

Latency

IN FINANCIAL MARKET NETWORKS

In pure engineering terms, latency is the time between an action and the results of that action. In the processing world, latency can be the time between an input action and the displayed results. In networking, latency is a synonym for delay or the time it takes a datagram to move from one network device to another. In the world of the internet, Latency is measured in round-trip time. There are several factors that can contribute to latency such as processing speed, DSP (Digital Signal Processing) time, routers changing the packet headers, traffic delays, bottlenecks and overall network performance/speed.

In the financial world, business results are measured by the bottom line. Contributing to the bottom line is the number transactions processed per minute (TPM). In computing we measure milliseconds and nanoseconds. However, if you take several thousand transactions per minute, in reality, transactions are processed in nanoseconds. Each transaction, whether it be a bank card authorization, a stock trade or other financial transaction must process an error free and in many cases in real-time with virtually zero latency. Global investment banks, stock exchanges, traders and hedge funds stand to gain the most from real-time market data and processing of that information.

TRANSACTIONS PER SECOND (TPS)

In order to benchmark transactions, the Transaction Processing Performance Council (www.tpc.org) was formed and defined the original TPC-A benchmark formula. Today, for benchmarking OLTP (Online Transaction Processing) by simulating a brokerage firm, the council has developed TPC-E. This model simulates a brokerage house where customers generate transactions related to trades, account inquiries and market research. Each transaction is then acted upon by the simulated brokerage firm that issues orders on behalf of those customers based on real-time market data and updates relevant account information. This draft, now in version 1.0, was developed to replace the older TPC-C (currently the only benchmark with published data) to more closely model the securities business. The TPC-E model is based on wholesale operations. Using the TPC-C models, the top tpmC is 4,092,799 with the average amongst the top 10 being 1,965,221 tpmC. Even a little latency here would be detrimental if it were a real financial institution.

LINK Interchange Network, Ltd., located in the UK, processes up to 12,000 shared online transactions over 36,000 ATMs throughout the United Kingdom on behalf of 47 of the UK's largest financial institutions, primarily banks. And SETS (The London Stock Exchange electronic order book) exceeded 500,000 per day in March of 2007, with the total value traded for the month represented 11.2 million trades for a total of 193.5 billion pounds.

Today, a simple PC can process about 8,000 transactions per minute — more than one of the largest US bank's total traffic in 1970, according to Microsoft®. According to Finextra Research (www.finextra.com), HSBC, with over 125 million customers (1/5 of whom are online customers) processes 100 million transactions daily and moves approximately \$1 trillion dollars in funds in that same daily period.

The NYSE and NASDAQ trade 2 billion shares each day with each transaction representing 3-400 shares each. This equates to roughly six million transactions per day, compared to 2 million just 5 years ago. This is due in large part to algorithmic trading which is HIGHLY sensitive to latency.

A representative from Visa International that chaired the International Payments Summit stated that they have laid 9 million miles of cable to support their global network. The networks that processes 6,800 peak transactions per second with a total of 24,960,000 per hour. In the 12 months prior to 6/06, they processed a total of 56.3 billion total transactions with a total volume of \$4.5 trillion.

GROWING CONCERNS OVER LATENCY

In a recent Wall Street and Technology magazine webcast, a poll was taken amongst attendees with the question: "What is the greatest limitation or challenge in your current infrastructure and organization around processing and analyzing real time market data?" Of the respondents, 43.1% stated that latency was their biggest concern.

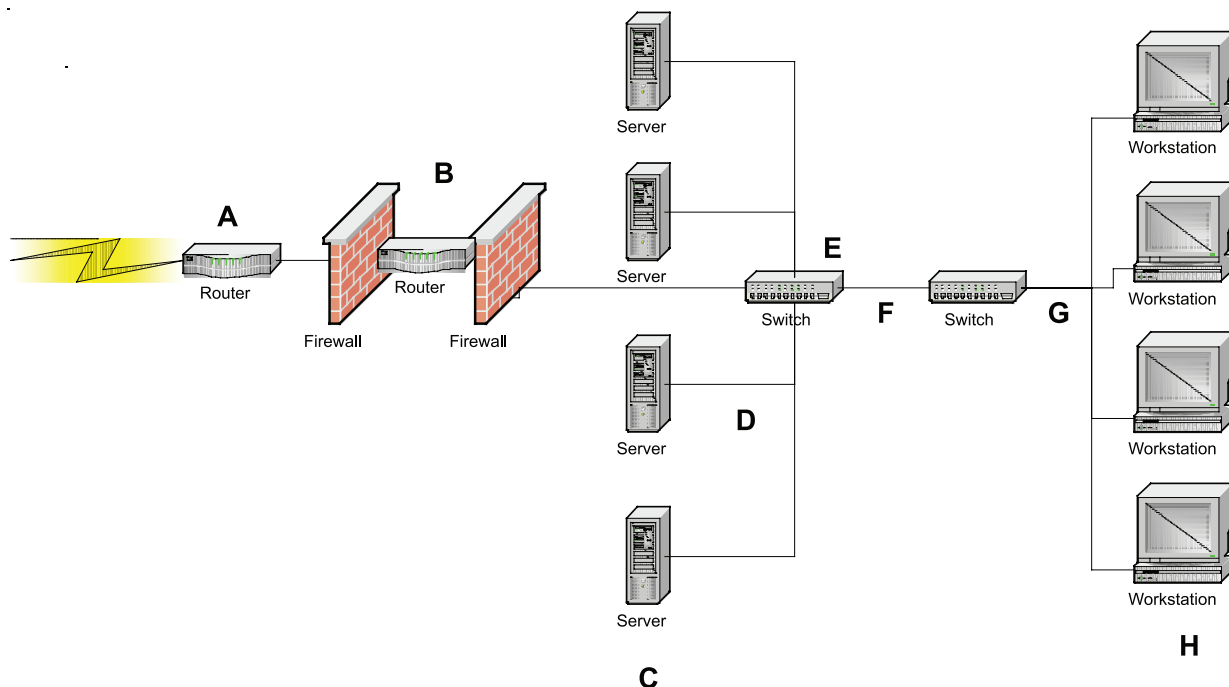
The hurdles to real time (virtually zero latency) data include:

- Traditional store and then process systems
- Processing a combination of live and historical data in tandem
- The inability to accurately predict and rely on bandwidth in the wide area
- Managing imperfections within a stream
- Load distribution/ load balancing and hardware performance
- Decimalization in real time
- Non-deterministic networks (such as Ethernet)
- Home grown and/or poorly performing infrastructures

From the enterprise to the local area network, the issues above will have different solutions. However, since the most controllable latency is within the local area, we will focus on the LAN and Data Center.

HOME GROWN AND POORLY PERFORMING INFRASTRUCTURES

There have been many mergers and acquisitions in the financial sector. Not only is compliance a concern, but the health of the inherited network can exacerbate problems. There are many points in the infrastructure that can cause latency. Using the following diagram, we will move through each section.



A. ROUTERS AND ROUTING PROTOCOLS

Each time a router changes a packet header (for instance the time to live counter), must encapsulate one protocol into another, alter packet size, or wait for retransmissions coming into or out of the network, latency is introduced. Maximizing routers with MPLS (MultiProtocol Layer Switching), large amounts of cache to buffer tables and entries while performing functions in RAM will improve performance on this link. In addition, verifying that all code revisions between all routers on the network are current will help. In most cases, it is advantageous to have multiple routers rather than multiple interfaces within one router to allow for load balancing of applications.

B. FIREWALLS AND SECURITY DEVICES

Security devices are outside the scope of this document, however with the newer appliance based devices, processing happens at wire speed rather than in software to help reduce latency. Any such devices should be checked against similar applications for an accurate accounting of transactions per second.

C. SERVER FARMS/STORAGE FARMS

A combination of hardware and network interfaces is available for each server and storage device. For years fiber has been a de-facto standard for storage towers. However, with the new 10GBASE-CX4 (Ethernet) standard's adoption, copper is rapidly gaining popularity in this space. New abilities to cluster servers and create grids to share processing power are increasing the transaction per second barrier to new limits. This also affords companies the benefit of load balancing and better transaction monitoring. Adding 10GBASE-T or similar 10Gb/s speeds will exponentially increase transaction power across the network and between grid based servers.

DATA STORAGE, RETRIEVAL AND PROCESSING CONCERNS

Traditional systems take data, store it to a disk then run processes against the stored data. The time to write to the disk and then process the data can be unacceptable in some applications. Further, with Storage Area Networks (SANs), the time to access the storage over the network and then retrieve the data to a processing machine can further lengthen the retrieval time (latency) if not implemented properly. Newer programs such as streaming SQL, enhanced cache abilities and increased memory can allow data to be processed first in faster memory and then written to disk. The speed of disk drives has also improved allowing for faster read and write access.

Data stores are also mirrored in real time rather than having to be read from tape in a failure. This hot mirror flexibility provides additional protection in case of a fault as well as the ability to load balance between disparate storage units.

D. INFRASTRUCTURE IN THE DATA CENTER ENVIRONMENT

The new 10GBASE-T standard for 100m of copper on Category 6A or above cabling is likely to drive increased copper acceptance in the storage community. 10GBASE-T copper will also allow switches and servers to communicate with a significant increase in speed at a fraction of the cost of fiber electronics.

10GBASE-CX4 uses twinax cables that cannot be field terminated and are available from various manufacturers. They do however have a distance limitation to 10-15m, depending on the manufacturer.

Fiber distances vary with the grade of fiber and the associated electronics. Fiber also has a minimum distance based on the speed of the equipment and various manufacturer requirements so that the laser at one end does not overrun the laser at the other. See the section G below for more information.

E, F. SWITCHES AND UPLINKS

Switches, of course, will have very similar issues to routers as far as abilities, load balancing and memory are concerned. However, hidden latency can reside in the switches if the cabling exceeds maximum distance or the uplinks are not of the proper fiber type. For instance, in gigabit transmissions, older 62.5 micron fiber channels over 220m will require mode conditioning patch cords at each end. If these are not provided, it is entirely possible that the links will not work error free for reduced latency. With fiber installations, link loss must be tested. Poorly performing fiber will cause retransmissions within the network which could significantly contribute to increased latency.

Likewise, poorly performing cabling channels or copper channels that exceed 100m in length can cause a switch to autonegotiate down to lower speeds, or in some cases even half duplex. It is important to run statistics on switches to be sure that cabling channels are not contributing to lowered switch performance in these situations.

G. HORIZONTAL NETWORK LINKS AND OVERALL CABLING HEALTH

It is important in any of these configurations to assure that the infrastructure performs solidly. All links should be installed, and tested with a tester that has recently been calibrated and loaded with the latest hardware and firmware. One mistake that companies make is to use periodic or random testing, or worse, to use the rating on the cable jacket to assume performance levels. This will not provide a level of protection needed for latency issues and error free operation. Improper matching of categories and/or grades of fiber can cause performance degradation below an expected speed. In a recent report from NewWorld Telecom, 83% of category 6 manufactured patch cables failed testing, with field terminated patch cables faring far worse.

Moves, adds and changes can also have an adverse affect on cabling installed channels, which may degrade their performance. Likewise, noise can be introduced if the channels are not properly routed around noise sources. The more the channel has been degraded, the more susceptible it will be to noise. Noise can introduce unwanted signals on the cable (although this is not an issue with screened F/UTP or fully shielded S/FTP systems). Sufficient noise immunity must exist to prevent unwanted coupling of noise onto the channel.

Furthermore, due to their improved alien crosstalk margins and the potential to support reduced digital signal processing (DSP), fully shielded systems could provide a significant latency benefit for future applications. In the IEEE 802.3an standard, the maximum transceiver latency is 2.56us. Many initial 10GBASE-T transceivers may have a difficult time meeting this mark with UTP systems due to the requirement to eliminate FEXT loss (far end cross talk) from each channel. In a recent Design Con presentation sponsored by KeyEye Communications and Siemon, category 7_A/class F_A cabling has FEXT loss levels at least 20dB lower than a comparable category 6_A/class E_A channel. This results in a reduced need for FEXT loss cancellation in the channel, which supports reduced latency during signal processing. Results of ongoing research and benefits related to the performance of screened and shielded systems will likely be used as incentives for these systems to be specified as the media of choice for the next generation of Ethernet applications.

IN CONCLUSION

If a financial institution owns a building or is likely to occupy the building for many years, it is highly recommended that a cabling system capable of 10Gb/s transmission speeds be utilized. To further improve latency performance, a screened category 6_A or fully shielded category 7_A/class F_A cabling system such as TERA® will provide enhanced performance due to noise immunity and a decreased delay (propagation and skew) within the cable. For more detailed information on these topics, please visit www.siemon.com.

The Americas

Watertown, Connecticut
Tel: (1) 866-548-5814 (US)
Tel: (1) 888-425-6165 (Canada)

Central & South America

Bogota, Columbia
Tel: (011) 571 317 2121

Europe/Middle East/Africa

Chertsey Surrey, England
Tel: (44) (0) 1932 571771

Asia Pacific

Shanghai, China
Tel: (86) 21 6390 6778

Japan

Tokyo, Japan
Tel: (81) 3 5437 1580