



Increase Storage Performance for Data-Intensive Applications

Pure Storage and Mellanox Technologies increase storage performance with an enterprise-class end-to-end NVMe solution

Executive Summary

A significant increase in the demand to access data is coming from resource-intensive applications that need to process transactions quicker, or gain rapid insight to large data-sets. No matter what the need, data centers will be built with infrastructure that can support these modern workloads. This in turn is driving requirements for new levels of storage and networking performance. Non-Volatile Memory Express (NVMe) offers a high-speed, low-latency and efficient transport for accessing high-performance storage media. Pure Storage and Mellanox deliver an enterprise-class end-to-end NVMe solution that provides a massive optimization between storage controllers and hosts engineered to handle the world's most demanding application workloads.

NVMe over High Speed Ethernet with RoCE

Before NVMe, flash-based technology had to leverage legacy SATA/SAS physical interfaces. However, these interfaces were not designed for high-speed storage media. NVMe is both a host controller interface and storage protocol purposely developed to accelerate the transfer of data between host systems and flash storage. NVMe was designed to connect directly to PCIe slots, which are directly connected to the CPU for memory-like access speeds. Pure Storage FlashArray™, with 100% NVMe DirectFlash™, delivers 3x higher bandwidth and 4x IOPS compared to SAS/SATA SSD and can provide local access speeds via a Storage Area Network (SAN). With these unique capabilities, NVMe is poised to deliver the next major transformation of shared, networked storage into the data center.

End-to-End NVMe over Ethernet

In theory, enabling NVMe-oF can be done on just about any Ethernet switch, however, the practical limitations of doing so will be clear when using an Ethernet switch that is not designed to handle storage traffic efficiently. Storage traffic differs from small and very large packets and is bursty by

Highlights

- Pure Storage FlashArray//X combines with Mellanox end-to-end Ethernet ConnectX Intelligent adapters to deliver the industry's first NVMe-oF RoCE solution with enterprise data services.
- FlashArray//X with DirectFlash Fabric benefits include:
 - Up to 50% latency reduction compared to iSCSI
 - Up to 20% latency reduction compared to Fibre Channel
 - Up to 25% host CPU offload
- DirectFlash Fabric application improvements include:
 - Up to 50% better Ops/Sec for MongoDB
 - Up to 30% better latency and Ops/Sec for Cassandra
 - Up to 33% increase in transactions/sec with Maria DB4x

nature. Bursty traffic requires adequately sized buffers along with intelligent buffer management to reduce congestion and deliver fair, predictable performance.

Mellanox Spectrum switches are designed to handle storage traffic efficiently. Optimized for ultra-low latency lossless fabrics, with dynamic shared buffers helps to support the Ethernet Storage Fabrics (ESF) requirements of today's performance-demanding storage applications. Spectrum switches are ideal top-of-rack (TOR) switches for various storage deployments including NVMe-oF.

Mellanox intelligent ConnectX Ethernet adapters provide offloading mechanisms such as Erasure Coding, RoCE, T10/DIF, TCP and UDP offloads and overlay offloads. Support for NVMe-oF is offered to bypass the CPU with RDMA (RoCE) to free up CPU resources and further reduce latencies, allowing for higher efficiencies within the data center.

FlashArray//X

Pure Storage FlashArray//X is the world's first enterprise-class, all-NVMe flash storage array that delivers major breakthroughs in performance, simplicity, and consolidation. The FlashArray//X, introduces a new architecture called DirectFlash™, which moves away from legacy SSD architecture and enables direct communication to raw NAND devices. Unlike traditional SSDs, there's no flash controller, just raw flash, removing over-provisioning and performance roadblocks that exist within legacy SSD architectures.

Utilizing end-to-end NVMe and NVMe over Fabric (NVMe-oF) allows the FlashArray//X to take advantage of high-speed Mellanox RDMA enabled adapters to balance the performance and simplicity of direct-attached storage (DAS) with the scalability, flexibility and manageability of shared storage. With NVMe-oF shared storage can achieve similar performance as local storage while also providing valuable data services that enterprises demand. This allows the FlashArray//X to:

- **ACCELERATE MISSION-CRITICAL APPLICATIONS**
Offer new levels of performance to mission-critical applications with latency as low as 200 µs.
- **HYPER-CONSOLIDATE YOUR CLOUD**
Enable unprecedented performance density for mixed-workload consolidation in private clouds.
- **UNIFY APPLICATIONS**
Run a mix of classic business apps and web scale apps on a shared accelerated storage architecture.

Optimize App Access with DirectFlash Fabric


Getting to super low latency isn't easy with enterprise class data services. Pure developed DirectFlash™ Fabric to extend DirectFlash™ family outside of the array and into the fabrics via RDMA over converged Ethernet (RoCE). DirectFlash removes the legacy protocols from the storage array, as this is where the biggest bottle neck exists. This brings a huge performance gain inside the array and massive benefits to applications, delivering on the requirements for flash disaggregation with super low latency (up to 200 microsecond access) and end-to-end NVMe which completely revolutionizes the way applications access a FlashArray.

Results are up to 50% lower latency compared to iSCSI and up to 20% faster than Fibre Channel, which is known as the fastest fabric for enterprise storage applications. With RDMA offloading there is up to 25% reduction in CPU processing which provides more CPU to the application for better application results and can potentially lower per CPU software license costs.

Network fabric optimization with DirectFlash Fabric and NVMe-oF is an evolution that is a great candidate for cloud-native applications such as MongoDB, Cassandra, MariaDB, Hadoop, and Splunk. These cloud-native applications can realize a serious boost in performance. In fact, test results showed:

WEBSCALE & CLOUD-NATIVE APPS

POWERED BY END-TO-END NVMe & NVMe-oF



First Enterprise NVMe-oF Using RoCE

Mellanox end-to-end Ethernet infrastructure complements the performance of the Pure Storage FlashArray//X. The combined solution is the first commercially available storage solution to support NVMe-oF on RoCE with enterprise data services. The solution offers the biggest performance jump with up to 50 percent latency reduction compared to iSCSI, 20% latency improvement over Fibre Channel, and in some cases, 30% improvement over DAS while delivering a massive optimization between storage controllers and hosts over Ethernet with support for 25, 50, and 100GbE speeds. This allows the solution to deliver all-NVMe performance for a variety of workloads.



Conclusion

The benefits of NVMe include significant performance advantages between storage and servers as well as more efficient CPU utilization. Together, Mellanox and Pure Storage leverage these benefits to enable shared networked storage that is capable of displaying similar performance as direct attached flash offering enormous value for customers looking to modernize their storage networking infrastructure including substantially higher performance over iSCSI and FC SANs.

Want To Learn More?

Pure Storage FlashArray//X enterprise-class, all-NVMe & NVMe-oF flash storage array:

<https://www.purestorage.com/products/flasharray-x.html>

Mellanox end-to-end Ethernet Storage Fabric:

<http://www.mellanox.com/ethernet-storage-fabric/>



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

©2019. Mellanox Technologies. All rights reserved.

Mellanox and Mellanox logo are registered trademarks of Mellanox Technologies, Ltd. Mellanox Spectrum is a trademark of Mellanox Technologies, Ltd.
All other trademarks are property of their respective owner.

MLNX-423558315-99363