

# High Speed Systems Fueling Low Latency Trading

Sponsored by Lenovo

Srini Chari, Ph.D., MBA  
March 2015

<mailto:chari@cabotpartners.com>

## Executive Summary

*Cloud, Social, Mobile, Internet of Things (IoT), Big Data Analytics and High Performance Computing (HPC) are transforming the speed of decision-making in the financial services industry. But these trends are also creating several information technology (IT) challenges. Leading-edge stock and securities exchanges are overcoming these challenges to deliver higher speed and lower latency to their clients by investing in and deploying an agile, cost-effective, high-performance and secure HPC infrastructure.*

*Lenovo offers a robust portfolio of enterprise servers, storage, software, and support, providing a compelling value proposition for the financial firms. Based on the recently acquired IBM System x business and legacy Lenovo systems family, the combined HPC portfolio has been optimized for low latency trading, providing financial firms a competitive edge in a market where speed and economies of access are critical.*

*Lenovo HPC solutions provide extreme low latency, reliability and fast performance enabling financial firms to detect and act on profitable trading opportunities faster than competitors. These systems also reduce power, cooling and real-estate costs while providing a highly secure and reliable IT environment.*

*Together with a robust ecosystem of business partners, Lenovo HPC solutions are helping Financial Services firms achieve single-digit microsecond latency for high-speed trading, which translates to ultra-fast trading for clients. Lenovo System x has set numerous performance records and will continue to be a major participant in Securities Technology Analysis Center (STAC), the financial services industry standards group.*

Copyright © 2015. Cabot Partners Group, Inc. All rights reserved. Other companies' product names, trademarks, or service marks are used herein for identification only and belong to their respective owner. All images and supporting data were obtained from Lenovo or from public sources. The information and product recommendations made by the Cabot Partners Group are based upon public information and sources and may also include personal opinions both of the Cabot Partners Group and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. The Cabot Partners Group, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your or your client's use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document. This paper was developed with Lenovo funding. Although the paper may utilize publicly available material from various vendors, including Lenovo, it does not necessarily reflect the positions of such vendors on the issues addressed in this document.

Cabot Partners Group, Inc. 100 Woodcrest Lane, Danbury CT 06810. [www.cabotpartners.com](http://www.cabotpartners.com)

**Cabot  
Partners**  
*Optimizing Business Value*

## Key Technology Trends Shaping Financial Markets

The relentless rate and pace of technology-enabled business transformation and innovation in the Financial Markets are stunning. Several fast-growing intertwined technology trends – Cloud, Social, Mobile, Internet of Things (IoT), Big Data Analytics and High Performance Computing (HPC) – are profoundly disruptive and are reshaping the industry's economics in developed and developing markets.

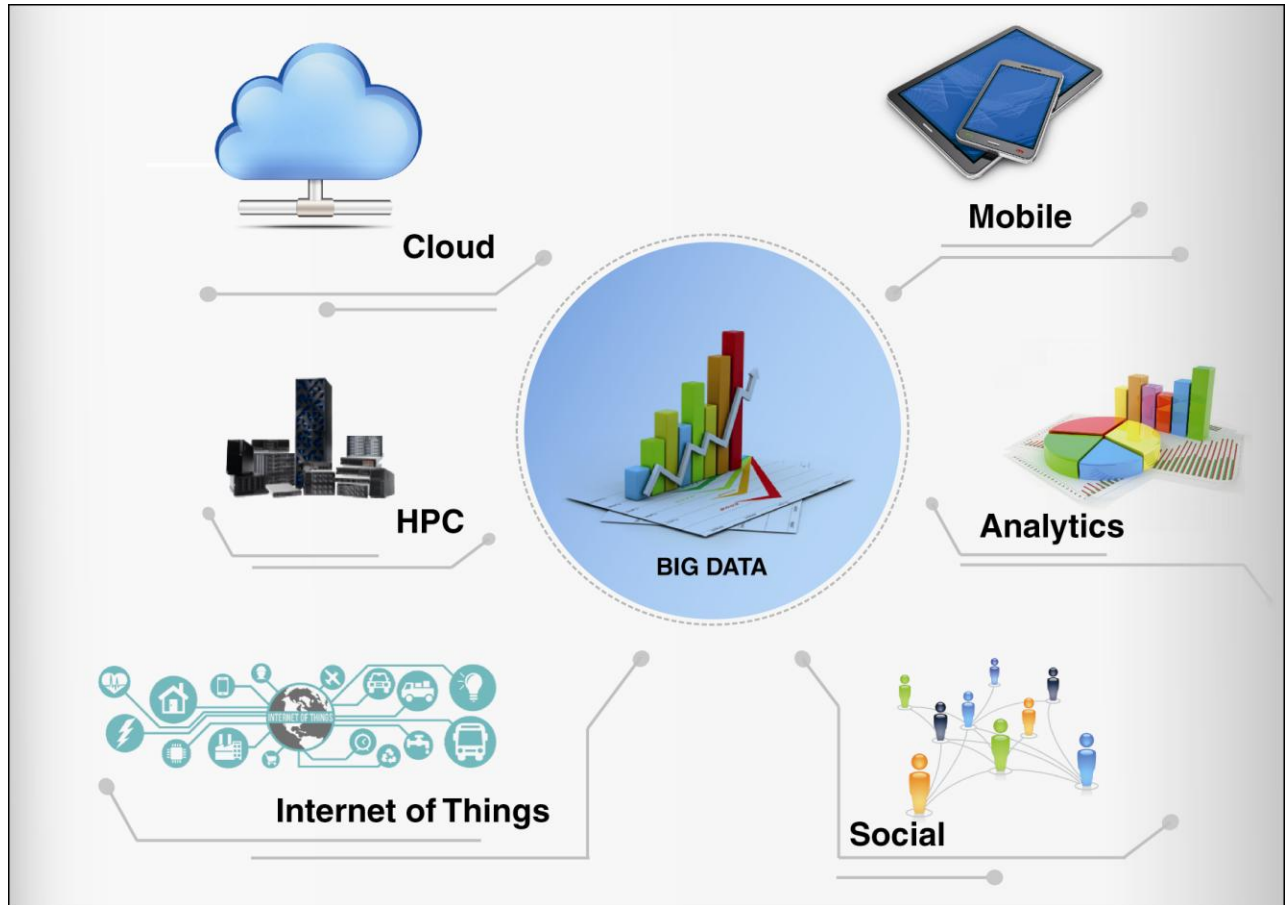


Figure 1: Intertwined Technologies of Cloud, Social, Mobile, IoT, Analytics and HPC

Financial institutions need to leverage these trends to get a competitive edge in finding liquidity, executing trades or calculating risk.

In addition, banks now have to compete with non-banks as digital channels and anytime, anywhere payments become the norm. New regulations are raising the cost of doing business. Securities exchanges are in an arms race to deliver higher speed and lower latency to their clients - investment banks, trading firms and hedge funds.

As a result, financial firms are increasingly investing in HPC infrastructure (servers, storage, networking, accelerators, software, etc.) for risk analytics and low latency trading to get a competitive edge. Consequently, HPC server revenues in Financial Services are expected to grow at 8.7 percent<sup>1</sup> annually much faster than the overall HPC market.

<sup>1</sup> IDC HPC Update at ISC'14

Cloud, Social,  
Mobile, IoT,  
Big Data  
Analytics and  
HPC key  
technology  
trends

Data is at the  
center and  
drives HPC  
and Analytics  
requirements in  
Financial  
Services

## HPC Provides Competitive Edge for Financial Trading

Financial firms are custodians of massive amounts of customer data but may not fully leverage it, partly because it is locked up in siloed databases. With the exponential growth of data from multiple new sources - for example, mobile, social, and web – new tools and technologies are needed to gather, analyze and visualize data in order to extract powerful insights from this data. These are some of the novel Big Data Analytics approaches. Many leading-edge financial firms are implementing HPC infrastructures to integrate real-time Analytics in the low latency trading process (Figure 2) to get a competitive edge.

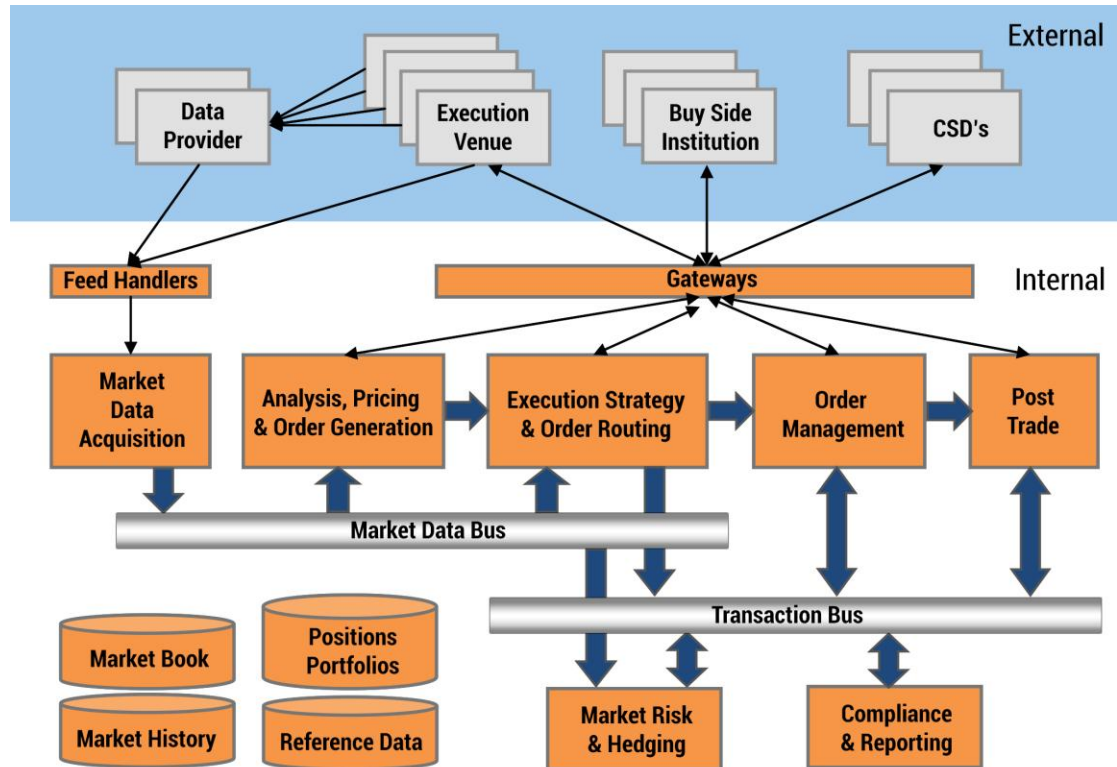


Figure 2: A Typical Technology-Enabled Trading Process

In financial markets, at the platform level, the NASDAQ stock exchange is able to consistently sustain an order rate of over 100,000 orders per second at sub-40 microsecond average latency.<sup>2</sup> Latency is the time taken to learn about an event, say a change in the bid, generate a response, and have the exchange act on the response.

In trading environments, both absolute and relative speed is important. Speed in absolute terms is important due to the inherent volatility of financial securities. Relative speed, in the sense of being faster than other traders, is also very important because it can create profit opportunities by enabling a prompt response to news or market activity. At a trader level, today's markets experience intense activity in the "millisecond environment," where computer algorithms respond to each other at a pace 100 times faster than the blink of an eye<sup>3</sup>.

<sup>2</sup> <http://www.nasdaqomx.com/technology/marketplacesolutions/trading/matching-execution>

<sup>3</sup> Low-latency trading, Joel Hasbrouck and Gideon Saar, Journal of Financial Markets 16(2013)646-679

*Analytics is being integrated into the low latency trading process*

*The end-to-end trading process requires an agile and secure HPC infrastructure*

*Absolute and relative speed of execution is vital*

*Exchanges are delivering higher speed and lower latency to their clients by investing in and deploying HPC infrastructure*

What is clear from all the above is that technology is reshaping the competitive landscape. The most likely winners will be those financial firms that embrace it. The challenge is that in most financial institutions, significant gaps exist between the business's aspirations and what they can realistically accomplish. Executives demand greater speed and agility, but their IT departments are unable to deliver, resulting in frustrations and mutual distrust.

Lenovo HPC solutions are being used worldwide at major stock exchanges, global investment banks, trading firms and hedge funds. The System x server infrastructure is powering exchanges, low latency trading, and massive compute grids.

## **Lenovo's Compelling Value Proposition for Financial Markets**

Lenovo offers a portfolio of servers, storage, software, services and support that provides a compelling value proposition for the financial services industry. Clients have access to Lenovo experts to work with them to design, tune, test and deploy low latency solutions.

In 2014, Lenovo acquired all of IBM's x86-based System x servers and System Networking portfolio (top of rack and blade switches, network OS, virtualization software). The deal included all the development and operating resources, including service and support. What is sometimes missed is that Lenovo also acquired all related intellectual property and that IBM will continue to provide maintenance services on Lenovo's behalf for a minimum of 5 years – as it did when it acquired IBM's Personal Computer division in 2005.

Lenovo also agreed to license and resell various enabling software for software-defined storage, enterprise file management, virtualization management, server administration software and firmware, and a bevy of systems management software and services such as IBM Platform Computing – including widely used IBM Platform Symphony and IBM GPFS – the leading enterprise big data management software. Lenovo also has an agreement to resell select IBM storage solutions tied to x86 server sales.

The System x acquisition complements an already existing Lenovo server business. While this business is insignificant in the US and Europe, Lenovo is a significant player in the server market in China. Lenovo offers dense, scale-out cloud systems and enterprise systems in China going as high as "4P" (4 socket) systems. They also have a joint venture with EMC, providing high-end storage; mid-range and lower-end storage<sup>4</sup>. The Lenovo acquisition moves the Lenovo server business forward significantly, making it the number three x86 server maker in the world.

## **Long-term Commitment to Earning Client Trust**

After the IBM System x acquisition, enterprise users can trust Lenovo to be committed to the financial services market for the long term despite recent security concerns related to the PC business. Lenovo has a track record of nurturing trusted client relationships.

Let's address the question of long term commitment first. The same concerns were raised back in 2005 when Lenovo acquired IBM's PC division. Ten years later, Lenovo's Think-branded PCs and laptops are still the gold-standard in enterprise PCs and notebooks.

*Lenovo's acquisition of IBM System x complements and enhances its existing server business*

*Lenovo has strong track record of excellence after the IBM PC acquisition*

<sup>4</sup> Lenovo's Plans To Acquire IBM's X86 Server Business Could Be Its Best Strategic Decision Yet, Patrick Moorhead, Forbes, 4/17/2014

*Substantial new investment in enterprise systems portfolio*

*Higher bar for ensuring product security*

*Lenovo has a wide and deep low latency portfolio of servers, adapters, switches, accelerators, and storage*

Further, Lenovo has retained much of the former IBM staff and expertise with excellent results. For instance, Lenovo won a record-breaking 61 awards at CES 2014 – the largest consumer electronics show. In short, Lenovo took a successful brand, and made it even better and more successful, recently celebrating the 100 millionth PC shipment.

Regarding its commitment to the server space, Lenovo has stated it has, and will continue to invest significantly in the System x brand to become the number one maker of servers. In fact, Lenovo has invested over three hundred million dollars in new Enterprise Business Group laboratories and facilities to innovate and enhance this portfolio.

### **Unmatched Product Security**

Lenovo takes security very seriously. System x has gone above and beyond the norm to harden its systems and development processes. This includes industry-leading features and practices to provide superior protection such as the System x trusted platform assurance, in which all code is securely compiled, assembled, linked, and digitally signed. This consists of two security processes with on-going rigorous validation to ensure a secure development process, digitally signed firmware, and a secure system execution environment based on secured hardware and secured firmware to establish and maintain a chain of trust. In addition, a host of other initiatives build on this foundation, including support for industry security standards – Trusted Computing Group and National Institute of Standards and Technologies (NIST) compliance – as well as other advanced features such as self-encrypting drives (SED's) and IBM Security Key Lifecycle Management (SKLM).

Further, since the IBM System x sale to Lenovo had to be cleared by the Committee on Foreign Investment in the U.S. (CFIUS), customers are assured of a rigorous security firmware development process. In fact, it can be argued that Lenovo set a higher bar for other vendors to meet.

And lastly, System x enables industry-leading advanced firmware attack detection in its supply chain to detect and mitigate tampering and future attacks.

### **Lenovo HPC Low Latency Portfolio**

The Lenovo HPC portfolio includes servers, the associated components, and a robust partner ecosystem to bring together the best technologies to address low latency and real-time analytics. A brief description follows.

#### **Servers**

Lenovo servers span four categories:

- Scale-up, enterprise-class X6 (X-architecture, 6th generation) servers. These are represented by award winning large-memory systems, the System x3850 X6, System x3950 X6, and the Flex System x880. These large-memory systems are optimized for Big Data Analytics workloads.
- Rack servers, including System x3550 M5 low-cost, high-performance server, System x3650 M5 server with integrated low-latency and high-capacity storage, and their



Lenovo ThinkServer counterparts - RD550 and RD650. The System x 3750 M4 offers great performance and core density. System x servers have additional RAS and manageability features, e.g., light path diagnostics, and additional memory options. The ThinkServers feature innovative and industry-unique features to support multiple storage types in the same drive bay and cost-effective high-speed connectivity options.

- Dense, scale-out, represented by the NeXtScale System family, featuring half-wide 1U NeXtScale dx360 M5 compute nodes with shared power and cooling for optimal density, power and cost efficiencies. Options include air or water cooling and the flexibility to add storage and accelerators for large-scale compute and data-intensive applications.
- Converged blade servers, represented by Flex System servers, bring together servers, networking, and enterprise storage – all under integrated management, including virtualization. These systems improve time to value and promote greater utilization and efficiencies.

### **High Speed Adapters and Low Latency Switches**

Low latency adapters include offerings by Mellanox and Solarflare. Mellanox offers both InfiniBand and Ethernet adapters including support for 100 Gbps EDR InfiniBand and 100 Gbps Ethernet. Solarflare's newest dual port 10 GbE low latency adapter is the SFN7122F which combines accurate host clock synchronization and precise hardware packet time stamping with the lowest latency at highest message rates.

These adapters are unique. In addition to being high speed and low latency, they also offer acceleration techniques such as Remote Direct Memory Access (RDMA) or kernel bypass.

For convergence, QLogic and Emulex provide converged network adapters (CNA) to enable the consolidation of Ethernet and Fibre Channel connections.

Lenovo offers several sub-microsecond low latency switches ranging from 24 ports to 64 ports. With these range of options, financial firms can tailor their low latency network environments to their specific environment and needs.

### **Hardware Accelerators**

There are many types of hardware accelerators including GPU's, FPGA's, co-processors, as well as network processors. FPGA is currently the leader in this portfolio of options to accelerate low latency applications.

GPU's are typically used in compute grids, often for risk calculations (many involving Monte Carlo or Black Scholes simulations), options and derivatives pricing, credit value adjustments (CVA), value at risk (VaR) as well as to accelerate backtesting. NeXtScale GPU support includes the NVIDIA Grid and Tesla K series, as well as Intel Xeon Phi coprocessors. ISV support includes Aon, Murex, MATLAB, and NAG and IBM Platform Computing.

*Key OEM  
Business  
Partners  
provide many  
high-  
performance  
adapters and  
accelerators  
for Lenovo  
switches and  
storage*

## Fast and Efficient Storage

Lenovo offers a rich portfolio of memory, flash and solid state offerings to meet specific needs of low latency and HPC applications. This portfolio is enhanced by complementary integrated and network attached storage (NAS) offerings from EMC and IBM.

## Innovation

System x has a long heritage of innovation and just recently opened its first HPC Innovation Center in Stuttgart, Germany<sup>5</sup>. The new HPC center will be used for industry and technology R&D in collaboration with industry leaders, a clear indication of Lenovo's ongoing commitment to innovation.

Other examples of innovative products include the X6 enterprise-class system that is designed specifically for mission-critical Big Data Analytics workloads and the introduction of the fastest Intel Xeon E5-2600 v3 processor sku in the market. This sku is featured in the NeXtScale server, and uses a unique water-cooled design based on extensive research done by System x thermal engineers. System x is engaged with FSS clients evaluating this technology which would avoid the need for a chilling unit (CDU).

## Industry-Standard Benchmarks

Given the expense and complexities of evaluating the various technologies, the Securities Technology Analysis Center (STAC) was created to establish industry standard benchmarks to enable easier comparisons for financial markets clients in the areas that matter to them. STAC has over 300 members, including over 250 technology end users (client firms) and over 50 technology or service providers. It is a major standards organization in the financial services industry.

STAC-M is for market data, STAC-A for analytics, and STAC-E for trade execution. Each of these has sub-areas, e.g., STAC-M3 is focused on "tick databases", where extremely fast queries of past trade data (potentially spanning multiple years), is important. Columnar databases such as kdb+ from Kx, is one popular highly optimized database used by many Wall Street clients, and is used in nearly all STAC-M3 benchmarks.

In January 2015, STAC used the X6 for the baseline STAC-M3 Benchmark suite using the kdb+ database from Kx Systems<sup>6</sup>. Compared to other systems using kdb+, the x3850 X6:

- Chalked up the fastest scores in seven response-time benchmarks,
- Proved 27%-97% faster than the next best performance in the month-, quarter-, and year-high-bid benchmarks, and
- Recorded 86% faster times than the next best performance in the market snapshot benchmark.

System x has been the leading server provider for STAC benchmarks, and that leadership is expected to continue with Lenovo's STAC membership.

<sup>5</sup> <http://news.lenovo.com/news/releases/first-global-hpc-innovation-centre.htm?linkId=13117741>

<sup>6</sup> <http://blog.lenovo.com/en/blog/x6-is-on-the-money-in-the-financial-industry-with-a-new-stac-m3-benchmark/>

## Customer Case Studies

Lenovo System x has well documented successes in the financial services industry.

### Redline Trading Solutions<sup>7</sup>

**Company:** Redline Trading Solutions creates ultra-low latency market data and order execution systems that enable firms to excel in today's equities, options, futures and FX markets.

**Need:** With a huge amount of money at stake, financial markets firms are in a constant arms race to minimize end-to-end latency and seize fleeting market opportunities ahead of the competition.

**Solution:** With Redline solutions on 32-core compact System x3750 M4 servers, firms can reliably achieve ultra-low latency and predictable performance. Clients can deploy their entire trading system on 32 cores in a single 2U System x3750 M4 server.

**Benefit:** Extreme low latency, reliability and fast performance enable firms to detect and act on profitable trading opportunities faster than competitors while reducing cost.

**Customer Testimonial:** *"With the x3750 M4, they can pack more processing into rented co-located rack space, reducing costs and latency."*, Lee Fisher, VP, Technical Marketing.

### Fluent Trade Technologies<sup>8</sup>

**Company:** Fluent Trade Technologies strives to be the world's fastest financial technology provider, offering an end-to-end, ultra-low-latency, turnkey solution for automated trading systems and brokers.

**Need:** Performance is essential to Fluent because it needs to process foreign exchange (FX) market data as fast as possible to boost performance for high-frequency trading, improve operational efficiency and enable companies to rapidly respond to changing market conditions.

**Solution:** Fluent deployed System x3750 M4 servers with Intel Xeon processors to deliver ultra-low latency for its high-frequency trading clients; reduce power, cooling and real-estate costs; and gain the flexibility to adapt to changing market conditions.

**Benefit:** The x3750 M4 solution helped clients achieve single-digit microsecond latency for high-speed trading, which translates to ultra-fast trading for its clients. The dense x3750 M4 helped reduce power, cooling and real-estate costs; and speed deployment of emerging technologies.

**Customer Testimonial:** *"We are always looking for new technologies that can trim off microseconds for our clients. System x3750 M4 can do many things at once to create the efficiencies needed for low latency."*, Moshe Roffe, CTO and head of IT and Connectivity.

<sup>7</sup> <http://blog.lenovo.com/en/blog/the-x3750-m4-is-a-proven-winner-in-the-financial-industry-new-lenovo-case-s/>

<sup>8</sup> <http://blog.lenovo.com/en/blog/x3750-m4-fuels-performance-and-reduces-cost-new-case-study/>

Profitable trading opportunities faster than competitors while reducing cost

Single-digit microsecond latency for high-speed trading



## Options-IT<sup>9</sup>

**Company:** Option-IT is a provider of independent high performance IT and support services for the investment finance sector.

**Need:** Options-IT needed to deliver high performance financial trading applications, where microseconds count. They were looking for a server platform with ultra-low latency, proven reliability with an efficient design

**Solution:** Option-IT rolled out x3750 M4 servers to its data centers located worldwide.

**Customer Testimonial:** “We needed to be able to host our customers’ trading applications on high-performance servers in data centers around the world. We needed a server platform that we knew was going to be reliable and efficient. We found all that and more with the System x3750 M4 platform.”, Ken Barnes, Senior Vice President of Corporate Development.

## Lenovo Offers a Wide High Value HPC Solutions Portfolio

The Lenovo portfolio for the financial services industry is wide and deep. Customers can choose from open systems or bundled enterprise systems. They can scale up with X6, Flex systems, and rack servers or they can scale out with HPC-developed NeXtScale systems. Customers can choose from a comprehensive portfolio of storage, networking, software, and services from Lenovo and Business Partners.

The Intelligent Cluster solution from Lenovo leverages decades of System x HPC experience to reduce the complexity of deploying scale-out systems and clusters. Pre-integrated, pre-tested, and delivered as a fully supported system with a single part number, this solution speeds time to value while matching best-in-industry components with optimized solution design.

### Intelligent Cluster



Figure 3: Lenovo Provides a Complete HPC Ecosystem

With Intelligent Cluster, clients can focus their efforts on maximizing business value, instead of consuming valuable resources to design, optimize, install and support the infrastructure required to meet business demands.

<sup>9</sup> <http://blog.lenovo.com/en/blog/x3750-m4-adds-speed-and-resiliency-to-financial-cloud-services-new-case-stu>

Reliable and efficient hosting of low latency trading environments for global clients

Lenovo's HPC solutions portfolio is comprehensive - open or bundled systems; scale up and scale out servers; and storage, networking, software, and services

## Conclusions

To get a competitive edge in risk analytics and trading, financial firms have been investing in technology for years. Low latency trading and real time analytics are critical to sustain this competitive edge especially in the midst of several intertwined technology trends that's fundamentally changing the rules of the game. Clients building scalable and secure HPC infrastructure must leverage the best technology that is optimized for low latency. These clients should seriously consider Lenovo HPC solutions that:

- Are being used worldwide at major stock exchanges, global investment banks, trading firms and hedge funds
- Include a rich portfolio of servers, storage, software, services and support optimized for low latency trading
- Harness a robust ecosystem of partner solutions that are pre-integrated, tested and fully supported
- Reduce power, cooling and facilities costs while providing a highly secure and reliable IT environment
- Ranks consistently higher on industry-standard STAC benchmarks
- Deliver single-digit latency for high-speed trading, providing a competitive edge.

Most importantly, financial services clients have access to Lenovo experts to work with them to design, tune, test and deploy low latency solutions.

For more information on Lenovo, see [www.lenovo.com](http://www.lenovo.com).

*Cabot Partners is a collaborative consultancy and an independent IT analyst firm. We specialize in advising technology companies and their clients on how to build and grow a customer base, how to achieve desired revenue and profitability results, and how to make effective use of emerging technologies including HPC, Cloud Computing, and Analytics. To find out more, please go to [www.cabotpartners.com](http://www.cabotpartners.com).*

*Clients should seriously consider Lenovo HPC solutions that are optimized for low latency trading*