FT.com
FINANCIAL TIMES

When microseconds really count

By ROSS TIEMAN

March 19, 2008 Wednesday
FT Digital Business

The need for speed is changing the geography of financial trading. If you drive through the Lincoln tunnel from New York, you emerge in the New Jersey suburb of Weehawken. There, six miles from Wall Street, in a nondescript data centre, hum the trading servers of five stock exchanges and electronic communications networks (ECNs).

This single server farm run by outsourcing specialist SAVVIS is home, among others, to the platforms of the American Stock Exchange, the Philadelphia Stock Exchange, and BATS Trading, an ECN that claims to process deals in less than a millisecond, one thousandth of a second.

But it also houses the trading servers of many of the banks and investment houses that trade shares on these exchanges. Co-hosting, or co-location of servers, is becoming commonplace in the equity trading industry, and some exchanges, such as Nasdaq, now regard it as a core part of their business.

The aim is to shave microseconds off the time lag - or latency - in which computers loaded with automated trading programs can react to momentary variations in share prices. Minimal latency enables the machines to slice a big "parent" order, to buy or sell shares, into small "child" orders that can be blasted into the market in quick succession, completing the order before the market moves against the dealing machine.

Kirsti Suutari, global business manager for algorithmic trading at information and data group Reuters, says: "You are starting to have machines competing against machines. The one that gets there first is able to exploit the opportunity. Traders are developing models to look for arbitrage opportunities that may only exist for a moment. That is why speed matters."

As Jeremy Badman, partner in the strategic IT and operations practice focusing on Investment Banks at Oliver Wyman, points out: "If data has to come a long way to the place where there is trading that can take milliseconds that really do matter. Since data travels at the speed of light, the difference in speed between a co-located server and one 200 miles away is about a millisecond."

That is about how much time a number of hedge funds save by putting some of their algorithmic trading programs (algos) in Weehawken.

But the traders don't need to be there. "You are often running several algos," says Mr Badman. "As a trader, you might want to stop one algo and run a different one. The message is going to take three or four milliseconds to get to the server and shut it down. Yet a human might take half a second to react to new data and send an instruction, so that bit is not critical."

The race for speed puts exchanges into constant competition to process more orders, more quickly.

The London Stock Exchange introduced a new electronic trading system, TradElect last June. In October, an upgrade improved TradElect's processing speed by 40 per cent and system capacity by 70 per cent. Now, it can execute trades in 6 milliseconds and process 4,200 orders a second. "Technology increases efficiency, lowers latency and cuts costs - driving trading volume growth," a representative says.

Varghese Thomas, vice-president of financial markets at SAVVIS, which operates the Weehawken facility and 29 others worldwide, says studies show "reducing latency by 1 millisecond can be worth up to Dollars 100m a year to a leading trading house".

BT Global Financial Services offers "proximity hosting", designed to put trading servers closer to multiple exchanges with connections offering up to 1 gigabit a second of bandwidth that cuts latency below a microsecond.

The race for speed, says Mark Akass, BT Global's chief technology officer, is "like turning up the pressure in a hose". Fix one bottleneck - for example by increasing bandwidth, and others appear that have to be tackled.

Says Ms Suutari: "The speed of light is only one constraint. Trading firms are also buying faster processors, investing in programs that run faster, and looking for unique content that the machines can digest more efficiently." Encrpyted news, now available from Reuters and Dow Jones, enables machines to react faster to price-sensitive announcements.

In a study last year, research firm Tabb Group estimated that investment banks are spending more than Dollars 300m a year on low-latency technology.

Mr Badman says: "The need for speed is putting a huge amount of pressure on technology vendors, especially system vendors, because they are in an arms race to keep up with the demands from clients."

Yet ironically, technology is also reducing the need for market players to cluster. Because trading servers can be co-located, traders can be tens, hundreds, or thousands of miles away.

In the US, investment house UBS has a huge trading floor in Stamford, Connecticut, which seems at no disadvantage. In Britain, some fund managers have located their offices close to the M25 highway that orbits London, avoiding a protracted daily commute into the City. It seems lower latency technology can have some lifestyle benefits, too.

Financial Times (London, England)