

Built for Speed

e-Forex asks **Mike Mistretta**, Senior Director, Infrastructure at FXall to tell us about their Latency Monitoring System.



Mike Mistretta

Speed of execution is key to maximizing foreign exchange performance and productivity. For banks streaming prices to a multibank portal, speed is central to winning and retaining client business, while for model-driven traders, it means that trades can be executed with minimal slippage.

Research from Greenwich Associates demonstrates that speed of execution is regarded as the single most important benefit of online trading. 68% of top-tier financial institutions – including banks, asset managers, and hedge funds – cited it as one of the most important advantages realized by trading foreign exchange online. Active traders, such as hedge funds and CTAs, place particular emphasis on lightning-fast pricing and execution – our clients in this sector often tell us that speed was a primary motivation for moving volumes online to FXall.

Accelerating the trading process

Speed of execution depends on two things – maximizing speed and minimizing latency. To increase the speed of trading, we have delivered:

Deep markets for one-click execution – FXall offers deep liquidity from more providers than any other portal.

Our diverse customer base delivers non-correlated flow, giving banks an incentive to put up more liquidity. This depth of liquidity means orders can get filled in one execution – far more efficient than managing risk over five or six transactions. What's more, our one-click execution capabilities ensure slippage is kept to a minimum.

Flexible connectivity options – Clients can interface to FXall seamlessly through FIX, API (including Java* and Microsoft COM*) or our STP solution QuickConnect, enhancing efficiency and eliminating the need for manual re-keying.

Powerful infrastructure – FXall is built for speed. Last year, we embarked on an ambitious project to design and deploy a new generation of software for FX trading.

Our new infrastructure is based around a grid technique of processing engines, which ensure the shortest code-path for maximum speed. It is highly scalable for increased throughput, meaning that as volumes rise it can be easily extended by adding individual software engines and hardware servers.

Minimizing latency

In the quest to deliver speed of execution, investing to make FXall the fastest, most effective trading platform is only the first step. Equally important is making sure it stays that way.

By monitoring latency both internally and externally, and by working with banks and clients to troubleshoot issues as soon as they occur, our company-wide Latency Monitoring System (LMS) ensures that we consistently achieve the sub-second transaction time targets we set when we built FXall. 99.99% of trades are executed in less than one second.

Transaction latency can be broken down into three components:

- Network latency – Latency in data being transmitted from a client or bank server to FXall over a WAN (wide area network). This can involve a number of issues, including latency in the user's internal infrastructure or the way that information is being routed over their ISP. The location of the end user also plays a part – if a trader is based in London, but his firm's Internet gateway is in New York, any latency in transmitting the trade information internally could add to the time it takes to trade.
- Application latency – Internal processing of the transaction at FXall's application layer
- Infrastructure latency – Latency in FXall's network and server

LMS continually tracks all these components. FXall's application and infrastructure performance are monitored and reviewed continually. Daily, weekly and monthly reports summarizing latency at component level are reviewed to ensure optimal performance is being achieved.

LMS also monitors WAN latency on the bank and client side, and automatically alerts the customer support team if any user appears to be affected by higher-than-average latency. Customer support staff are also provided with tools to review WAN and application latency on a real-time basis. As soon as a latency bottleneck is identified, our network staff will work closely with network staff at the user organization to address the issue.

As one of the fastest-growing online trading platforms, the daily volume of transactions FXall handles is rapidly increasing. Capacity planning is therefore a central part of our latency reduction strategy. We regularly test our servers and applications to ensure that they are capable of handling not just the volumes we are seeing today, but the volumes we expect to see next week, next month or next year.

For FXall, delivering speed of execution is an ongoing process. We regularly set ourselves – and better – ever-lower latency targets to ensure that we are among the very best in the business. We are continually reviewing not only the performance of our systems, but also those of our users, and are engaged in a constant dialogue with banks and clients to ensure that they are operating at optimal levels. We have also gone beyond execution to enhance speed and productivity across the entire trade lifecycle. Only by doing this can we provide our users with the speed and efficiency they demand.

*Java is the registered trademark of Sun Microsystems; COM is a trademark of Microsoft Corp.