



Bandwidth's silver bullet

Date: 1 September 2007

The expectations of regulators and investors of best execution are forcing broker/dealers on both sides of the Atlantic to evaluate their connectivity and supporting infrastructure. *John Panzica* argues that ethernet technology has the potential to address the connectivity, latency, and scalability challenges faced by firms

Let's face it – the trading world has grown in complexity and it is not very generous. Today's markets use statistical information along with complicated algorithmic trading and execution strategies. Trading is powered by high-capacity computer systems and the first firm to the market makes all the profit, leaving little behind for the next guy.

Complicating things further are current industry regulations such as the US Securities and Exchange Commission's National Market System (NMS) regulations, which are being rolled out in phases between June and October this year, and the Markets in Financial Instruments Directive (Mifid) within the European Union, which will begin to be rolled out in November.

According to a July 2007 study from TowerGroup, these regulations could potentially result in a 900% increase in the volumes of market data published by 2012. This prediction captures the need to further automate and increase the speed at which data is being processed. Regulations like The Order Protection Rule, The Access Rule, and The Data Market Rules, three tenets upon which Reg NMS is based, are supposed to even out the trading field, although it seems these regulations have sent firms scrambling.

The Order Protection Rule (Rule 611) in particular is the most critical for firms to implement and is deserving of further discussion. In concept, Rule 611 is redefining best execution for all investors by providing equal access to prices and is mandating that trades be executed at the best price. The trade must be executed at the best price regardless of whether the price is displayed, not published, or resides on another exchange.

This requirement has forced firms to look at private linkages, or connectivity. In a game where speed determines winners and losers, firms must consider the need for connectivity that provides bandwidth scalability, low latency, and low total cost of ownership. Technologies such as ethernet have become more and more accepted as the financial industry plays this speed game. Today, exchanges, liquidity pools, and buy- and sell-side firms are all moving to establish ethernet connectivity because of its robust feature set.

The need for bandwidth scalability

To know which exchange is offering the best price for best execution compliance, financial institutions and exchanges need to receive market data so they can analyse and route orders accordingly. With the significant growth in the volume of electronic trading and the bandwidth of market-data feeds expanding at unprecedented rates, trading is challenging the capabilities of legacy infrastructure. In the traditional telecoms world, firms have been relegated to T1 (or 1.5 megabits per second [Mbps]), DS3 (45Mbps) and OC3 (155Mbps) connectivity, which all limit the size of bandwidth increments, have long installation intervals, and at higher bandwidths require expensive router cards that cost anywhere from \$5,000 to \$100,000.

Consider the following example: A typical firm trading equities may receive the Nasdaq Totalview ITCH 3.0 market data feed at 27Mbps. However, because this feed is 'un-throttled', the feed tends to grow larger as trading volumes increase. We have seen examples at market open and closings where in a microsecond snapshot this feed spikes to in excess of

45Mbps. With this type of explosive growth, planning can quickly go out the window – firms need a solution to manage their increased data consumption, which is where ethernet technology proves its value. Ethernet lets firms buy bandwidth as they need it. For example, Nasdaq offers ethernet connectivity. Firms can buy a 27Mbps TotalView ITCH 3.0 feed from an ethernet provider to start with, and if they find over time that the feed is growing or their trading applications are seeing microsecond data bursts, they can increase the bandwidth with one phone call – and receive the bandwidth in minutes. This kind of bandwidth scalability gives firms the flexibility they need to keep up with massive growth.

The value of low latency

If we are playing a game of speed, it is no wonder that firms are dissecting the components of speed to get the best execution times. Simply stated, there are two components to speed: bandwidth availability and latency. Latency is the amount of time it takes a data packet to travel from source to destination.

RedSky Financial, the Chicago-based broker-dealer acquired by ITG in June this year, used a DS3 connection from its Chicago data centre to the Nasdaq data centre in New Jersey. It then switched to an ethernet connection and measured its latency performance to Nasdaq over a two-month period. The broker-dealer reported a 35% decrease in latency, which yielded a 30% increase in trade-order fill ratio to Nasdaq. For RedSky, increased speed equalled millions of dollars.

Now consider latency in a Reg NMS best-execution model. Many brokerages are taking matters into their own hands – technically they can be reg NMS-compliant by connecting to the main markets and allowing the exchanges to do the routing. Plus, the brokerages save money implementing this solution.

Buy side implications

Here is another scenario: a hedge fund customer sends an order that is 'traded through' because of a better price. But the order is cancelled because of latency between trading venues. Not only would your firm have to rely on trade through 'private linkage', but to add insult to injury, if that trade had been executed, your firm would have been hit with additional trade-through fees executing on a different venue. The need for speed is great not only for sell-side execution but also for buy-side clients.

Firms have recognised this, and many are implementing gigabit ethernet ports and reg NMS connectivity bundles to protected venues, along with a range of alternate trading systems. This may seem like a large task: more market-data feeds to process, more connectivity to manage, and far costlier implementations and maintenance. However, using a strategic ethernet provider would certainly be less costly, and taking RedSky as an example, the profit gains can far outweigh any additional operational expense.

Total cost of ownership

Implementing an ethernet solution has a hidden silver lining – total cost of ownership. Because ethernet technology and equipment has been around for years – most corporate networks run Ethernet and have ethernet switches on their premises – costs are reasonable. A standard 48-port ethernet switch retails for \$5,000, or \$104 per port. Hardware for a complete Reg NMS-compliant network with redundancy can be built for a total of \$10,000, using two switches. Network costs would be separate but each switch can support gigabit handoffs and variable-speed port rates.

Whether firms need a 1Mbps order-flow connection or a 300Mbps Options Pricing Reporting Authority (OPRA) market-data feed, their switches – and therefore their costs – will be the same. Compare this to a router solution from any manufacturer – line cards from most manufacturers cost about \$800, and the OPRA market-data feed requires an optical carrier 12 card, which retails for about \$100,000.

Winners and losers are lining up every day to compete in the speed game. With the convergence of technology, regulations and globalisation, the landscape of the financial industry continues to change and get more complex.

John Panzica is general manager, Financial Services Sector, Yipes Enterprise Services, a San Francisco-based provider of Ethernet and application delivery services.

Source: