



## Exchange connectivity

By Dave Rhodes, Director Business Development, TransactTools Inc.



**Electronic trading has grown steadily over the last several years, but rarely have we seen as much evidence of that than in the last two months. The merger of the NYSE and Archipelago clearly shows the commitment of even the largest and oldest US exchange to electronic trading. The merger of NASDAQ and INET also highlights the importance of boosting capacity and performance, even for an exchange that has been electronic since its inception.**

Electronic trading volumes have also soared in the last few months. In Chicago, all three derivatives exchanges announced record volumes. April was CBOE's busiest month in history, setting 18 different trading records. CME topped 100 million contracts in April, and electronic trading accounted for 70% of volume for the month, up 48% from a year ago. CBOT set new monthly records for electronic trading of several instruments including futures and options on mini-sized Dow and 5-year US Treasury notes. The Archipelago Exchange, which has always been fully

electronic, has seen its message volumes soar in recent months as it has expanded order types, and as algorithmic and "black box" trading in equities has taken off.

While these records are exciting, they are challenging at the same time. The more contracts or shares traded, the more customers are attracted – liquidity attracts liquidity. The more diverse the customer base becomes, the more diverse the order flow patterns become. One customer may only need a few electronic messages to trade a certain

number of contracts or shares, while another uses an algorithm that results in hundreds, even thousands, more messages to affect the same traded volume. The increased message volumes tax the trading systems as well as market data feeds. Yet customers demand that not only do the systems handle these loads today with low latency and a high degree of consistency, but that performance continues to improve even as they increase message volumes. In derivatives, the pressure is on to drive down order acknowledgement times, a common measure of order flow performance, well into the double-digit millisecond range, while in equities, the battle is already being fought in single-digits. Exchanges can find themselves running just to keep up.

**Some algorithms can end up with thousands of orders per second being generated, followed by thousands of cancels in less than a second, and only a handful of trades being done**

### Order flow challenges

The variety of trading strategies, especially those that result in the highest message volumes, presents a big challenge to many exchanges. The best case for an electronic exchange is a small message volume resulting in a high number of shares traded. Although significant blocks of shares will always be traded electronically, the explosion of algorithmic and “black box” trading has a very different profile. Some algorithms can end up with thousands of orders per second being generated, followed by thousands of cancels in less than a second, and only a handful of trades being done. At these trade-to-message ratios, an electronic exchange must be capable of handling massive message volumes.

Understanding the order flow is key to managing it. Each trading engine and each protocol engine (commonly FIX, but can be other protocols) must provide visibility into the message flow to and from customers.

- **Rates per message type.** How many orders, cancels, cancel/replaces or order status requests are coming in, and how many types of execution reports – acks,

fills, outs, status responses – are going out. This helps to better understand the trading strategies used by the customers on each engine, and can highlight the more challenging order flow – high rates of cancels and low rates of fills.

- **Overall message rates.** How many messages per second, average and peak, are coming in to and going out of both the protocol engines and trading engines. This information is critical for capacity planning, as well as correlating with overall system performance, both historically and in real-time.
- **Bandwidth.** Although order flow doesn’t present the same order of magnitude challenge that market data does, a few hundred FIX messages per second can be enough to fill up a T1 and cause customers to experience latency. Not all network providers can offer continuous real-time views into bandwidth used between the exchange and its customers, and sampling may not be sufficient to highlight spikes. However, message rate monitoring can provide a good proxy for bandwidth monitoring, using the protocol message size (e.g. Body Length for FIX, fixed length for some other protocols) as an estimate.

**In today’s electronic trading environment, order acknowledgement time is the most common ruler used by customers**

With a better understanding of order flow, correlations can be made to performance. In today’s electronic trading environment, order acknowledgement time is the most common ruler used by customers. Therefore, it is also critical to monitor these statistics at every level – trading engine, protocol engine, and even individual customer connection.

With this information in hand, exchanges can make better decisions about resource allocation, capacity planning, and even software enhancements. For example, if certain instruments tend to attract higher rates of message flow, they may be allocated to a separate trading engine if order ack times begin to trend up. If certain customer

connections exhibit higher rates of cancels than others and order ack times on that protocol engine begin to trend up, some connections may be shifted to another protocol engine. If outgoing rates of execution reports on connections exceeds a few hundred per second, the customer may want to consider adding bandwidth.

**In situations where there is potentially a long delay between the time a request to cancel or replace an order is received and the time the cancel or replace is actually accomplished, the two-step reply of ack-then-out is valuable**

For most exchanges, reducing message volume while keeping the same trading volume is an important goal. This reduces the strain on trading engines, protocol engines, the network, and even the customers' systems. There are several options for accomplishing this:

- **Suppress cancel and cancel/replace acknowledgements.** In situations where there is potentially a long delay between the time a request to cancel or replace an order is received and the time the cancel or replace is actually accomplished, the two-step reply of ack-then-out is valuable. However, for trading strategies that result in very high rates of cancels, avoiding the extra ack can be a significant saving.
- **Provide a bulk cancel message.** FIX 4.3 introduced the OrderMassCancelRequest type, enabling customers to cancel many orders at a time (e.g. all, by symbol, by side). This can significantly reduce message volume, since only one incoming cancel is sent, and only one outgoing cancel ack is sent, followed by all of the individual cancel out messages.
- **New order types.** A new order type can achieve the same behavior that would otherwise require a larger message volume. For example, a pegged order can track with the NBBO automatically. Without a native

pegged order in the trading engine, a customer who wants to use a pegged strategy ends up sending cancel/replaces, or even cancels followed by new orders, each time the NBBO moves.

### **Market data challenges**

The order flow challenges above directly impact market data. Each incoming order can generate a quote, and each cancel can generate an update or delete to a quote. These increased message rates can have impacts not only on top-of-book quote feeds, but can have devastating impacts on quote feeds that include depth-of-book. The case is exacerbated even more in equity options, due in part to the number of instruments quoted on a single underlying symbol. A new, updated or deleted quote on an equity can result in multiple adjustments to quotes on its options.

In late 2004, the Archipelago Exchange book data had peaks of over 6,000 messages per second, while OPRA (Options Price Reporting Authority) had one-minute peaks of close to 40,000 per second. These rates are expected to exceed 100,000 in the near future.

**Gone are the days when a fractional or even full T-1 is sufficient for receiving an exchange's full market data**

Clearly bandwidth is a challenge for both exchanges and their customers, since messages per second translates directly to bandwidth. Gone are the days when a fractional or even full T-1 is sufficient for receiving an exchange's full market data. Archipelago recommends at least 3Mbps (equivalent to two T1s), and OPRA volumes are already approaching the limits of a DS-3 at 45Mbps.

Market data protocols are also challenges to exchanges and their customers. For order flow, FIX is generally the standard, but using FIX and a TCP/IP transport layer for market data is not realistic at these volumes. It's too bandwidth intensive, and too taxing on the servers sending out the market data to maintain individual connections and pump out that much data to each one. Therefore, each exchange has come up with its own protocols. Some are ASCII, some are binary, and some even offer compression.

Some use TCP/IP as a transport, some use IP multicast, and some use middleware like TIBCO Rendezvous.

**Bringing together exchanges, sell-sides, buy-sides, networks and vendors, the group's charter is to define standard services for dissemination of FIX market data information**

In early 2005, FPL kicked off a working group focused specifically on market data optimization. Bringing together exchanges, sell-sides, buy-sides, networks and vendors, the group's charter is to define standard services for dissemination of FIX market data information. This won't be FIX as we know it – tag-value messages over a point-to-point TCP/IP connection. Instead, the group is exploring a concept called “implicit tagging”, which reduces message size, but does not give up flexibility as can happen when moving to fixed-length protocols. The group is also considering compression techniques, multicast transports as well as point-to-point, and overall market data and book management best practices.

**As algorithmic and black-box trading grows in usage, some customers are demanding no network-induced latency, in either order flow or market data**

#### **Other trends**

As algorithmic and black-box trading grows in usage, some customers are demanding no network-induced latency, in either order flow or market data. For this reason, some exchanges are beginning to offer co-location, allowing customers to place their systems in the same data center as the protocol engines. While the customer can still remotely access their systems, the applications receiving the market data and driving the order flow are now one high-speed local area network connection away from the exchange, rather than several miles and a T1 away.

As order flow message volumes continue to increase, there may be limitations to the current FIX session layer, especially related to bandwidth. INET already offers proprietary, fixed-length ASCII protocols called OUCH and RASH for order flow. These messages consume much less bandwidth than their FIX message equivalents. Archipelago is also exploring a proprietary, fixed-length protocol, offering both ASCII and binary versions, designed to reduce bandwidth. However, if the recommendations of the FPL Market Data Optimization Working Group get traction in the industry, it's quite possible that some of the same techniques could be applied to order flow.

The exchanges that are already fully electronic, along with those who are seeing their percentage of overall trading being done electronically skyrocket, are already feeling the pain and must take measures to address these concerns. The maxim “liquidity begets liquidity” still holds true, but the real challenge is with the corollary to this – “message volume begets message volume.” **FIX**

#### **Any thoughts on this or other articles?**

Please send any comments, referring to this article as Vol 1 Issue 6 AM1, direct to Edward at [edward@fixglobal.com](mailto:edward@fixglobal.com)