



Breaking the Low Latency Trading Barrier with Next-Gen Intelligent Interconnect

Mellanox SN2000 series Ethernet switches with Cumulus Linux

LOWEST LATENCY OVER 10GBE AND FASTER

With the advent of algorithmic trading, the need for low latency trading platform becomes very crucial. In these market data exchange environments, where, for example, when executing arbitrage strategies, every millisecond lost may result in ~\$100M lost opportunity. Hence, speed becomes the key differentiator, making it critical to establish the lowest possible latency between processing environments. The ultimate competitive advantage in businesses that supports low latency trading lies in the ability to access the liquidity pool faster than the competition, outperforming competitors during peak periods of intense activity. Low latency trading firms require an agile and high-performing infrastructure that addresses the enormous volumes and surges of high rate data speeds which are typically associated with periods of market volatility.

The Mellanox Spectrum™-based SN2000 series switch with Cumulus Linux is the industry's best solution to meet the demanding operational needs of low latency trading firms. The joint solution includes:

Hardware support:

- Lowest latency cut-through switching from 10GbE to 100GbE systems
- 16MB shared buffer to handle micro-burst traffic

Software support:

- Robust and scalable protocol support for low latency traffic (multicast with PIM, SM and SSM)
- Enhanced buffer monitoring with histogramming and watermark-based triggers

Uniform operating model:

- Enhanced command-line support (with NCLU, the Cumulus Linux network command line utility)
- Native integration with Linux system tools/applications

Automation: Extending the Linux server ecosystem into switches

- Native support for all automation systems, including Ansible, Puppet and Chef

HIGHLIGHTS



Lowest latency, 300ns



16MB shared buffer to handle Micro bursts



Real time buffer monitoring



Advanced traffic mirroring



Multicast at scale



Zero Packet Loss
www.zeropacketloss.com



IEEE 1588 PTP



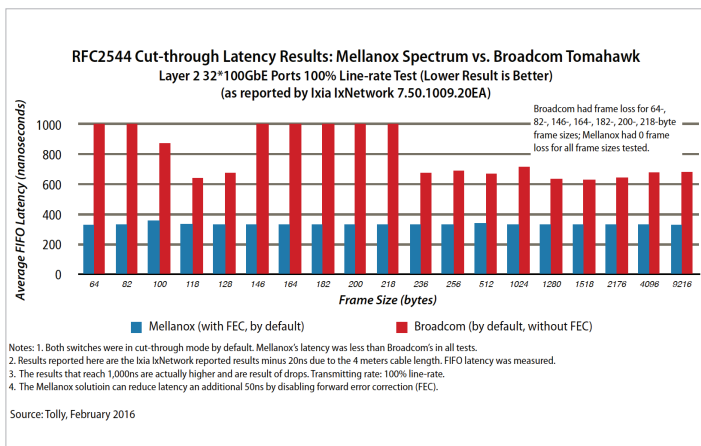
Automation ready platform

The messaging nature of financial information exchange creates bursty traffic which can overwhelm network elements. Cumulus Networks and Mellanox technologies collaborate closely to overcome microburst effects. Spectrum's superior shared buffer design enables absolute latency fairness with predictable wire speed performance and zero packet loss, combined with unmatched resilience to microbursts. Cumulus Networks couples this with unprecedented flexibility in buffer configuration and monitoring to ensure that full system capabilities are available to the consumer.

EVERY NANOSECOND COUNTS...

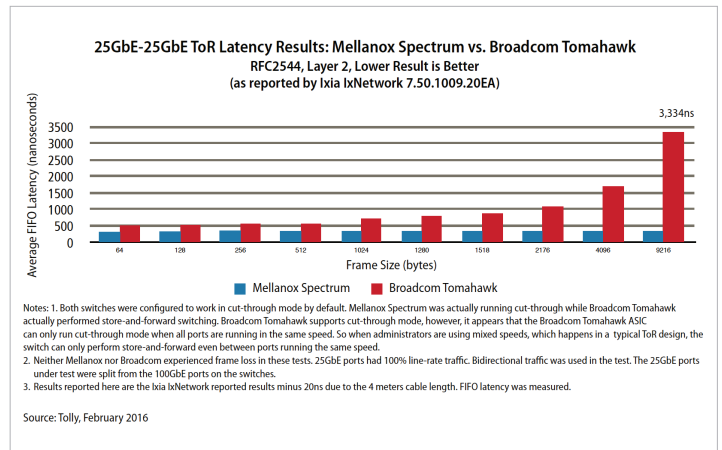
A switch has one job – to move every frame as rapidly as possible. Dropping frames and/or introducing excess latency can have a crucial impact on an application, as microseconds and nanoseconds represent a competitive edge, both for the demand and the provision of electronic execution services.

Mellanox Spectrum-based platforms demonstrate sustainable cut-through latency of 300ns with zero frame loss at all frame sizes and at all link speeds. Moreover, this latency is preserved even in mixed environments, combining slower and higher speed ports.



One of the fundamental premises for building a data center is that network infrastructure needs to be predictable in the way it performs. Predictability is measured in the consistency of the throughput and latency, regardless of the packet size or how fairly traffic is divided among subscribers when needed. For low latency trading, unfair distribution may lead to poor performance for tenants who lose the ability to forecast and control the traffic behavior.

The capability to build and maintain a synchronized, accurate timing solution is the basis for successful provisioning and management of High Frequency Trading (HFT) networks and applications. Using IEEE 1588 PTP, switches can deliver highly accurate precision time synchronization to applications within existing network infrastructure with no need to invest in and deploy a separate timing network. The Spectrum and Cumulus Linux low latency solution offers support for PTP together with the Cumulus Linux software stack.

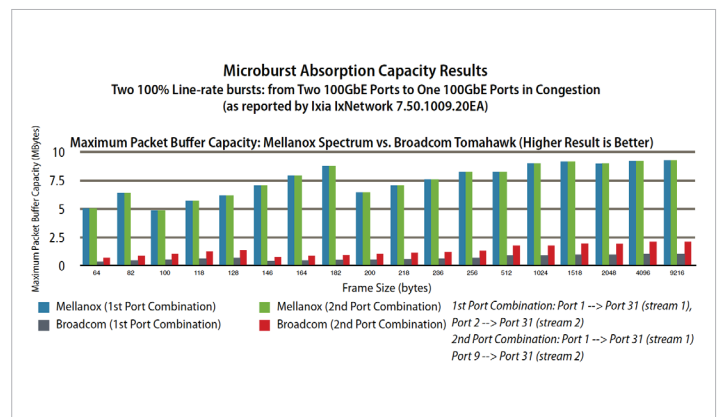


PREDICTABLE PERFORMANCE, EVEN DURING MICROBURSTS

A microburst is a traffic pattern that causes short-lived congestion in the network. This is often due to a period of high activity that causes network endpoints to send bursts of traffic into the network — for example, during high-volatility periods in the market. Even dropping a single packet of information because of a microburst that exceeds the network bandwidth can dramatically increase trading time, requiring retransmissions and possibly resulting in complete data loss and, therefore, transaction loss.

Due to the bursty nature of market data and the need to consume and aggregate various sources, every retransmission can translate into significant revenue loss. The challenge is to minimize trading latency end-to-end while handling peaks without packet loss.

Microburst absorption tests measure how many frames a switch can hold in its buffer while waiting for the output port to become available. The greater the buffer, the less traffic is dropped, thus avoiding performance degradation.



As you can clearly see, the variations between port combinations is minimal and clearly shows the benefit of the architecture.

ENHANCED BUFFER MONITORING AND VISIBILITY

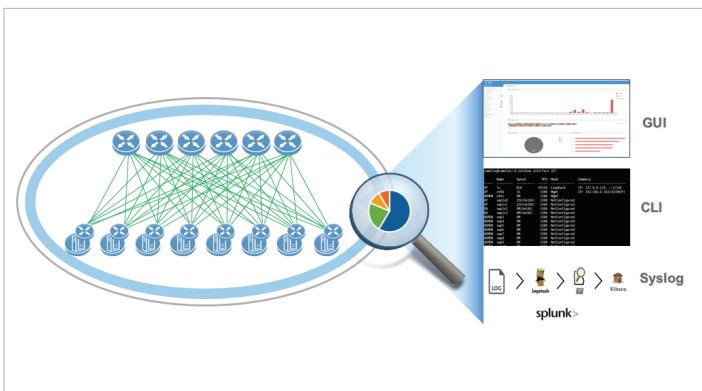
A microburst can overwhelm a non-optimized network and cause transactions to be lost or require retransmission, which has significant negative impact on a firm's financial performance. With the new 100GbE switches, 1 second can translate to more than a hundred million messages!

Rather than wait for the effects of network congestion to trickle up to the application layer, Congestion Monitoring Event can proactively detect impending congestion events, allowing preemptive capacity planning before the effect is seen at the application layer.

Spectrum provides real-time hardware buffer monitoring. The architecture enables the switch to detect delays that occur in traffic flows when packets are buffered. Detecting the presence of any congestion at the hardware level provides useful visibility into the switch. Awareness of the location, depth and duration of the congestion provides for additional optimizations to occur.

The ability to configure alerting thresholds on every switch port enables the creation of a syslog entry whenever ingress or egress queues for traffic exiting that port cause traffic to be delayed more than a few microseconds. The recorded syslog enables the trader to analyze the historical capacity performance of the network during high traffic/congestion events and allocate additional bandwidth where needed.

Cumulus Linux builds on that set of capabilities and enables Linux applications to register for and monitor buffer depths, highwater threshold crossings and full counter complements. This is available as a streaming, asynchronous interface, presented in JSON format for easy consumption by higher level tools. This combination enables a simple, fast and powerful way to monitor, analyze and pinpoint traffic issues in the network whenever they occur.



The ability to present, preserve and indeed prove latency when traversing a switch is integral to a good solution for high frequency trading. As show above, the combination of Spectrum with Cumulus Linux provides the best solution for predictable and invariant latency, guaranteed fairness and throughput.

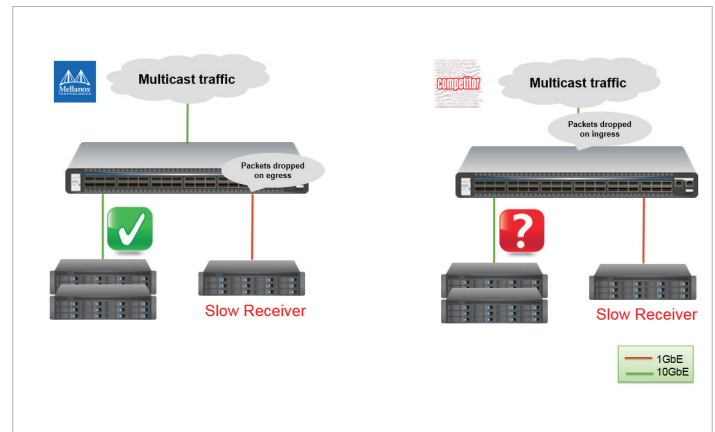
In addition, with innovative buffer monitoring techniques, buffer control tools and a high quality PTP implementation, this solution provides the ability to verify those capabilities and offers insights into the actual traffic patterns and behavior in effect.

MULTICAST BOOST...

IP multicast is a commonly used technology that was designed for "few to many" IP data transmissions, and is used prominently in financial market data feed networks. Through a physical connection, traders can receive market data (typically via TCP or UDP multicast) and send orders (typically via TCP or UDP unicast). Low latency trading platforms enable traders to execute millions of orders and scan multiple markets and exchanges in a matter of seconds. Thus, low latency traders are heavily dependent on multicast traffic at large scale.

Spectrum switches are powered by the Cumulus Linux software stack, with full support for multicast protocols, including Protocol Independent Multicast-Sparse Mode (PIM-SM), PIM-Source-Specific Multicast (PIM-SSM) and Multicast Source Discovery Protocol (MSDP). Cumulus Linux supports multicast at scale with more MACs, IPv4 routes and multicast groups than any other solution available in the market. All combined with extensive validation and interoperability testing to ensure fast mitigation from current 1/10GbE infrastructure to the 25/100GbE next generation deployments.

On top of the Cumulus Linux multicast stack support, Spectrum is optimized to multicast head-of-line (HOL) blocking, where a single slow receiver can cause congestion in a switch where the feed is pushed faster than the receiver's ability to absorb. For systems running multicast traffic, where the drop is done at the ingress, the switch "punishes" the fast servers as well, since the multicast packet is dropped before being replicated. Spectrum switches, on the other hand, drop only at the slow receiver egress port; thus, fast servers (part of the multicast group) keep on running intact.



SUMMARY

Mellanox's Spectrum, together with Cumulus Linux, represents the only solution in the market that delivers lowest latency and jitter at speeds of 10Gbps and faster. This joint solution delivers Spectrum's best performance regarding latency, jitter, microburst absorption and low power with the rich layer 3 stack and monitoring capabilities of Cumulus Linux. It's a perfect fit for customers with a need for low latency and high frequency multicast suite at scale.

Specifications

Switch Model	SN2700	SN2410	SN2100
Max 100GbE ports	32	8	16
Max 40GbE ports	32	8	16
Max 25GbE ports	64	64	64
Max 10GbE ports	64	64	64
Throughput	3.2 Tb/s	2 Tb/s	1.6 Tb/s
Packet Per Second	4.77 Bpps	2.98 Bpps	2.38 Bpps
Latency	300ns	300ns	300ns
CPU	Dual-core x86	Dual-core x86	ATOM x86
System Memory	8 GB	8 GB	8 GB
SSD Memory	32 GB	32 GB	16 GB
Packet Buffer	16MB	16MB	16MB
100/100 Mgmt Ports	1	1	1
Serial Ports	1 RJ45	1 RJ45	1 RJ45
USB Ports	1	1	1 Mini USB
Hot-Swap Power Supplies	2 (1+1 redundant)	2 (1+1 redundant)	No
Hot-Swappable Fans	4 (N+1 redundant)	4 (N+1 redundant)	No
Reversible Airflow Option	Yes	Yes	Yes
Typical Power (ATIS)	150W	165W	94W
Size (W x H x D)	1.72" x 16.84" x 27" (43.9mm x 427.8mm x 686mm)	1.72" x 17.24" x 17" (43.9mm x 438mm x 436mm)	1.72" x 7.87" x 20" (43.9mm x 200mm x 508mm)

About Mellanox

Mellanox Technologies is a leading supplier of end-to-end InfiniBand and Ethernet interconnect solutions and services for servers and storage. Mellanox interconnect solutions increase data center efficiency by providing the highest throughput and lowest latency, delivering data faster to applications and unlocking system performance capability. Mellanox offers a choice of fast interconnect products: adapters, switches, software, cables and silicon that accelerate application runtime and maximize business results for a wide range of markets including high-performance computing, enterprise data centers, Web 2.0, cloud, storage and financial services.

Learn more about Mellanox products and solutions:

www.mellanox.com

About Cumulus Networks

Cumulus Linux embodies native Linux networking. Supercharged versions of the kernel and other networking-related packages encompass the latest industry thinking in networking while retaining compatibility with the full range of software available in Debian. The SN2000 series running Cumulus Linux provides standard networking functions such as bridging, routing, VLANs, MLAGs, IPv4/IPv6, OSPF/BGP, access control, VRF and VXLAN overlays.

Learn more about Cumulus Networks operating system:

www.cumulusnetworks.com



350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085

Tel: 408-970-3400 • Fax: 408-970-3403

www.mellanox.com