

Ethernet Storage Fabric (ESF)

Deliver performance and efficiency for scale-out storage and hyperconverged infrastructures

EXECUTIVE SUMMARY

An Ethernet Storage Fabric (ESF) leverages the speed, flexibility, and cost efficiencies of Ethernet to provide the foundation for the fastest and most efficient way of networking storage. An ESF is run on purpose-built switches which are optimized to deliver the highest levels of performance, lowest latencies and zero packet loss, with unique form factors and storage aware features. Best suited for scale-out storage and Hyperconverged infrastructures, an ESF is capable of simultaneously handling compute and storage traffic and is future proofed with support for the NVMe over fabric protocol. An ESF supports file, block, and object storage while delivering 3 times the performance, and lower latencies at a third the cost, compared to Fibre Channel. Mellanox Spectrum Ethernet switches were designed to handle the rigors of storage and are proficient in providing a true Ethernet Storage Fabric.

DATA CENTER CHANGES

The need for high-performance storage networking is being fueled by rapid deployments of multi-core servers densely packed with virtual machines and a move to all-flash storage arrays capable of blazing speeds. Communication between virtualized servers and ultra-low latency flash-based storage is a challenge for legacy IP and Fibre Channel switching architectures. As a result, there is an increasing demand for an efficient, reliable, storage-aware, high-performance storage network to connect servers and storage. Recent innovations in Ethernet switches are addressing these demands and are keeping pace with new storage trends such as the growing adoption of Cloud and Object storage, an area Fibre Channel does not address. Similarly, an explosion of file and object storage is being driven by distributed local storage from Big Data applications like Hadoop. Furthermore, as IT budgets continue to shrink, storage tiers based on performance requirements and access

SOLUTION HIGHLIGHTS

- Faster storage needs faster networks
- Most efficient 25/100GbE network solution for SDS and HCI
- Support for NVMe-oF with storage aware features
- Support lossless technology to ensure data delivery
- Ideal for All-flash storage and Hyperconverged infrastructure
- Support for file, block, and object storage
- Delivers 3 times the performance of Fibre Channel at a third the cost

patterns are becoming more common practice as we are now seeing about 80% of data on secondary storage. The increase in storage types and tiers makes a huge difference in storage planning, because all the file and object storage solutions must use Ethernet (can't use Fibre Channel). Secondary storage would also go on Ethernet because of the sheer expensive of a Fibre Channel SANs. For these reasons, we are seeing Ethernet storage growing rapidly.

Ethernet Storage Fabric

Everything a traditional SAN offers, but.... Faster, Smarter, & Less Expensive

PERFORMANCE

- Highest Bandwidth
- Latest Latency
- RDMA and Storage Offloads
- Native NVMe-oF Acceleration

INTELLIGENCE

- Automated Discovery & Provisioning
- Security & Isolation
- Monitoring, Management, & Visualization
- Storage-aware QoS

EFFICIENCY

- Just Works Out of the Box
- Flexible: Block, File, Object, HCI
- Converged: Storage, VM, Containers
- Affordable: SAN without the \$\$

Figure 1: Ethernet Storage Fabric characteristics

ETHERNET STORAGE FABRIC

It's important to make the distinction between a standard data center switch from one that is designated for an Ethernet Storage Fabric. First, storage switches must support lossless technology. This is to ensure data delivery and Quality of Service (QoS). Furthermore, advanced congestion control mechanisms must be employed to avoid buffer overruns and remove bottlenecks. Finally, it must be storage aware and offer sophisticated telemetry, monitoring, and management tools to simplify deployment and management while supporting a variety of industry-standard Layer 2 and Layer 3 protocols. Due to these, an ESF can support any storage architecture or protocol—including a hyperconverged infrastructure—at speeds from 1 to 100Gb per second. An ESF should be optimized for both primary & secondary storage environments and capable of delivering flash-ready performance and agility to keep pace with the most intense storage and data center workloads, not every data center switch can provide this.




Figure 2: Ethernet supports any kind of servers, storage, and data type, as well as different tiers of storage.

OPTIMIZED FOR STORAGE TRAFFIC

A true Ethernet Storage Fabric (ESF) is optimized for storage networking including buffers designed to handle bursty storage traffic, low latency and predictable performance to maximize data delivery and scale-out storage architecture support—all crucial attributes for today's business-critical storage environments. Must support block, object, and file storage, along with storage connectivity from iSCSI and all-flash to the newest NVMe over Fabric arrays. Additionally, an ESF must provide support for storage offloads, such as RDMA, to free resources and increase performance. Not only are ESF switches specifically optimized for storage, but they also provide better value than traditional enterprise storage networking at a third the cost of traditional Fibre Channel networks.

SPECTRUM – THE BEST ESF SWITCH

The Mellanox SN-series of Spectrum-based Ethernet switch family is designed and optimized for high performance, flexibility and value for all storage environments. SN-Series switches allow maximum flexibility with unique port configurations and speeds ranging from 1GbE to 100GbE. Ideal for storage with support for all storage types; file, block and object, and include support for HCI and Cloud infrastructure and software-defined storage by maintaining high-bandwidth at low latency to ensure data delivery and support for new media types such as NVMe.

All the SN-series models are non-blocking and allow enough uplink ports to build a fully non-blocking fabric. The latency is not only the lowest (300 nanoseconds port-to-port) of any generally available Ethernet switch, but the silicon and software are designed to keep latency consistently low across any mix of port speeds, port combinations, and packet sizes.

SN-series switches use an intelligent buffer design to ensure buffer space is allocated to the ports that need it most. This ensures data I/O is treated fairly across all switch port combinations and packet sizes. Other switch designs typically segregate their buffer space into port groups which makes them up to 4-times more likely to overflow and lose packets during a traffic microburst. This buffer segregation can also lead to unfair performance where different ports exhibit wildly different performance under load despite being rated for the same speed.

The SN-series family has features specific to optimizing current and future storage networking. These include support for Data Center Bridging (DCB), including DCBx, Enhanced Transmission Specification (ETS), and Priority Flow Control (PFC). iSCSI traffic can be specifically classified and prioritized using iSCSI-TLV. And the switches are designed to integrate with storage and network management tools as well as running containers on the switch to provide storage specific services.

STORAGE OPTIMIZATIONS

- Support for scale-out architectures
- DCB, including DCBx, ETS, and PFC
- ISCI-TLV
- Lowest latency of any mainstream Ethernet switch
- Differentiated storage ports - ports can be locked down for storage traffic
- The highest levels of (predictable) performance
- Zero avoidable packet loss regardless of packet size
- Performance fairness across any combination of ports
- Conjestion control through shared smart buffers
- Smart cut-through, allowing rapid packet forwarding even with mixed speeds



SN2010

- 1/2 width ToR
- 18x 10/25GbE + 4x 40/100GbE



SN2100

- 1/2 width ToR
- 16x 40/100GbE
- 32x 50GbE or 64x 10/25GbE



SN2410

- 48x10/25GbE + 8x 40/100GbE



SN2700

- 32x 40/100GbE
- 64x 10/25/50GbE

Figure 3: Mellanox Spectrum-based SN-Series of Ethernet Storage Fabric Switches.

CONCLUSION

Storage networks built on the Mellanox SN-Series of Ethernet Switches provide an ideal Ethernet Storage Fabric (ESF), optimized to deliver the highest levels of performance, industry-best latency, zero packet loss, and unique storage aware features and form factors. By leveraging the SN-series switches it's possible to build an Ethernet Storage Fabric or converged network capable of simultaneously handling compute and storage traffic. Support for iSCSI storage networks as well as for file, block, object and the latest NVMe over fabric storage is possible due to unique storage networking optimizations built into each switch, including dynamically shared buffers and predictable performance, making the SN-series ideal for storage environments. Future-proof your storage environment with support for faster speeds, new protocols and the ability to add storage services by either running them on the switches or integrating with the Switch OS. Put your storage on the fast track with Mellanox SN-series Ethernet switches and an Ethernet Storage Fabric.

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.

To find out more, visit our website: www.mellanox.com



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