

Solution Brief

NVMe-oF on NetApp E-Series and EF-Series

A complete portfolio with NVMe/IB, NVMe/RoCE, and NVMe/FC

Key Benefits

- Offer superior price/performance value, with less than 100-microsecond latency
- Run over a high-speed 100Gbps InfiniBand (IB), RDMA over Converged Ethernet (RoCE), or 32 Gbps Fibre Channel (FC) transport layer and network connections
- Support mission-critical big data analytics workloads in which data access response times are critical
- Support persistent reservations that allow cluster software on the host side to access shared namespaces on the array

The Challenge

Increasingly, organizations like yours are looking for ways to improve the speed and responsiveness of the applications that control their key operations. Because the performance of these applications is tightly linked to time to value and to stakeholder 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 and accelerate efficiency.

The Solution

The midrange NetApp® EF600, EF570, and E5700 storage systems turbocharge access to your data and increase its value. Requiring just 2U of rack space, the systems combine extreme IOPS, less than 100-microsecond response times, and up to 44GBps of bandwidth for EF600 and up to 21GBps for EF570 and E5700. You also get industry-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
- NVMe over Fabrics (NVMe-oF) support, providing faster performance and investment protection

Combined, these capabilities give you industry-leading price/performance, configuration flexibility, and simplicity in a compact package to help you make decisions quickly and securely.

NetApp E-Series and EF-Series systems with NVMe-oF are optimal for workloads that demand extremely low latency, such as high-performance computing, Oracle, Microsoft SQL Server, and data analytics.

Introduction to NVMe-oF

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

The introduction of NVMe-oF makes NVMe more scalable while still benefiting from low latency and small overhead. NVM Express specifications outline support for NVMe-oF over Ethernet, over remote direct memory access (RDMA), and over FC.



NetApp Technology Partner



NetApp is a Promoter Member of the NVM Express group

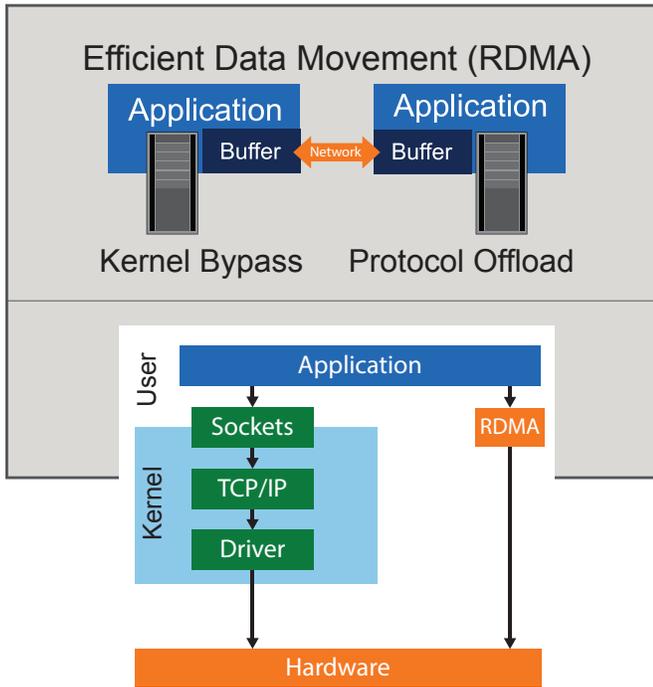


Figure 1) RDMA with Mellanox ConnectX InfiniBand Adapters.

NVMe-oF Versus SCSI

One of the advertised advantages of NVMe (and NVMe-oF) versus SCSI is that it can support lower-latency I/O because the devices are faster and because the host OS driver stack has inherent benefits. As a result, I/O spends less total time getting from the application to the storage, thus reducing response times.

RDMA

RDMA is a technology for high-performance compute environments that enables a low-latency transfer of information between compute or storage nodes at the memory level. With Mellanox ConnectX InfiniBand Adapters (Figure 1), this function is offloaded to the network adapter to bypass the operating system network stack. Adapters therefore work directly with the application memory, eliminating the need to involve the CPU while providing a more efficient, faster way to send data.

RoCE

RoCE is a standard protocol that enables RDMA efficient data transfer over Ethernet networks. RoCE enables transport offload with hardware RDMA engine acceleration and gives you superior performance.

With a 100Gb NVMe over RoCE (NVMe/RoCE) host interface for EF600, EF570, and E5700 systems, you can achieve consistent microsecond response. Both the 100Gb NVMe over InfiniBand (NVMe/IB) and NVMe/RoCE host interface cards are powered by Mellanox technology.

NetApp E-Series and EF-Series

To optimize the price/performance ratio, NetApp E-Series and EF-Series offer two primary products for NVMe-oF:

- EF600 is a complete end-to-end NVMe system that supports NVMe-oF from the host to the front end and all the way to the NVMe drives in the back end. (See Figure 2.)
- EF570 and E5700 systems support NVMe-oF from the host to the front end. The back end is still SCSI based with SAS drives. (See Figure 3.)

NetApp EF600, EF570, and E5700 systems support NVMe/IB, NVMe/RoCE v2, and NVMe over Fibre Channel (NVMe/FC) with benefits that include but are not limited to the following:

- E-Series and EF-Series storage systems already support FC as a transport layer for SCSI protocol commands. NVMe/FC adds a new protocol over this well-established transport layer.
- The same hardware that runs FC can run NVMe/FC, although not at the same time.
- Both protocols (FC and NVMe/FC) can coexist on the same fabric and even on the same FC host bus adapter (HBA) port on the host side. With this capability, your existing fabric that runs FC can connect your EF600, EF570, or E5700 system that runs NVMe/FC to the same fabric.
- IB and RoCE have RDMA built into them. E-Series and EF-Series systems already have a long history of supporting SCSI-based protocols over RDMA, such as iSCSI Extensions for RDMA (iSER) and SCSI RDMA Protocol (SRP).
- The same hardware on EF570 and E5700 systems that runs iSER or SRP can run NVMe/IB or NVMe/RoCE, although not at the same time.
- iSER, SRP, and NVMe/IB can coexist on the same fabric and even on the same IB host channel adapter (HCA) port on the host side. With this capability, your existing fabric that runs iSER or SRP can connect your EF600, EF570, or E5700 system that runs NVMe/IB to the same fabric.
- Both iSCSI and NVMe/RoCE can coexist on the same fabric on the host side.
- NetApp EF600, EF570, and E5700 systems support NVMe/RoCE v2 (which is routable), and they are also backward compatible with RoCE v1.
- Components in the fabric (NetApp EF600, EF570, and E5700 storage systems; switches; and HCAs) can negotiate the speed down as needed. A lower speed makes it easier to connect to legacy components.

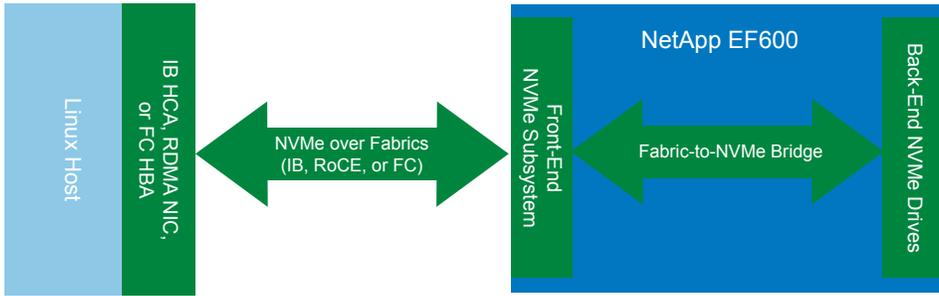


Figure 2) NVMe end to end on an EF600 system.

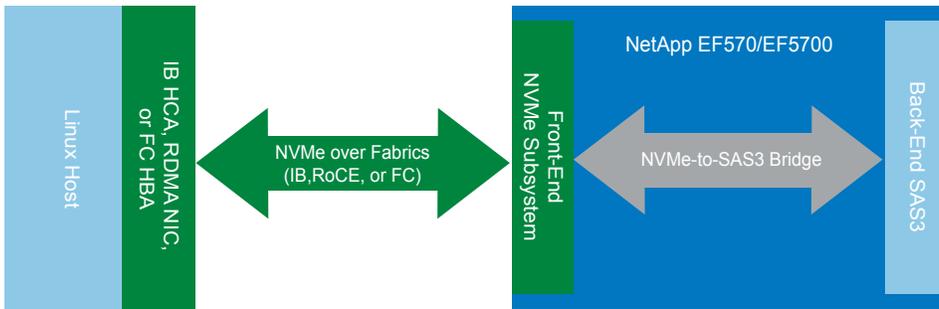


Figure 3) NVMe front end on EF570 and E5700 systems.

E-Series and EF-Series Interoperability with Third-Party Components

For a comprehensive list of supported components, see the NetApp E-Series and EF-Series Interoperability Matrix Tool (IMT). Following is an example:

- Operating systems: SUSE Linux Enterprise Server 12 SP4
- Host adapters: for NVMe/IB, Mellanox FDR and EDR HCAs; for NVMe/RoCE, host adapters from other vendors such as Cavium (Marvell); for NVMe/FC, Broadcom
- Switches: Mellanox, Cisco, and Broadcom
- Topology: fabric-attached and direct-attached

Coexistence Between NVMe/IB, iSER, and SRP

All three of these protocols can coexist on the same fabric and even on the same IB HCA port on the host side. Therefore, if your existing fabric runs iSER and SRP, you can connect an EF570, EF600, or E5700 system that runs NVMe/IB to the same fabric. Figure 4 shows an example.

Coexistence Between NVMe/RoCE and iSCSI

On the host side, iSCSI and NVMe/RoCE can run on the same network adapter (if it supports RDMA) at the same time, but they must be connected to different E-Series and EF-Series systems. Figure 5 shows an example.

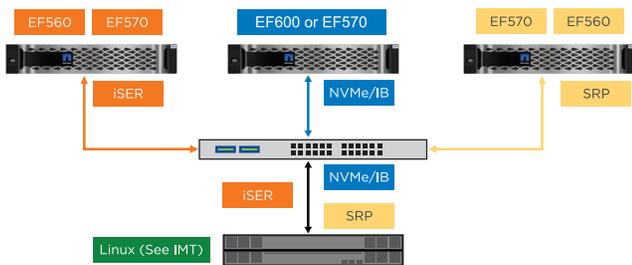


Figure 4) Coexistence between NVMe/IB, iSER, and SRP.

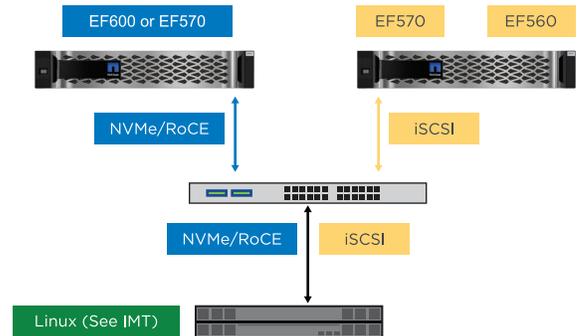


Figure 5) Coexistence between NVMe/RoCE and iSCSI.

Coexistence Between NVMe/FC and FC

On the host side, NVMe/FC and FC can run on the same HBA at the same time, but they must be connected to different E-Series and EF-Series systems. Figure 6 shows an example.

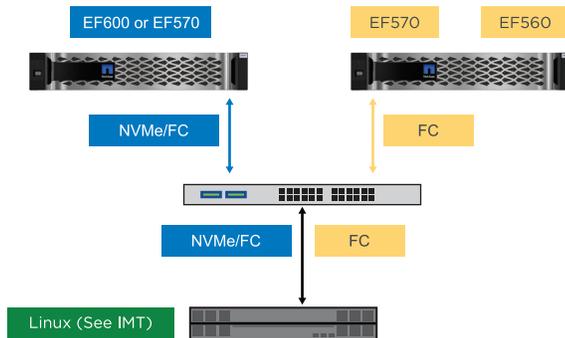


Figure 6) Coexistence between NVMe/FC and FC.

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. #DataDriven