

# ULTRA-LOW-LATENCY NETWORKING

Facilitating real-time equities, futures, options and other financial transactions nanoseconds faster can be worth millions.

## Executive Summary

Latency is such a hot topic on Wall Street that some believe it's the next big thing.

It's so big that well-known IT veteran Saro Jahani, currently CIO of the Direct Edge Stock Exchange, publicly stated at a conference that he wouldn't be surprised "if the newest title on The Street was that of a chief latency officer."

Although the "CLO" concept may be yet in its infancy, the drivers for Jahani's remarks are easily understood. With the most recent estimates on the impact of latency showing that every millisecond lost can result in \$100 million per year in lost opportunity, it's no wonder IT departments at securities firms are making latency a top priority.

In the global electronic financial marketplace, trading systems are only as good as the infrastructure upon which they execute. Securities firms, regardless of size, are adopting low-latency networks to remain competitive in an industry increasingly driven by computerized algorithm trading.

## Table of Contents

- 
- 2 The Situation

---

  - 2 Types of Network Latency

---

  - 2 WAN Design

---

  - 2 Trading Network Design

---

  - 3 Other Components to Reduce Latency

---

  - 4 Low-Latency Storage

## The Situation

When it comes to algorithmic trading, everyone's getting into the game. No longer limited to a select group of large Wall Street players, now securities firms of all sizes use algos to gain competitive advantage and improve profitability.

For example, a 2012 survey, conducted by *Automated Trader* magazine, found that more than 80 percent of buy-side firms currently use systematic algorithms to trade certain instruments. And that number is expected to rise to 97 percent within the next three years.

Further, the subset of algo trading called "high-frequency trading" (HFT) now accounts for more than 50 percent of all market volume, according to estimates posted by the TABB Group, a financial markets research and advisory firm, in July.

In a world where so much trading is systematic, millisecond delays matter. Low-latency, high-performance networking, therefore, has become a competitive imperative.

## Types of Network Latency

When it comes to networking latency in the securities industry, there are two general types: partner-based and firm-based latency.

**Partner-based latency:** This type of latency is the time required by partner systems, such as an exchange or market data provider, to process messages within their systems. A firm clearly can't control the performance of its trading partners' networks, but it can select partners based on network performance. To do so, insist on actual latency measurements, not just service-level agreement (SLA) numbers.

**Firm-based latency:** As the term implies, this type of latency originates from within a firm's networking systems. The source of latency – and how it's driven out – depends upon the type of network. These fall into three categories:

- Wide area networks (WANs), which communicate messages to and from trading and campus networks and external providers
- Trading networks, which are dedicated to trading and optimized for the purpose
- Campus networks, which are for all business functions not related to the actual trading of equities

## WAN Design

Communicating trading information to partners frequently occurs over a WAN. Because few small and mid-sized firms lay their own fiber-optic networks, what's most important is asking the service provider whether it utilizes low-latency architectures. These include:

- **Passive optical components:** In general, a low-latency WAN is constructed of passive optical components that

do not introduce any significant delay. Among these are multiplexers, demultiplexers and optical add-drop multiplexers as well as either erbium-doped fiber amplifiers (EDFA) or Raman amplifiers.

- **DCM:** To remove pulse distortion over long distances, dispersion compensation modules (DCM) are superior to dispersion compensation fiber (DCF) due to DCM's enormous latency advantages.
- **Non-FEC transponders:** Transponders send, receive and optimize signals over a fiber network. Forward error correction (FEC) transponders introduce higher latency value than non-FEC transponders. Therefore, the latter are preferred.

Additionally, it's worth considering the benefits of leasing dark fiber that the firm can light itself versus a "lit" wavelength provided by a communications service. All things being equal, dark fiber will provide a lower-latency solution. If a lit wavelength makes more sense, however, require the provider to supply actual latency measurements rather than SLA metrics.

## Trading Network Design

In recent years, it has become common to build a network dedicated to algorithmic trading as a separate environment from other business functions. This enables the deployment of purpose-built, ultra-low-latency networking components.

One benefit of a separate trading network is cost-effectiveness. Ultra-low-latency components can require more upfront investments, making a dedicated trading network more affordable than deploying low-latency gear systemically.

Also, trading environments typically are not virtualized at this time. Current virtualization technology introduces latency, which defeats any benefits achieved by increasing either CPU or storage utilization efficiency. By segregating the trading environment, a firm can still leverage virtualization in the campus network and gain all the benefits virtualizing offers.

Additionally, creating a dedicated trading network permits moving the environment to a colocation facility, eliminating the latency discussed above when sending data across a WAN.

### A look at the factors affecting latency over the last three decades

	1983	2011	Improved
<b>CPU Speed</b>	1x10Mhz	4x3GHz	> 1,000x
<b>Memory Size</b>	≤2MB	8GB	≥ 4,000x
<b>Disk Capacity</b>	≤30MB	2TB	> 60,000x
<b>Net Bandwidth</b>	3Mbps	10Gbps	> 3,000x
<b>Round-trip Time (RTT)</b>	2.54ms	80µs	32x

Source: *It's Time for Low Latency*, Stephen M. Rumble, Diego Ongaro, Ryan Stutsman, Mendel Rosenblum and John K. Ousterhout, Stanford University, 2011

Although sell-side firms have a track record of benefiting from colocation, buy-side firms are increasingly using colocation sites as their utilization of algorithmic and HFT trading grows.

## Other Components to Reduce Latency

To achieve algorithmic trading goals, securities firms also need speedy processors, accurate clocks and efficient applications to run on top of low-latency network infrastructures.

**Ultra-fast processors:** Six-, eight- and 12-core ultra-fast processors enable complex calculations and create accurate, detailed forecasts, simulations and data visualizations. By decreasing input/output latency and improving I/O bandwidth, today's ultra-fast processors help reduce application latency from microseconds to nanoseconds.

Ultra-fast processors rev up not only trading applications, but also other securities firm apps such as real-time risk management. Some manufacturers offer additional hardware features for enhanced data security and improved management capabilities.

In addition, chip companies continue to push the envelope, pointing to an even faster future. Intel began shipping experimental 48-core processors to research institutions in 2010. Last fall, AMD shipped 16-core processors based on its new Bulldozer architecture. Even startups are selling 16-, 34-, 64- and 100-core processors.

**Precise time synchronization:** To ensure precise time-stamping of market data, transactions, risk management and other applications, cutting-edge precision time appliances provide accuracy to the microsecond – some, even to the nanosecond – improving both regulatory compliance and competitiveness for capital markets firms.

The latest time synchronization appliances gather data from global positioning system data and feature precise clock management software that works in conjunction with synchronization hardware. They support industry standards such as the Network Time Protocol (NTP) and Precision Time Protocol (PTP) as well as high-speed network interfaces such as 10 Gigabit Ethernet and InfiniBand.

**Efficient applications:** In the wake of the global financial crisis, Wall Street firms are more focused than ever on cost-effectiveness, which extends to the return on investment their apps provide. Trading software makers therefore strive to design increasingly efficient transaction processing platforms that lower transaction costs while providing high-quality customer service.

Application optimization adds another layer of efficiency by automating the provisioning of computer resources in real time to respond to demand. This not only improves a firm's ability to respond to sudden increases in trading volume, but also reduces the total cost of operations.

## Anatomy of a Low-Latency, High-Performance Trading Network

Although networks at individual capital markets firms will contain technologies most appropriate for the needs of the organization, the following components can be essential to successfully meeting low-latency goals.

- **High-speed interconnects:** To ensure wire speed data rates, build trading networks using gear that supports 10 Gigabit Ethernet, InfiniBand or both. This reduces serialization delay while protecting networks against microbursts – short-term events that create network-clogging congestion.
- **Ultra-low-latency switches:** Mainstream switch manufacturers such as Arista Networks, Brocade, Cisco Systems, Juniper Networks, Mellanox Technologies and NetScout Systems all offer ultra-low-latency switches designed for high-performance trading networks. These specialized switches offer capabilities specific to such environments, such as latencies measured in nanoseconds. The three methods to measure latency are: first in, first out (FIFO); last in, first out (LIFO); and last in, last out (LIFO).
- **Message acceleration:** Although numerous reliable providers offer message acceleration software, a new breed of purpose-built, hardware-based acceleration appliances are now available from manufacturers such as IBM. They offer sub-millisecond latencies at volumes ranging up to millions of messages per second. Frequently, both Ethernet and InfiniBand networking are supported.
- **Next-generation security:** Similar to switches, purpose-built low-latency firewalls from manufacturers such as Fortinet and SonicWALL supply multiple functions for unified threat deterrence without impeding network traffic. Important characteristics include the ability to control applications and also prioritize traffic.
- **Network management technologies:** To ensure a network built to support low-latency fulfills its potential, it's also critical to take a holistic systems-based approach to network management. From that perspective, two essential technologies for low-latency environments are network packet brokers (NPBs) and application-aware network performance management (aaNPM).
- **NPBs:** These purpose-built appliances, from manufacturers such as NetScout, are intelligent switches that connect network monitoring and security tools to the underlying network. NPBs capture network behavior, provide packet-processing services (such as deduplication) and distribute actionable data back to security and monitoring tools, such as network performance management (NPM) devices.
- **aaNPM:** Although there are many NPM solutions, for trading networks it's vital to adopt a solution with the intelligence of an aaNPM. These products provide the end-to-end visibility into the performance of the critical applications that reside in low-latency, high-performance trading networks. Providers include Cisco, NetScout and Riverbed Technology.

## Low-Latency Storage

No discussion of low-latency, high-performance networking is complete without addressing latency contributed by hard disk-based storage.

It's well established that the performance of traditional disk storage is restricted due to the limitations of electro-mechanical devices, including access time, rotational latency and data transfer rates. In combination, such limitations can add significant delays into a network that is otherwise engineered for low latency.

Fortunately for the high-speed needs of capital markets firms, solid-state storage offers an efficient and cost-effective alternative.

Solid-state storage solutions use integrated circuits, rather than rotating magnetic media, for storing data and are sometimes also called "flash-based storage." Solid-state solutions come in familiar form factors, including drive-shaped devices or PCI Express cards. As a class, all types of solid-state storage devices are typically referred to as solid state-drives or SSDs.

To understand the orders of magnitude in performance that SSDs offer, a good metric is random input/output operations per second (IOPS). According to the Storage Networking Industry Association (SNIA), an independent industry organization, enterprise SSDs offer a 475 percent improvement in IOPS over comparable enterprise hard drives.

Once considered out of reach for small and mid-sized firms, solid-state storage solutions are now attractively priced, while costs continue to drop.

Best of all, adopting SSDs doesn't necessarily mean abandoning current hard disk investments. Some SSDs are deployed as "Tier 0" in an existing storage subsystem. Others types of SSD solutions can serve as a caching resource, allowing them to transparently run front-end current systems.

**To learn more about CDW's low-latency solutions, contact your CDW securities and investments account manager, call 888.706.4239 or visit [CDW.com/financial](http://CDW.com/financial)**



Mellanox interconnect solutions are some of the most deployed high-speed interconnects for large-scale simulations, replacing proprietary or low-performance solutions by providing low-latency, high-bandwidth, high message rate with complete transport offload for extremely low CPU overhead, remote direct memory access (RDMA) and advanced communications offloads. Mellanox's scalable high performance computing (HPC) interconnect solutions are paving the road to Exascale Computing by delivering very high scalability, efficiency and performance for HPC systems today and in the future.

[CDW.com](http://CDW.com)

## ARISTA

Arista Networks has adopted an architectural approach to building low-latency cloud networking. Its wire-rate, nonblocking and ultra-low-latency 1Gb/10GbE switches, along with Extensible Operating System (EOS), enable resilient cloud networks for transporting data, multimedia, storage and compute traffic. Architecting cloud networks, based on Arista's 7xxx series switch fabric, offers ways to do the following:

- Reduce network latencies across all four axes
- Future-proof the network infrastructure for latency demands
- Implement standard Ethernet, which is widely deployed and operationally well understood

[CDW.com](http://CDW.com)

## Myricom

Myri-10G, for high performance computing, provides the benefits of Ethernet with performance formerly available only in specialty interconnects. Myri-10G network adapters for HPC support 1.2Gbps MPI ping-pong data rates, with available network adapter bonding boosting that rate to 4.5Gbps with four network links. Kernel-bypass mode consistently delivers low latency, even on multicore systems. Zero copy implementation dramatically reduces host-CPU utilization, freeing up computing cycles to focus on computing tasks. Additionally, the technology can carry TCP/IP and UDP/IP traffic at line rate for legacy applications.

[CDW.com](http://CDW.com)



The information is provided for informational purposes. It is believed to be accurate but could contain errors. CDW does not intend to make any warranties, express or implied, about the products, services, or information that is discussed. CDW®, CDW-G® and The Right Technology. Right Away® are registered trademarks of CDW LLC. PEOPLE WHO GET IT™ is a trademark of CDW LLC. All other trademarks and registered trademarks are the sole property of their respective owners.

Together we strive for perfection. ISO 9001:2000 certified  
108144 - 121016 - ©2012 CDW LLC

