Managing latency

The rapidly increasing volume of market data and trade flow is creating new challenges in managing latency. At the same time latency is becoming one of the key factors in determining the best strategy for trade execution. So, how do we extract accurate and timely latency information from this sea of data? **John Smith**, **Mihaela Risca**, and **Helen Ojha** at Cisco offer their view*.

Challenges:

Server Load

As the volume of market data and trade order messages grow, firms are upgrading their networks from Gigabit Ethernet (1GE) to 10 Gigabit Ethernet (10GE). With the increased network's capacity, latency-monitoring tools are being flooded with huge amounts of data. When the monitoring is done on servers, the server CPU cycles required to process this incoming data compete with the CPU cycles required by the business application that the server should be running. On the other hand, if probes are used for monitoring at multiple points in the network, they will capture multiple copies of the message traffic, but processing it down to the interesting information becomes unmanageable in real-time.

Granularity

Another challenge in accurately identifying the real causes of latency is being able to measure latency with sufficient granularity. Sampling times matter. Traditional network monitoring tools operate with minutes or seconds of granularity. Next-generation trading platforms, especially those supporting algorithmic trading, require latencies of less than 5 milliseconds and extremely low levels of packet loss. Let us consider an example: on a 1-Gigabit LAN, a 100 ms microburst can cause 10,000 transactions to be lost or excessively delayed. In these environments, latency must be measured at a much more granular level. Additionally, measurement of latency must break down the end-to-end transaction so that each step can be properly measured.

"Latency must be looked at end to end. The quest to reduce latency could mean shaving off milliseconds by upgrading network infrastructure, or deploying trading systems closer to the execution values."

JOHN SMITH

"Next-generation trading platforms, especially those supporting algorithmic trading, require latencies of less than 5 milliseconds..."

HELEN OJHA

Time Synchronisation
Another challenge in this environment is time synchronisation. Time stamping in network and server devices is typically handled by Network Time Protocol (NTP). NTP is accurate to 1ms to 50ms at best. But latency measurement requires accurate timing commensurate with network and application processing speeds making NTP inadequate for hop-by-hop measurements. In addition, server-based time-stamping with NTP time synchronization is often subject to clock drift greater than 1 to 50ms.

Clock drift can occur in three areas:

- Most server or blade hardware use inexpensive crystal clocks and are subject to physical environmental changes on the crystal – heat for example. This causes clocks to drift.
- Internal clocks are updated via NTP time stamps – this update is typically a low priority compared to other traffic managed by the server operating system. Under heavy CPU loading, the task to process this is placed on a lower priority – causing clock drift.
- Jitter in the network – causes variability in the delivery of time clock information.

Solution:

Where to Measure Latency
Latency can be measured end-to-end, or hop-by-hop; in the network or at the server; at the packet level or at the transaction level; using active probes, or in passive mode. A key issue in any latency monitoring implementation is to provide relevant and accurate information in a timely manner and also identifying exact location of delays i.e. whether in the trading application, core network or at the interface to counterparties. The main components of a truly effective latency monitoring solution comprise multiple points of data collection and the central application that performs the data reduction, correlation and reporting.

Network as the Platform
The network's day job is to forward packets and information. The same hardware characteristics within network devices that achieve this can be deployed towards a ubiquitous data collection infrastructure with reduced footprint and power requirements. In effect, the same type of hardware that pushes packets forward can be used to "reflect" upon their forward movement. The capture, filtering and analysis of this information is carried out on copies of the traffic. This "out of band" technique, ensures that we do not add to the overall latency.

Network engineers have already invented many tools to gain insight upon the flow of packets as they traverse the network. One example is in the analysis of network traffic for security purposes such as intrusion detection or DDOS (Dynamic Denial of Services) attacks prevention. The tools use bandwidth usage, latency and microburst measurements to identify a threat and then, based on the appropriate policy, modify the network behaviour. This is in fact a "feedback loop" or a "reflection" point in the network.

The network is location-aware, in the sense that it can collect data at intermediate points as the

"Network engineers have already invented many tools to gain insight upon the flow of packets as they traverse the network."

MICHAELA RISCA

messages are carried end to end. Server-based collection tends to be at the edge, only, thereby having a more limited view of what is being measured.

The network has hardware enabled filtering. In addition to packet capture and packet inspection, the network is capable of pre-processing and refining at the message level, before sending the information to the central analysis engine. The issues surrounding accurate time-stamping, can be solved by using hardware-assisted time synchronization network modules based on IEEE 1588 Precise Time Protocol.

Suitably implemented, these capabilities make possible continuous and near real-time monitoring of latency characteristics across the trading cycle.

Reflective Architecture

In addition to being an intelligent packet forwarder, the network can also be a "reflector" of intelligent information. This information can be used as feedback not only to the network devices, but also to the business applications themselves and improve the applications decision-making quality. At Cisco, we see these characteristics enabling the development of a Reflective Architecture.

Programmable refining on captured traffic creates new information streams that can be used immediately by any analysis application, for example: micro-burst transaction detectors or end-to-end performance analysers. Examples in the area of financial markets include companies like Corvil and OpTier.

Visibility into end-to-end infrastructure is essential for operations. Taking this to the next level, the network could provide latency insight to trading decision-making. For example, a smart order router searching for best execution of trades could be provided with latency information from the network, pertaining to different execution venues.

Summary:

Latency must be looked at end to end. The quest to reduce latency could mean, shaving off milliseconds by upgrading network infrastructure, or deploying trading systems closer to the execution venues. This is all well and good. If however, the trading applications themselves are not efficiently coded, or the servers upon which they run are not well optimized, then the real gains that could be made with a holistic approach could be lost. The same principles hold true in identifying the causes of latency with an effective capture, analysis and action approach. As the application and network engineers together with enterprise architects work to build a reflective architecture, then the "voice" of the network can be used to dynamically affect the use of the applications and improve the applications decision-making quality. With this co-ordinated action, we can look forward to more effective market data and trading systems that are fast, robust, and more adaptable to the increasingly volatile environment in which they must perform. ■

*The authors are John Smith, Business Development Manager, Financial Markets; Mihaela Risca, Industry Solutions Manager, Financial Markets; Helen Ojha, Business Development Manager, Network Services, Cisco Systems.