



**InformationWeek**  
THE BUSINESS VALUE OF TECHNOLOGY

Without fast data, the most advanced trading systems are useless.

By Jim Middlemiss, [InformationWeek](#)

March 01, 2005

URL: <http://www.informationweek.com/news/60404306>

Investment management firms can spend millions of dollars on fancy trading software and complex algorithms, but, if the system managing their market data is slow, then they're wasting their time and money. Data latency, the lag experienced when data is moved from one place to another, robs firms of their ability to leverage those investments and react quickly to market changes.

**ADVANCED  
TRADING**  
Spring 2005

While it might be only a millisecond - one-thousandth of a second (think: a blink of an eye) - data latency can be the difference between getting a fill or being shut out of the market.

It is evident that cloud will save money; but less obvious is what implementing cloud will mean over the next few years.

[Review research that reveals IT pro's perception of the cloud.](#)

Danny Moore, COO of Incline Village, Nev.-based Wombat Financial Software, which makes software that allows for direct exchange connections, estimates that 60 percent to 70 percent of high-performance brokerage firms don't even know what latency is. "If you don't know what it is or where it's coming from, you can't address it," he says.

But there are things firms can do to reduce data latency. Following are five ways to tackle the problem.

**1. Build a direct feed to an exchange.** By connecting directly to an exchange, as opposed to a data consolidator, investment firms can gain between 150 milliseconds to 500 milliseconds in transmission times. Direct-exchange feeds eliminate the data hops and thus reduce latency, Moore explains.

A direct-exchange feed, however, is more complex and more expensive. Firms need to build and maintain a system that cleans, formats and processes data so that it can be distributed throughout the organization.

**2. Colocate servers near an exchange.** This is an option for firms whose main office is located far away from the exchange. Five miles is one thing; 50 miles might be something else. Market data can't travel faster than the speed of light, which is 186 miles a millisecond. So, if servers are 3,000 miles away, a firm may be sacrificing 16 milliseconds or more, notes Moore.

However, Guy Tagliavia, president and CEO of InfoDyne Corp., a Chicago-based firm that makes technology that connects exchanges to firms, disputes this assertion. He says that based on his firm's testing with brokerages, it "doesn't really matter where the server is located."

**3. Make IT physically part of the trading floor.** By placing IT professionals and traders side by side, IT experts can monitor workflow among traders, applications and systems in real time.

For example, firms need to be prudent in the way they manage data, says Mike Alsip, head of product

management at Reuters Market Data Systems in Chicago. "Limit the amount of data elements and value-added processing to what is absolutely needed to support pre-trade analysis and execution."

**4. Re-evaluate your middleware structure.** Internally, ticker plant software handles the incoming data, which flows through a firm's middleware to downstream applications. Depending on that technology, firms could be sacrificing as much as three seconds, according to Wombat Financial's Moore. "There are still a lot of firms out there with internal systems that are doing much more damage to latency than the aggregate data feed," he says.

"It's not just about latency; it's about latency at high through-put rates," adds InfoDyne's Tagliavia. For example, if you can send 1,000 messages in a millisecond but it takes one second to send 40,000 messages, that's a problem. "You need to be able to support high through-put and still maintain flat, low-latency," Tagliavia points out.

**5. Get applications closer to the ticker plant.** Locating critical applications closer to the ticker plant also reduces data latency. "It's all about eliminating hops," says Tagliavia. "Processing hops is expensive in terms of latency."

Reuters' Alsip adds that firms need to look at "flat client-site distribution architecture. You must be able to distribute from the direct feed directly to applications, minimizing the local area network hops in between." Such architecture, he says, "needs to be able to scale to handle applications subscribing to the low-latency direct data feed."

**VISIT GO PARALLEL**  
PROGRAMMING FOR PERFORMANCE

EXPERT TIPS AND TRICKS: C++,  
OPEN MP, THREADING AND MORE



Copyright © 2011 [United Business Media LLC](#). All rights reserved.