regulation IT, Reference Data Review, Market Data Insight, AsiaMarketsIT.com, Low-Latency.com, FinTech-Infrastructure.com, and A-Team IQ magazine. Find out more at www.a-teamgroup.com/publishing

A-Team Group's research division collates data, performs analysis and delivers results for internal product development and marketing, as well as publishing in-depth reports for external use, examples of which include:

- OTC Valuations - Pricing Assets in the Post Credit Crunch World
- Buy-side Firms Take a Hard Look at Data Practices: Derivatives Push the Data Management Envelope
- Buyer Persona - The Influence Behind Data Management Decisions, and more...

Find out more at www.a-teamgroup.com/research

A-Team Group's events division partners with companies to produce highly-targeted seminars on specific business challenges, attracting senior-level industry practitioners, such as its recent series of Algorithmic Trading seminars with Dow Jones across Europe and North America, and upcoming events with GoldenSource and Sybase.

For more information visit
www.a-teamgroup.com

A-TEAMGROUP

For more details on NYSE Technologies' Market Data Platform V5, please contact us:

Email: **NYSE-Technologies-Sales@nyx.com**
Web: **www.nyse.com/technologies**

Tel:
Europe: + 44 20 7987 5400
Americas: +1 (212) 510 3600
Asia Pacific: +65 6501 4100

NYSE Technologies℠