



25 Gigabit Ethernet

Accelerating Network Efficiency and Performance while Lowering Costs and Easing Future Speed Transitions

Executive Summary

Operational efficiency and continuous availability are essential requirements for network operations in mission-critical business environments. When it comes to the network,

oversubscription, inefficiencies and congestion can result in delays and ultimately in a loss of productivity that can severely impact business operations and revenue for organizations. In order to remain competitive, IT infrastructure must enable network bandwidth to be cost-effectively scaled in support of next-generation server and storage solutions that are requiring considerably more throughput and efficiency. For this reason, the dominant next-generation server connection speed is moving to 25G. When faced with a network upgrade, why consider 10 GbE when there is something considerably faster that utilizes the same infrastructure for cabling? 25 GbE is a “single-lane” technology that can provide a 250 percent increase in bandwidth with greater switch port density than 10 GbE and is a better fit for the eventual shift to 100 GbE.

Increasing Bandwidth Requirements Demand 25GbE

The need for greater operational efficiency and increased automation in data centers is driving a requirement for a consolidated footprint and an increase in the density of network components being offered. With increasing levels of performance coming out of multi-core servers, more virtual machines are being packed per server and an increasing number of servers per rack. To keep up with

the bandwidth demand this is creating, organizations have turned to 10GbE. However, with increases in East/West traffic from dense virtualization deployments and new compute models like containers and flash storage, the demand for more bandwidth has increased in just a few short years. 25GbE is the ideal solution, providing two and a half times more bandwidth than 10GbE per-port. 25GbE accelerates workloads including database, virtualization, video streaming, and high-frequency trading (HFT). 25GbE is also perfect for accelerating the next-generation of data center infrastructure, including public and private cloud, hyper-converged, software-defined storage, and big data. And with the full advantages of 25GbE being offered today, communications between servers, servers and top-of-rack switching and servers to storage can reach their bandwidth requirements while reducing complexity, power usage, and Total Cost of Ownership (TCO).

The Only End-to-End Ethernet Solution

Mellanox is the only networking vendor who makes their own NICs, cables, optical transceivers, and switch silicon. By offering a comprehensive end-to-end Ethernet solution, Mellanox is able to reduce the risks associated with the deployment and management of high-performance networking. Designed to overcome current data center challenges by providing a highly flexible and scalable solution, allows businesses to deploy the hardware/software combinations that are best suited to meet their unique needs. This places Mellanox in a unique position to help answer organizations pent-up bandwidth demands by delivering networking infrastructure solutions that can handle the dramatically increasing requirements, while maintaining cost-effective from both an operational and capital expense perspective.

The Mellanox 25GbE Interconnects Effectively Deliver:

- Backward compatible with existing 10GbE
- Future Proof Upgrade Path:
 - Today – 1/10/40GbE
 - Tomorrow – 25/100GbE
- 2.5X Bandwidth of 10 GbE at 1.5X the Cost
- Greater Switch Port Density
- Reduction in CAPEX and OPEX
- 60% Lower Cost per Gb
- Up to 40% Lower Cabling Complexity
- Zero Packet Loss, Fair Behavior & Predictable Performance

High-performance, Efficient Ethernet Switching

Optimized for ultra-low latency fabrics, the Mellanox SN2000 series of Ethernet switches provide support for today's high-performance data center requirements and work as an ideal top-of-rack or spine switch for various data center deployments. With unique form-factors that include half-width switches that allow for two switches to be deployed side-by-side in single RU space, to easily meet high-availability requirements. SN2000 series has the ability to process data packets at full line rate, without dropping packets, to ensure network and application performance does not suffer. To better handle I/O contention, Mellanox employs a superior buffer to predictably manage I/O across all ports on the switch. This allows the switch to divide buffers resources fairly, similarly as it does bandwidth, allowing the full use of switch capacity. For example, when a microburst or incast (many-to-one) broadcast occurs, the network cannot allow one application or client a majority of the network capacity and accidentally allow others to starve. The SN2000 series switches provide fair and predictable performance to prevent these sort of predicaments.

Furthermore, in virtualized environments, a few hundred servers can host tens of thousands of VMs. This places a significant strain on the switches which require access to forwarding tables to complete delivery. This can cause severe network performance degradation when the network is operating under load. Mellanox switches offer full L2/L3 switching, routing, and data center bridging capabilities that remove network virtualization penalties and increase overall network efficiency.

Lastly, the SN2000 series of switches include the ability to support open APIs, standard interfaces and popular open source operating systems. This gives the switch the ability to disaggregate hardware from software allowing for the freedom of choice with different network operating systems. A switches ability to complement Open Ethernet, all the way down to the switch silicon level, can free an organization from vendor lock-in strategies and increase flexibility. The capability of "mixing and matching" allows for independent selection regarding best of breed hardware and network operating systems from different equipment and software vendors. This opens up the ability to achieve an optimal configuration that allows for better control of both capital and operating expenses.

Delivering Superior Utilization and Scale

Mellanox's 25GbE ConnectX-4 Lx Ethernet adapters are equipped with intelligent ASICs providing sub-microsecond latency and offloading mechanisms such as RDMA, Erasure Coding, TCP, UDP, as well as overlay network and OVS offloads. By bypassing the CPU, the server resources are freed up,

leaving more CPU cores available to analyze data or other compute tasks. This allows for higher scalability and efficiency within the data center. By doing this the adapter reduces application runtime and offers the flexibility and scalability to make infrastructure run as efficiently and productively as possible. Enabling data centers to leverage the adapter to increase their operational efficiency, improving server utilization, and maximizing applications productivity, while reducing total cost of ownership (TCO).

Improving efficiencies in today's virtualized environments is imperative when migrating to high-performance networking and the ConnectX-4 provides dedicated adapter technology such as SR-IOV and ASAP2 to guarantee isolation and protection for virtual machines (VMs) within the server. I/O virtualization gives data center administrators better server utilization while reducing cost, power, and cable complexity, allowing more Virtual Machines and more tenants on the same hardware.

Reducing Infrastructure Costs and Complexity

In efforts to reduce costs and ensure interoperability, Mellanox introduced new 25Gb Ethernet breakout cables and optical transceivers. Breakout cables (cables that connect to higher speeds at the switch and "fan out" to multiple lower speed links) allow deployment of four 25G connects at the servers to be aggregated to the top-of-rack switch at 100G connections. This lowers costs and simplifies cable management by reducing the number of cables required. New QSFP28 transceivers (quad-port 25Gb) offer lane speed increases, from 10G to 25G, from its QSFP+ predecessor. Utilizing the QSFP+ familiar form factor, QSFP28 transceivers support either copper Direct-Attach Cables (DACs) or Active Optical Cables (AOCs). QSFP28 ports ensure data centers can scale to 50 or 100GbE networks with simplicity. Mellanox offers QSFP28 transceivers that either uses VCSELs for shorter distances on multi-mode fiber or silicon photonics for longer distances. Silicon photonics enables QSFP28 transceivers to support any data center reach up to two kilometers or more, and due to stringent interoperability testing with Mellanox switches and adapters, provides a high degree of reliability and guaranteed integration.

Future Proof Networking

While 10GbE is well established, 40GbE is just making its way into the Enterprise data centers for Top-of-Rack uplinks, and provides a compelling upgrade path but based on "yesterdays" 10G line rate. Other market segments are moving much faster. Hyperscale systems will be moving to 50G (2x25G) to the server by early next year. HPCs will be moving to 200G to the server end by the end of 2017 and 400G is on the horizon.

Why deploy 25GbE? There are several reasons to consider. First, the price differences between 10G and 25G equipment is relatively small and yields a 2.5X increase in throughput. Second, a SN2700 32-port 100G switch can operate as 64 50G-ports and is much less expensive compared to two 40G 32-port switches (64-ports). 25GbE provides a simple and efficient migrate path to 100GbE. Upgrading a 50G network adapter with 50G QSFP28 transceivers to 100G is as simple as changing the optical fiber cable (<\$30) and reconfiguring the software. Mellanox offers 25G based interconnects in copper and optics that support 1x25G, 2x25G, and 4x25G 100G speeds, providing a level of upgradability unmatched in the industry. All the 25G line rate based cables and transceivers are backwards compatible to 10G and 1G line rates.

Network Acceleration

Mellanox adapters, switches and interconnects provide the network acceleration needed to run an efficient and scalable network infrastructure. The flexible framework offered by Mellanox high-performance Ethernet solutions work together to run at maximum performance and efficiency. This guarantees overall application performance while supporting web-scale IT now and into the future.

Conclusion

Networking is in the midst of a transition from the edge to the core with requirements to support more and more bandwidth-intensive applications. This is demanding organizations to look at higher-performance networking solutions to provide the end-to-end connectivity at bandwidths greater than 10GbE speeds while the requirement to reduce infrastructure costs and complexity continue to be high priorities. Mellanox is the first company to offer 25GbE solutions and the only company to provide a true holistic end-to-end solution. With a heritage in ultra-low latency and high-performance, Mellanox can help you accelerate your network and meet future scalability requirements. The transition to 25GbE has already begun and can utilize existing cabling infrastructure so it won't cost your organization a fortune to implement. By adopting 25Gb today will help lower future costs than a 10GbE implementation can, and will provide for a reduced TCO over time.

About Mellanox

Mellanox Technologies (NASDAQ: MLNX) is a leading supplier of end-to-end Ethernet and InfiniBand intelligent interconnect solutions and services for servers, storage, and hyper-converged infrastructure. Mellanox intelligent interconnect solutions increase data center efficiency by providing the highest throughput and lowest latency, delivering data faster to applications and unlocking system performance. Mellanox offers a choice of high performance solutions: network and multi-core processors, network adapters, switches, cables, software 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, network security, telecom and financial services. More information is available at www.mellanox.com.



350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085
Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com

MLNX-31-428