

Co-Location Catches On

By Bennett Voyles

The pit is back. Just a few years since the concept of a commodity exchange as a tangible “place” had begun to seem hopelessly old-fashioned, many traders now want to be at the heart of the action once more.

At Eurex, customers that until recently were scattered all over the globe are moving closer to the exchange, “forming a physical community like a pit again,” says Matthias Kluber, head of networks and infrastructure operations at Deutsche Börse Systems, which builds and operates the Eurex trading and clearing systems.

But no one is practicing the old hand signals. The desire to be closer to the exchanges is driven by the need for speed. Even when traveling at the speed of light, an order message transmitted from hundreds of miles away from the exchange will take slightly longer to arrive than a message transmitted from only a few hundred yards away. The pressure to be the first in the queue is so great that more and more firms are installing their servers as close as possible to the exchanges where they trade.

The demand is especially strong from the “black box” trading firms that engage in very high frequency trading. Using one of the new “co-location” services offered by the exchanges allows automated trading systems to shave a few milliseconds off the time it takes for their order messages to reach the exchange. Every millisecond counts in a world where traders and market makers are trying to squeeze as many as 10,000 order/answer pairs into a second.

Instead of using a third-party network or telecom carrier, which typically introduces latency as order messages are transferred from one node to another, the firms using these services can connect their automated trading systems directly to the exchange’s electronic system via a dedicated high capacity line. In some cases, the servers are located in the same building as the

exchange, bringing the physical distance that these order messages must travel to the absolute minimum.

Martin Koopman, head of Orc Software in North America, says one of his customers is so focused on achieving the lowest possible latency in its trading at the Chicago Mercantile Exchange that the firm moved its servers into the CME’s co-location facility, even though the servers were already in Chicago. “For a high-frequency trader, latency is one of the most critical factors in determining the profitability of his trading strategy. You have to have the fastest possible connection to the exchange, because if you don’t, somebody else is going to get the trade before you every time.”

The catch is that the advantage of being closer to the exchange gets smaller and smaller as more and more firms put their

servers into these facilities. “The joke is that just as soon as everybody does it, you’re back to competing on an even playing field,” says Eric Goldberg, chief executive officer of Portware, a software company specializing in algorithmic trading solutions.

CME

The Chicago Mercantile Exchange began sponsoring a co-location service in October. The service, CME Local Network or LNet, went live in December and gives trading firms direct access to the Globex network through 40mb redundant connections to two co-location facilities in Chicago, one operated by Equinix, the other by Digital Realty Trust. The CME charges \$6,000 per connection, per month, not counting co-location facility charges,

and requires LNet customers to maintain a back-up site with a separate connection to the CME.

CME spokesman Allan Schoenberg says using LNet brings the network latency down to less than two milliseconds. That is dramatically faster than the average round trip time across the entire Globex network. Total round-trip time, including the time it takes to process an order message and send it through the Globex gateway, can be as low as 31 milliseconds for futures and 11 milliseconds for options.

"No question, it's a lot faster than WAN (wide area network) routing into the Globex platform," says Diane Saucier, director of product marketing at Trading Technologies. "We have customers for whom speed is the absolute priority. They want the fastest connection to the exchange, and that's the way to get it."

Like several other independent software vendors, TT offers a hosting service called TTNET that provides fully managed connectivity to numerous exchanges around the world. TTNET's servers are co-located with the CME, Saucier says, and TT is looking to do the same with other exchanges including Eurex.

The CME will not release detailed information about the growth of the LNet service or the number of users, but firms familiar with the service say the facilities are already crowded with high-speed trading firms. One industry source involved in algorithmic trading said some technology companies are now offering to rent rack space in the facilities as part of the services they offer to the high-speed trading community.

Eurex

Eurex began looking into the co-location idea in early 2006 and launched what it calls a "proximity service" in August of that year via a facility in Frankfurt operated by IXEurope. In February 2007, Eurex signed an agreement with a second provider, Colt, to increase capacity and ensure a backup to the IXEurope facility.

By the end of 2006, 13 customers had signed up for co-location. Six months later, the number had risen to 36, according to Heiner Seidel, a spokesman for the exchange. Round-trip time from within one of these facilities averages less than 10 milliseconds, versus 29 milliseconds for the

fastest connection from London, and 128 milliseconds from Chicago. For market makers and proprietary traders outside Germany, the difference can be critical to the success of their trading strategies.

One firm using this service is Nico Trading, a proprietary trading firm in Chicago that specializes in trading options on fixed income futures. In February, the firm said that moving its Eurex trading system from Chicago to Frankfurt resulted in a net reduction of latency from approximately 130 milliseconds to less than 30 milliseconds, measured from the firm's own servers through the gateway to the Eurex system, into the matching engine, and back out again to the firm's servers.

"Our automated trading systems benefit from, among other things, low message latency so this upgrade is making the difference between a successful and an unsuccessful trade, especially in high volume markets," Ricardo Naon, the firm's chief information officer at the time, said in a statement issued by Deutsche Börse highlighting the value of the proximity service. The move also allowed Nico to close its London office, the previous base for its Eurex trading, which saved money as well as time.

Euronext.liffe

Euronext.liffe does not offer co-location yet, but it's coming soon. Spokesman James Dunseath says that he expects customers will be able to start bringing in equipment in the fourth quarter of 2007, when the exchange completes two new data centers. The exchange is building centers in Paris and outside London. Half of the space will be dedicated to co-location cabinets with separate power supplies, according to Dunseath. Each data center will act as a backup for the other in offering high-speed connectivity to both Euronext.liffe and Euronext cash markets.

ICE

ICE Futures first began offering co-location service in summer 2006. "We were getting increasingly more requests from different firms that wanted to co-locate and it seemed that a lot of exchanges were already offering co-location as an option," says Edwin Marcial, chief technology officer for IntercontinentalExchange in Atlanta.

Marcial says demand has been largely driven by black-box traders, especially customers who want to avoid the 20-25 milliseconds of latency involved in transmitting order messages to the Atlanta-based platform from an office in New York or Chicago. The network latency is even greater for firms operating in London, the base for many of the traders in the exchange's Brent crude oil futures complex.

Marcial says demand is also coming from trading firms that use third-party algorithmic trading tools. "We are starting to see more interest from proprietary traders who are using independent software vendors," he adds.

Having a server in the ICE co-location facility in Atlanta brings the transmission time down to about one millisecond, Marcial says. That does not include, however, the 27 milliseconds that it takes for ICE's trade engine to process an order message. ICE is currently installing a new trading engine that will cut the processing time down to seven milliseconds, and Marcial expects that to be implemented by the end of August across all of its markets.

Later this year, ICE plans to move its technology operations to Chicago, with the hosting of the trading operations scheduled to move from Atlanta to Chicago in January 2008. This will bring the exchange closer to its customers in the Chicago trading community, but may set off a scramble among firms based in London, Houston and New York to move their automated trading engines to the new co-location center, which will be in the same building as the telecommunication hub the exchange currently has in Chicago.

The exchange is also trying to make the processing time more consistent. Marcial estimates that today, roughly 8% of the order messages that the exchange receives take more than 50 milliseconds to process, versus the average of 27 milliseconds. The new trading engine will not only be faster, it will also eliminate these latency spikes. He predicts that less than 0.1% of the messages will be processed in more than 50 milliseconds once the new trading engine is fully installed. ■

Bennett Voyles is an independent business reporter based in New York.