



ESG WHITE PAPER

Cisco Nexus Ultra-low Latency Solutions Accelerate High Frequency Trading

Design Innovation Enhances Efficiency Where Every Nanosecond Matters

By Bob Laliberte, ESG Senior Analyst; and Leah Matuson, Research Analyst

November 2020

This ESG White Paper was commissioned by Cisco and is distributed under license from ESG.



Contents

Introduction: Every Nanosecond Matters	3
Network Challenges.....	3
Optimizing HFT Environments	4
Utilizing a Next-generation Platform Is Key.....	4
Intent-based Network Logic Is Essential	5
Eliminating the Potential for Network Latencies	5
Cisco Ultra-low Latency (ULL) Portfolio Gains Strength with Exablaze Acquisition	6
Cisco High Performance, Flexible Connectivity Switches.....	7
Cisco Nexus 3548XL and Switch Management via Cisco Data Center Network Manager	7
Cisco Nexus 3550 Series (formerly Exablaze ExaLINK Switches)	7
Cisco Nexus SmartNIC (formerly Exablaze ExaNIC)	8
The Bigger Truth	8

Introduction: Every Nanosecond Matters

The speed of business continues to rapidly accelerate, especially for organizations in the financial sector, and particularly in the high frequency trading (HFT) market. HFT is an ultra-fast, computerized trading strategy using complex algorithms to analyze multiple markets and execute orders based on market conditions.¹ So, it's no surprise that every nanosecond matters because any increase in network latency could potentially result in millions of dollars in lost revenue and profit.

With an average transaction taking a fraction of a second to complete, financial organizations are focused on streamlining every aspect of the HFT process. To get a sense of how fast a trade is transacted, blink your eyes. The blink of an eye takes a fraction of a second. If each second comprises 1,000 milliseconds, a blink takes approximately one-third of a second.²

Any delay between receiving market data and making a trade could potentially result in millions of dollars in lost revenue and profit.

Ensuring the transaction process remains lightning fast from beginning to end means understanding the factors contributing to latency throughout the process. These factors include compute, data processing, algorithms, and network. With this in mind, it is imperative for organizations to optimize each area.

While network typically represents approximately 10% of the HFT process, it is a key foundational layer and can even improve compute capabilities. Organizations engaging in HFT must leverage innovative ultra-low latency (ULL) solutions to minimize latency and accelerate efficiency—but there are a variety of challenges that must be addressed.


Network Challenges

High frequency trading applications must be able to handle a number of requirements that include ultra-low latency feed and rapid analysis and data correlation. Consequently, when it comes to deploying an ultra-low latency network to meet their unique business needs, organizations must be poised to tackle any issues they may encounter. These include:

- **Proximity to the financial exchange.** It is vital to understand that proximity dictates latency. Organizations have to be colocated with the exchange or suffer a competitive disadvantage. Today, most exchanges have expanded their data center locations to accommodate and provide each organization with a latency-neutral connection to level the playing field (see Figure 1). This requires organizations to look elsewhere to gain an advantage.

¹ Source: Investopedia. [High-Frequency Trading \(HFT\), October 2019](#).

² Average duration of a single blink = 0.1-0.4 seconds. Source: Schiffman, H.R., *Sensation and Perception: An Integrated Approach*, New York: John Wiley and Sons, Inc., 2001. <http://faculty.washington.edu/chudler/facts.html>

Figure 1. The Value of Colocating with the Exchange


Source: Cisco

- **Latency due to outdated switches.** Due to the intense competitive pressures, these environments are constantly being upgraded to ensure the lowest possible latencies. Legacy environments have to make way to innovation and the latest generation of ultra-low latency switches.
- **Disparate vendor solutions.** Networks comprising multiple, disparate vendors could impact efficiency if they restrict the ability for organizations to have holistic visibility and control across the end-to-end network environment. As a result, network vendors with comprehensive portfolios will have an advantage.
- **Reliance on legacy compute architectures.** In ultra-low latency environments, pushing all processing to the central processing units (CPUs) is not an option. Organizations have to leverage solutions that enable time-sensitive processing to take place closer to the exchange; essentially, this means adopting field-programmable gate arrays (FPGAs) to execute certain functions at the edge of the server.
- **Inability to rapidly modify applications.** Organizations can't be held back when trying to change an algorithm or modify applications. To remain competitive and adapt to rapidly changing market dynamics, organizations need to make changes quickly, regardless of where the applications reside (CPU or FPGA). In these highly dynamic and time-sensitive environments, changes need to be made as quickly and easily as possible.

Optimizing HFT Environments

To gain a competitive advantage, organizations must understand the importance of selecting a network platform that can optimize their high speed trading environments, offering ultra-low latency, high availability, and security with ease of management.

Utilizing a Next-generation Platform Is Key

Accuracy and network visibility are essential for high frequency trading organizations, where every transaction must be precisely time-stamped to meet compliance and provide digital forensics.

It's no surprise that legacy platforms are incapable of supporting next-generation data centers and complex applications. While Network Time Protocol (NTP) has provided millisecond-level timing in packet-based networks, in the financial industry (and especially in the high frequency trading arena), millisecond-level timing is inadequate for organizations to achieve visibility into every process occurring on the network. A platform must be able to deliver the lowest possible latency and a high performance switch environment, incorporating nanosecond (one-billionth of a second) switches and picosecond (one-trillionth of a second) time-stamping.

High frequency trading requires a next-generation platform that can deliver ultra-low latency, high availability, and security with ease of management.

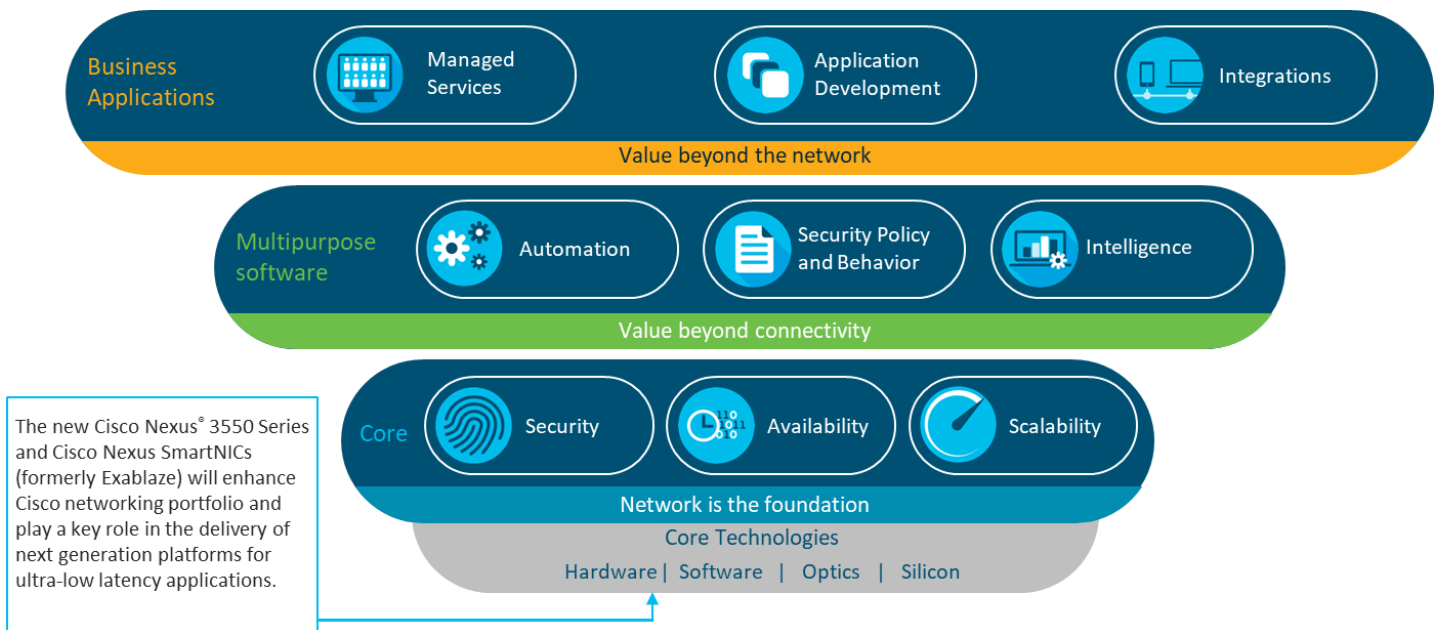
Intent-based Network Logic Is Essential

For high frequency trading environments, an intent-based network (IBN) could be a significant upgrade. The ability to implement closed loop automation that does not require human intervention (other than to define the intent or policies required) could dramatically improve operational efficiency.

Based on a combination of machine learning (ML) and orchestration, IBN allows IT to easily define policies detailing business intent to physical and virtual devices across the network—without having to rely on time-consuming, manual command line interfaces (CLI). Once defined, IBN automatically applies those policies to every device on the network, ensuring policy and orchestration operate as defined by business intent.

The ability to extend IBN models to the field-programmable gate array level could significantly improve HFT environments, efficiently orchestrating or altering processes automatically. See Figure 2 to understand how these technologies fit into an intent-based networking strategy.

Figure 2. Cisco's Intent-based Strategy Diagram



Source: Cisco

Eliminating the Potential for Network Latencies

Instantaneous timing is imperative for high frequency trading. Thus, being able to eliminate the potential for network latencies from the start is crucial. Given that, the expanse of space between decision-making (i.e., trading) activities and the financial exchange should be minimal. For example, the NYSE EuroNext Exchange expanded its data center from 75,000 to

380,000 sq. ft to accommodate high frequency trading processes within the exchange, eliminating the potential for WAN latencies.

With this in mind, where can financial organizations turn for help in achieving and maintaining success with their HFT activities? Enter Cisco and Exablaze.

Cisco Ultra-low Latency (ULL) Portfolio Gains Strength with Exablaze Acquisition

A worldwide network technology leader, Cisco continues to evolve and innovate across all areas—from hardware and software, to optics and silicon. Cisco's recent acquisition of Exablaze FPGA-based network products, technology, and engineering staff will enable Cisco to deliver a more comprehensive portfolio of next-generation, ultra-low latency solutions.

Exablaze, a designer and manufacturer of advanced network devices that reduce latency, will augment Cisco's strong domain expertise in the financial markets, strengthening Cisco's existing portfolio of ULL solutions, and complementing Cisco's current switching technology with ultra-low latency network switches and FPGA-based smart network interface cards (NICs). Cisco's intent-based networking strategy (see Figure 2) comprises investments that span in-house innovation, strategic acquisitions, and partnerships. Cisco's focus on ASIC and programmable FPGA-based next-generation platforms will be significant in building future platforms for ultra-low latency applications—enabling networking and security at scale, all powered by analytics.

To that end, Cisco will continue to integrate the ULL technology developed in the HFT proving grounds—think high performance computing, artificial intelligence (AI), and machine learning (ML). In addition, Cisco is committed to supporting Exablaze products and services, and is extending the roadmap for all existing customers.

With Exablaze, Cisco offers a comprehensive line of enhanced ULL solutions to assist financial and HFT customers, helping them achieve their unique technology and business objectives, and supporting their business-critical workloads. These include:

- Switches, smart network interface cards (NICs), development boards, high precision timing products, IP, and applications, encompassing:
 - ✓ Ultra-low latency hardware, software, and firmware design.
 - FPGA-based network switches and ULL NICs.
 - Next-generation architectures for future AI/ML activities.
 - High bandwidth network capture, analytics, and NVMe SS drives.
 - Firmware development kits (FDKs).
- Time measurement testing capabilities, including:
 - ✓ Global synchronized time distribution.
 - High accuracy clock synchronization.
 - Laboratory grade network measurement <70ps.

Cisco High Performance, Flexible Connectivity Switches

With the acquisition complete, the integrated lineup of next-generation, ultra-low latency solutions comprise the following:

Cisco Nexus 3548XL and Switch Management via Cisco Data Center Network Manager

Built from switch-on-a-chip architecture (SoC), the Cisco Nexus 3548XL switch is part of the Cisco Nexus 3500 series of switches offering line-rate Layer 2 and 3 switching at ultra-low latency, and high performance with flexible connectivity. The switch runs the Cisco NX-OS network operating system, offering comprehensive troubleshooting capabilities, scale, and Layer 3 feature sets for network session management, gateway traffic management, and multicasting needed by exchanges and high frequency traders.

Cisco Data Center Network Manager (DCNM) enhances data center implementations by offering organizations comprehensive management for Cisco Nexus, and multilayer distributed switching (MDS) data center LAN and SAN switching deployments (including control, automation, monitoring, visualization, and troubleshooting).

Introduced in 2011, the Nexus 3548 included innovative features for high frequency trading environments, namely the Cisco Algorithm Boost technology (Algo Boost) built into the switch application-specific circuit (ASIC) to realize Layer 2 and Layer 3 switching latencies of less than 200 nanoseconds(ns). While Layer 3 switching at scale with a “normal mode” latency of 250ns remains of utmost need today, Layer 1 technology introduced by Exablaze is now much better suited for traffic replication to and aggregation from trading servers, naturally displacing the appeal of WARP features on Nexus 3548 for nanosecond chasers.

Cisco Nexus 3550 Series (formerly Exablaze ExaLINK Switches)

According to Cisco, the Nexus 3550 series offers the lowest latency among marketplace solutions. Formerly Exablaze switches, the Nexus 3550 series devices deliver ultra-low latency switching and multiplexing functions, precision time-stamping, and programmability to facilitate mission-critical network applications. Nexus 3550-F and 3550-H Layer 1-only products allow for data monitoring and distribution applications to run with port-to-port latencies of less than 5ns, and switches with embedded FPGA components include additional “smart” features. Aggregation and Layer 2 features including MAC learning, VLAN tagging/trunking, and IGMP snooping are made possible, with packet latencies remaining below 100 ns in most cases.

All Nexus 3550 switches have timing synchronization and clock syntonization capabilities. The Nexus 3550-F HPT model is dedicated to time-stamping functions and is the highest precision network aggregation and monitoring solution available on the market. These capabilities and functions help traders design their networks with deterministic timing schemes in order to avoid costly congestions and packet collisions over non-IP circuitry.

The next-gen Nexus 3550 series offers:

- Layer 1 data distribution 5 nanoseconds (ns)
- Layer 2 packet aggregation 39 ns
- Layer 2 packet switching 95ns
- High-precision time-stamping 70 picoseconds (0.07ns)
- FPGA-based technologies.

Cisco Nexus SmartNIC (formerly Exablaze ExaNIC)

High frequency trading requires high performance computing solutions. The Cisco Nexus SmartNIC (formerly Exablaze ExaNIC) complements Cisco ASIC and FPGA investments, enhancing intent-based networking—key for financial and HFT organizations that require high performance and instantaneous connectivity with ultra-low latency to run trade algorithms, submit transactions, and receive stock ticker feeds.

The Nexus SmartNIC ultra-low frequency network interface card incorporates its own silicon, allowing certain network functions and trade processing to be offloaded from the CPU to reduce latency. The solution is x86-compliant, and optimized for extremely high performance, offering up to 10x latency performance when optimized with Cisco software and tools.

Organizations can enjoy high precision time-stamping and time synchronization features, vital for packet capture. And because all Nexus SmartNIC adapters are FPGA-based, adapter life can be extended by allowing field upgrades, while users can take advantage of its software-programmable functionality leveraging firmware development kits (FDKs).

The Bigger Truth

The HFT market demands the lowest possible latencies, where every nanosecond matters, and any delay will trigger a negative impact to the bottom line. Cisco understands the criticality of this delicate balance, and is committed to delivering next-generation, FPGA-based switching platforms and smart network interface cards to address the rapidly growing demand for ultra-low latency network performance optimization for mission-critical systems—but they're not going it alone. An indispensable part of this commitment is the Cisco acquisition of Exablaze.

Together, Cisco and Exablaze will address the growing need for performance-optimized networking with next-generation switching platforms and smart network interface cards. To that end, Cisco is ensuring support for Exablaze products and services, and extending the roadmap for all existing customers.

Cisco's innovative ULL Nexus 3550 and Nexus SmartNIC platform offer the benefits, support, and engineering resources to enable organizations to accelerate the time to value. By incorporating next-generation technology developed in the HFT proving grounds into its networking platforms, Cisco (with Exablaze) will expand the reach of this technology to enable high performance computing, artificial intelligence, and machine learning environments.

To learn more, please visit Cisco online:

<https://www.cisco.com/c/en/us/products/switches/nexus-3550-series/index.html>

<https://www.cisco.com/c/en/us/products/interfaces-modules/nexus-smartnic/index.html>

[What is low latency?](#)

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.



Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.



www.esg-global.com



contact@esg-global.com



508.482.0188