



NYSE Technologies Expands Access to High-Speed Market Data Delivery Over 10 Gigabit Ethernet using Mellanox Solution

Standards-based infrastructure opens up new level of access to low-latency trading

Financial Services Industry Challenges

The volume and pace of market data have grown considerably in recent years, caused by the surge in electronic and automated trading – more computers are placing more trades, taxing networks and challenging their ability to provide timely quote and trade data.

“Electronic trading and its resultant market data have been increasing; perhaps doubling every 12-18 months in recent years,” said Feargal O’Sullivan, Head of Trading Solutions, Americas, for NYSE Technologies, the commercial technology division of NYSE Euronext (NYX). “This growth arguably outpaces Moore’s law, which says that the number of transistors on a CPU chip doubles every two years.”

In addition, rules have been put in place by the SEC to ensure the efficient and fair availability of price information across securities markets. Regulation NMS (or “Reg NMS”) mandates that all stocks be made available on any exchange, and that quotes are simultaneously shared on all exchanges.

These trends place ever higher demands on market data distribution platforms, the technology that transports market data.

NYSE Technologies’ Middleware Solution

To address these trends, NYSE Technologies was seeking a way to expand access to the performance boosting capabilities of Data Fabric™, its market leading platform for low-latency, high-throughput application messaging. Data Fabric bypasses the biggest bottleneck in legacy publish/subscribe middleware platforms – data input/output (I/O) through the Operating System (OS) and IP stack – by using Remote Direct Memory Access, fundamentally changing the common approach to middleware.

Over two years ago, driven by the need to stay ahead of burgeoning growth of market data and the need to provide timely information to the trading community, NYSE Technologies launched Data Fabric running with Mellanox ConnectX® cards on InfiniBand networks. This year, the company also turned to Mellanox for its ConnectX®-2 EN with RoCE adapter cards. ConnectX-2 EN with RoCE brings the same I/O boosting performance and latency jitter reduction – spurious networking activity that can disrupt and confuse trading applications – to a much wider audience of 10 Gigabit Ethernet users.

Compared to using IP-based middleware, bypassing the operating system and IP stacks with remote direct memory access (RDMA) over Converged Ethernet (RoCE) saves up to 25% or more on CPU overhead, thus reducing latency spikes from an average of 2-3 ms to 50-60 microseconds.

The Feed Handler Linchpin

Market data distribution platforms need to move huge volumes of information extremely quickly. “Market data is the starting point in the trade life-cycle,” Feargal says. “The ability to process quote data is the key piece – systems take data in and distribute it out. When information is first sent out, it is extremely valuable, but its worth quickly diminishes with time.”

The critical path in the information transport chain is the feed handler. Feed handlers are middleware that facilitate the delivery of real time market data from the exchanges to the end users. They translate source data, or normalize it, aggregate data sources and make quote and trade data available across exchanges and trading platforms. Feed handlers allow customers to obtain multiple direct feeds with less latency than consolidated feeds.

Boosting Feed Handler Performance by Tackling I/O Bottlenecks

For optimal performance, feed handlers run on hardware platforms consisting of powerful servers with multi-core processors. "It is tempting to just add more servers to boost performance, but this quickly becomes a complex management headache; that's the challenge," says Feargal.

In its efforts to improve the performance of feed handlers on its platform, NYSE Technologies explored factors that contribute to delays and bottlenecks, and determined that I/O operations are a critical factor. "The data traffic is unpredictable, thereby causing significant CPU context-switching which is inefficient and introduces latency," he said, "therefore off-loading I/O work from the processors is key to reducing latency."

Any given server can receive up to 1 million messages per second. NYSE Technologies servers use multiple cores – feeds are split between the cores, so any single core needs to handle over 100K updates/second. From the input side, raw data must be normalized and published, and sent on the wire without delay.

The Mellanox Connectivity Solutions

NYSE Technologies boosted I/O performance by taking advantage of the remote direct memory access (RDMA) capabilities of InfiniBand connectivity solutions from Mellanox before adding support for ConnectX-2 EN with RDMA over Ethernet (RoCE).

Mellanox's ConnectX-2 EN implements the IBTA's RoCE (RDMA over Converged Ethernet) standard and utilizes the popular OpenFabrics Enterprise Distribution (OFED) and application programming interfaces (APIs). ConnectX-2 EN with RoCE supports both 10 and 40 Gigabit Ethernet, making it the only Ethernet NIC future-proofed for 40 Gigabit Ethernet fabrics of the near future.

ConnectX-2 EN with RoCE adapter delivers up to 1/10th the latency, saves power consumption and delivers improved economics compared to other standards-based 10 Gigabit Ethernet adapter solutions in the market today. Networking solutions from Mellanox complemented middleware solutions from NYSE Technologies to deliver lower latency even at higher volumes and sustain the performance through market data traffic spikes.

The Mellanox connectivity option allowed them to bypass the IP stack and the OS on the systems that process feed handlers. It uses RDMA to free up huge amount of CPU time and reduces context switching, allowing the feed handlers to operate at optimal efficiency without OS intervention.

Compared to using an IP-based middleware, bypassing the operating system and IP stack with RoCE save up to 25% or more on CPU overhead, and reduces maximum latency spikes from an average of 2-3 ms to 50-60 microseconds.

"The performance in comparison with TCP is much better," said Feargal. "This opens up a whole new level of low-latency trading with standards-based infrastructure."



350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085
Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com

© Copyright 2010. Mellanox Technologies. All rights reserved. Preliminary information.
Subject to change without notice.