

## Solution Brief

# NetApp E-Series NVMe-oF (InfiniBand)

One of the first Nonvolatile Memory Express over Fabrics enterprise targets in the market

### Key Business Benefits

- Best price/performance value, with sub-100-microsecond latency
- Running over high-speed 100Gbps InfiniBand transport layer powered by Mellanox technology
- For mission-critical big data analytics workloads where data access response times are critical
- Support for persistent reservations that allow cluster software on the host side to access shared namespaces on the array

### The Challenge

Increasingly, organizations are looking for ways to improve the speed and responsiveness of the applications that control their key business operations. Because the performance of these applications is tightly linked to time to market, revenue, and customer satisfaction, it is critical that they operate at the highest levels with maximum efficiency. Getting value and insights quickly and reliably from a range of mixed workload environments can differentiate your organization from the competition and accelerate time to market.

### The Solution

The midrange NetApp® EF570 all-flash array is an all-SSD storage system that can turbocharge access to your data and increase its value. Requiring just 2U of rack space, the EF570 all-flash array combines extreme IOPS, sub-100-microsecond response times, and up to 21GBps of bandwidth with leading, enterprise-proven availability features, including:

- Redundant components with automated failover
- Intuitive storage management with comprehensive tuning functions
- Advanced monitoring and diagnostics with proactive repair
- Nonvolatile Memory Express over Fabrics (NVMe-oF) support, providing faster performance and investment protection
- SANtricity® Cloud Connector, enabling backup to the cloud and data mobility across NetApp systems

Combined, these capabilities provide the best price/performance, configuration flexibility, and simplicity in a compact package to help you make decisions quickly and securely.

### Introduction to NVMe-oF

NVMe has become the industry standard interface for PCIe SSDs, with a streamlined protocol and command set and fewer clock cycles per I/O. NVMe supports up to 64K queues and up to 64K commands per queue, which make it more efficient than the existing SCSI-based protocols such as SAS and SATA.

The introduction of NVMe-oF makes it more scalable while still benefiting from low latency and small overhead. The NVMeexpress.org specifications outline support for NVMe-oF over Ethernet, Remote Direct Memory Access (RDMA), and Fibre Channel (FC).



NetApp Technology Partner



NetApp is a Promoter Member of the NVM Express group

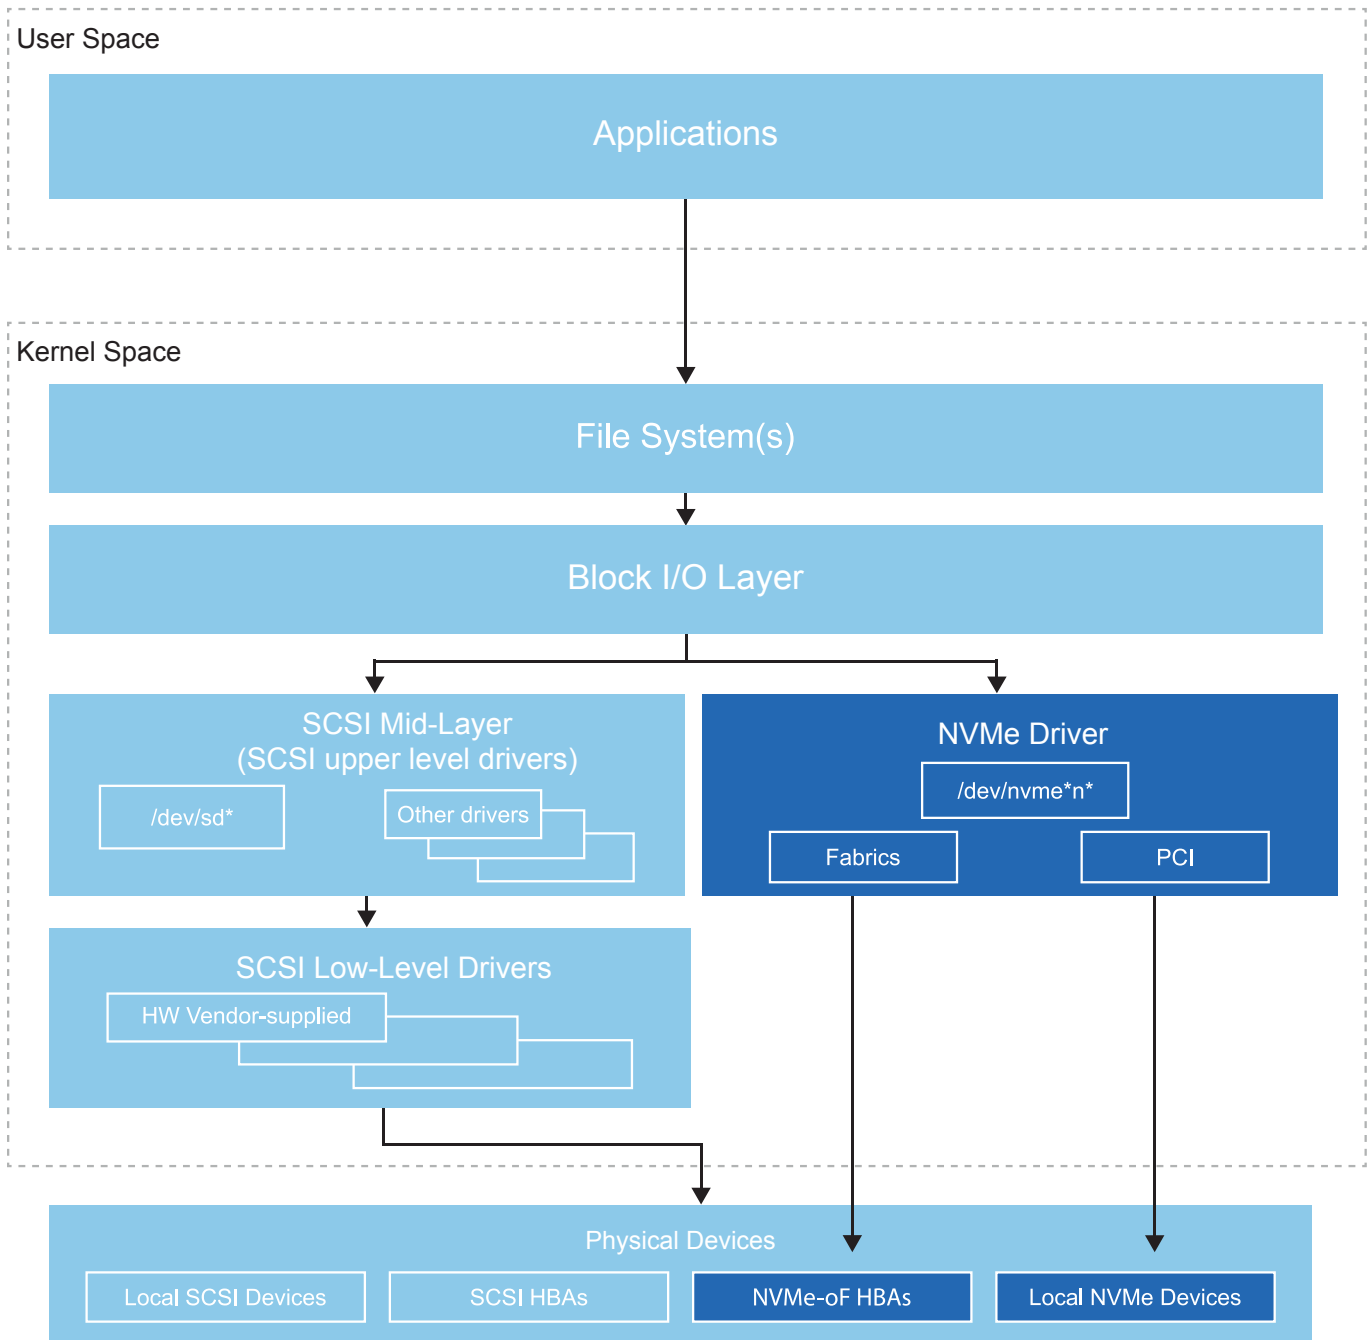


Figure 1) Linux OS driver stack.

### NVMe-oF vs. SCSI

One of the advertised advantages of NVMe (and NVMe-oF), as compared to SCSI, is that it can support lower latency I/O not only because the devices are faster but also because of some inherent advantages in the host OS driver stack (Figure 1). Because of this, I/O spends less total time getting from the application to the storage, thus reducing response times.

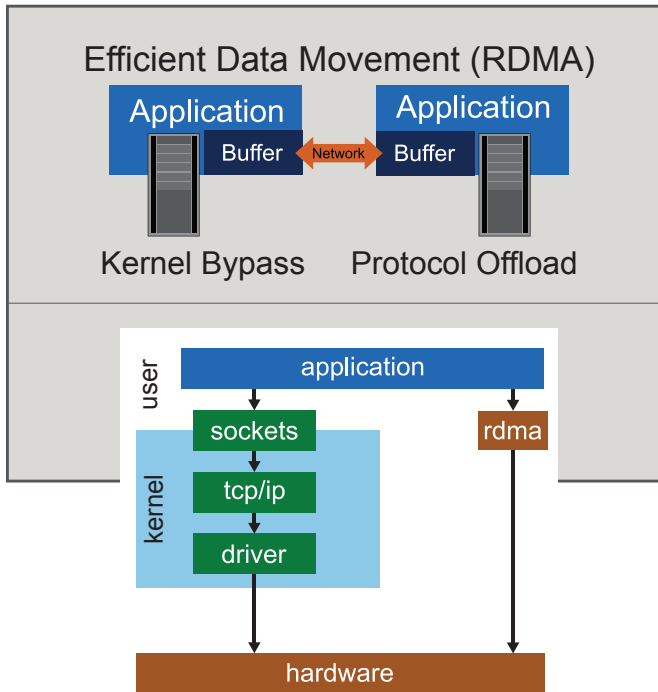


Figure 2) RDMA with Mellanox ConnectX™ InfiniBand Adapters.

### Remote Direct Memory Access (RDMA)

RDMA is a technology designed for high-performance compute environments that enables a low-latency transfer of information between compute nodes at the memory level. With Mellanox ConnectX™ InfiniBand Adapters, this function is offloaded to the network adapter to bypass the operating system network stack. This allows the adapters to work directly with the application memory, eliminating the need to involve the CPU while providing a more efficient, faster way to send data.

### NetApp E-Series

Although most implementations on the market today focus mainly on just adding NVMe drives to the back-end storage while keeping the front end to the host SCSI based, NetApp E-Series has taken a separate approach. The NVMe-oF is supported from the host to the front end of the EF570/E5700, while the back end is still SCSI based with the SAS drives. (See Figure 3.)

The NetApp E-Series implementation is done over InfiniBand for several reasons:

- InfiniBand has RDMA built into it.
- E-Series already has a long history and experience of supporting other protocols over InfiniBand (SCSI based): for example, iSCSI Extensions for RDMA (iSER) and SCSI RDMA Protocol (SRP).
- The same hardware on EF570/E5700 that runs iSER or SRP can run NVMe-oF (but not at the same time).
- All three protocols can coexist on the same fabric and even on the same InfiniBand host channel adapter (HCA) port on the host side. This allows customers with existing fabrics running iSER and/or SRP to connect the EF570/E5700 running NVMe-oF to the same fabric.
- All the InfiniBand components in the fabric (NetApp EF570/E5700, switches, HCAs) can negotiate the speed down as needed (enhanced data rate [EDR] 100Gbps, 14 data rate [FDR] 56Gbps, quad data rate [QDR] 40Gbps), which makes it easy to connect to legacy components with lower speeds.

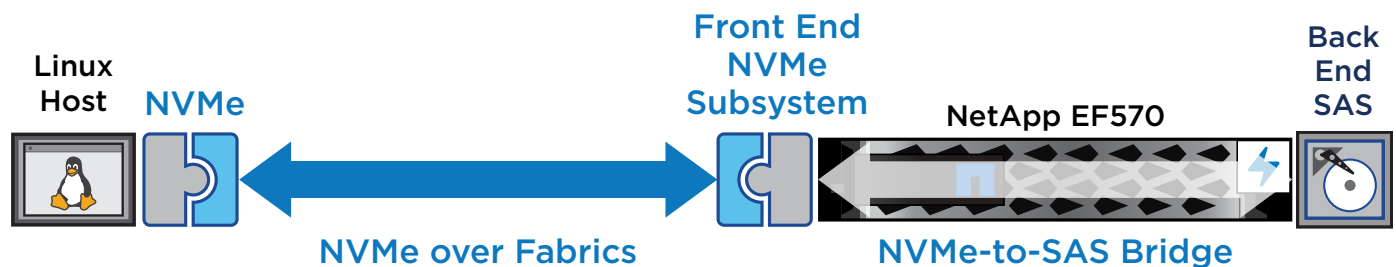


Figure 3) NVMe-oF front end on E-Series.

## E-Series Interoperability with Third-Party Components

Supported components:

- Operating systems: SLES 12 SP3, RHEL 7.4
- HCAs: Mellanox FDR and EDR HCAs
- Switches: Mellanox FDR and EDR InfiniBand switches
- Topology: fabric attached and direct attached




Figure 4) E-Series interoperability.

## Coexistence Between NVMe-oF, iSER, and SRP

All three protocols can coexist on the same fabric and even on the same InfiniBand HCA port on the host side. This allows customers with existing fabrics running iSER and/or SRP to connect the EF570/E5700 running NVMe-oF to the same fabric.




Figure 5) NVMe-oF, iSER, and SRP on one fabric.

## About Mellanox

Mellanox 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 multicore 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. [www.mellanox.com](http://www.mellanox.com)

## About NetApp

NetApp is the data authority for hybrid cloud. We provide a full range of hybrid cloud data services that simplify management of applications and data across cloud and on-premises environments to accelerate digital transformation. Together with our partners, we empower global organizations to unleash the full potential of their data to expand customer touchpoints, foster greater innovation and optimize their operations. For more information, visit [www.netapp.com](http://www.netapp.com). #DataDriven