First Look

Simplifying Data Center Infrastructure with Dell EMC and VMware End-to-end NVMe/TCP Solution

Date: September 2021 **Author:** Alex Arcilla, Senior Validation Analyst

Storage Infrastructure Challenges:[1]

74% The percentage of organizations that have deployed or are planning to deploy **on-premises NVMe-based solid-state (or flash) storage technology**.

70% The percentage of current and planned NVMe users that have deployed or are planning to deploy **on-premises NVMe over Fabrics** in storage area network fabrics.

End-users continue to demand high application performance to perform their tasks in today's rapidly changing business environment. Organizations initially addressed this issue by implementing NVMe SSD-based storage to mitigate storage performance as a potential bottleneck. However, scaling became an issue, as these performance gains required NVMe to be implemented as direct-attached storage (DAS) for every server host. Because NVMe acted as DAS, organizations potentially had unused storage that could not be easily shared by multiple hosts.

To achieve scalability of NVMe storage while maintaining high levels of performance, organizations began implementing NVMe over Fabrics (NVMe-oF) technology. NVMe-oF extends the NVMe protocol to transport protocols such as Fibre Channel (FC) and remote direct memory access (RDMA), facilitating faster connectivity between storage and servers. However, implementing NVMe-oF using such protocols can increase both network configuration costs and complexity.

As organizations seek to consolidate their data centers, leveraging an NVMe IP storage area network (SAN) provides a viable option. This approach not only enables a pool of NVMe-based storage to be shared amongst servers, but also does not require extensive network configuration to existing data center infrastructure, decreasing both capital and operational costs. More importantly, application performance is not compromised when compared to alternative NVMe-oF implementations.

Dell Technologies End-to-end NVMe/TCP Solution

To help organizations achieve the high application performance end-users demand without adding unnecessary capital and operational costs, Dell Technologies has released its end-to-end NVMe/TCP solution with VMware, comprised of:

- Dell EMC PowerEdge servers with VMware ESXi hypervisor.
- Dell EMC PowerStore, an all-NVMe storage array.
- VMware ESXi hosts on Dell EMC PowerEdge managed via VMware vCenter.
- Dell EMC SmartFabric Storage Software, a standards-based centralized discovery controller.

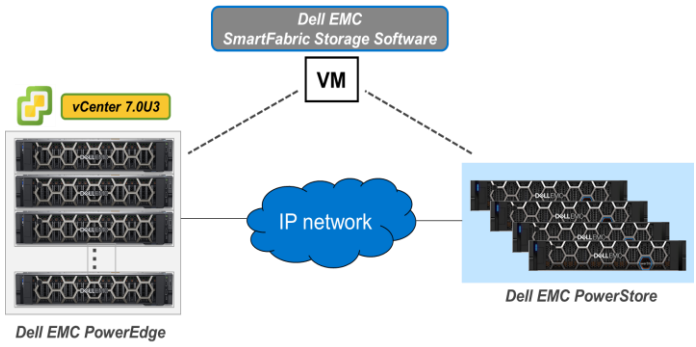

The solution is ideal for organizations scaling out their storage infrastructure cost-effectively, typically experienced when using SCSI, while achieving comparable performance experienced with storage area networks (SAN) using the FC protocol.

At the heart of the solution is Dell SmartFabric Storage Software, a centralized discovery controller that helps to automate NVMe/TCP end-to-end configuration, connectivity, and management of a server host and storage subsystem. The controller

[1] Source: ESG Master Survey Results, *2021 Data Infrastructure Trends*, September 2021.

is delivered via containerized services embedded within a virtual machine (VM) on an on-premises server. Dell EMC SmartFabric Storage Software implements NVM Express standards related to centralized discovery.[2]

What makes Dell EMC SmartFabric Storage Software attractive is its ability to discover hosts and storage devices automatically. Typically, discovering compute and storage devices in IP networks requires complex and manual configuration. Organizations can then only identify and configure storage subsystems independently of each other; the amount of manual work hinders organizations from scaling out an NVMe-oF infrastructure efficiently.

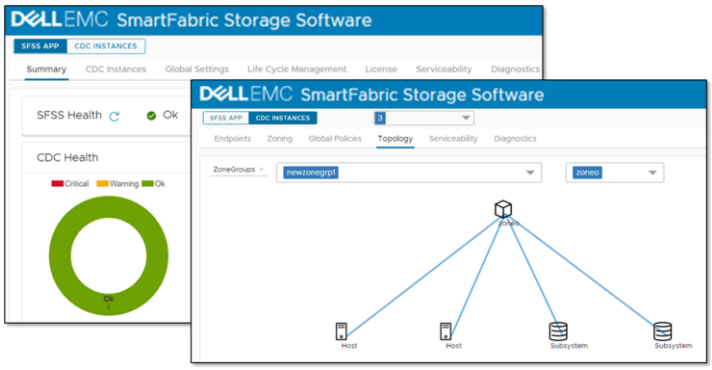

With Dell EMC SmartFabric Storage Software, organizations can automatically discover both hosts and subsystems added to an NVMe IP SAN. As host and storage subsystem interfaces become visible to Dell EMC SmartFabric Storage Software, the compute and storage devices self-register with the controller. Administrators can use the registered information to enable access of host interfaces to specific storage subsystem interfaces. Organizations can also automatically log fabric-generated events and notifications, such as topology changes, to keep resource visibility up to date. The amount of manual work to account for, register, and track changes of hosts and subsystems decreases when using Dell EMC SmartFabric Storage Software, leading to lower management costs.

ESG Highlights

Storage Performance Using NVMe/TCP Comparable to FC/NVMe and SCSI/FC

ESG first reviewed test results comparing NVMe/TCP performance with the performance typically seen in FC and iSCSI SANs. The test bed consisted of:

- Two PowerEdge R640 servers installed with VMware ESXi 7.0 Update 3 (7.0U3). Each server contained a dual port Dell QLogic 32G FC adapter and a dual port Mellanox CX4 25GbE adapter. One server ran a software initiator for NVME/TCP storage traffic on the Mellanox adapter, while the other server supported SCSI/FC and iSCSI storage traffic. A Dell EMC PowerStore (v2.1.0.0) acted as the storage subsystem for both servers.
- Two Cisco Nexus 9000 C93180YC-FX chassis connecting the servers and storage, used for both FC and Ethernet traffic. Interconnect between the switches was active.
- Windows 2019 as the guest OS on the ESXi hosts. Eight storage LUNs were mapped as VMFS datastores.

VDBench, an I/O workload generator, measured storage performance achieved with NVMe/TCP, compared with that of FC/NVMe and iSCSI. After generating a random read workload with a 4KB block size, increasing the queue depth from five to 20 threads (in increments of five threads), IOPS and response time were measured (see Figure 1, Figure 2, and Table 1).

ESG observed that I/O performance and latency achieved with the NVMe/TCP SAN fell between performance and latency observed with the FC and iSCSI SANs. Similar trends were observed when we increased the block size to 8KB. Based on the results, ESG noted how a NVMe/TCP SAN can combine the cost-effectiveness associated with SCSI without sacrificing the higher performance expected with FC. We also noted the advantages of using NVME/TCP as networks transition to higher speeds. Organizations with Ethernet networks that seek to implement NVMe/TCP will transition from 25 Gb/sec to 100

[2] The implemented NVM Express standards include TP 8009 – Automated Discovery of NVMe-oF Discovery Controllers and TP 8010 - Central Discovery Services. These have yet to be ratified as of the time of this publication.

Gb/sec, helping to improve overall performance. Currently, those opting for FC SANs can increase network speeds from 32 Gb/sec to 64 Gb/sec.[3]

Figure 1. IOPS with NVMe/TCP Comparable to FC and iSCSI

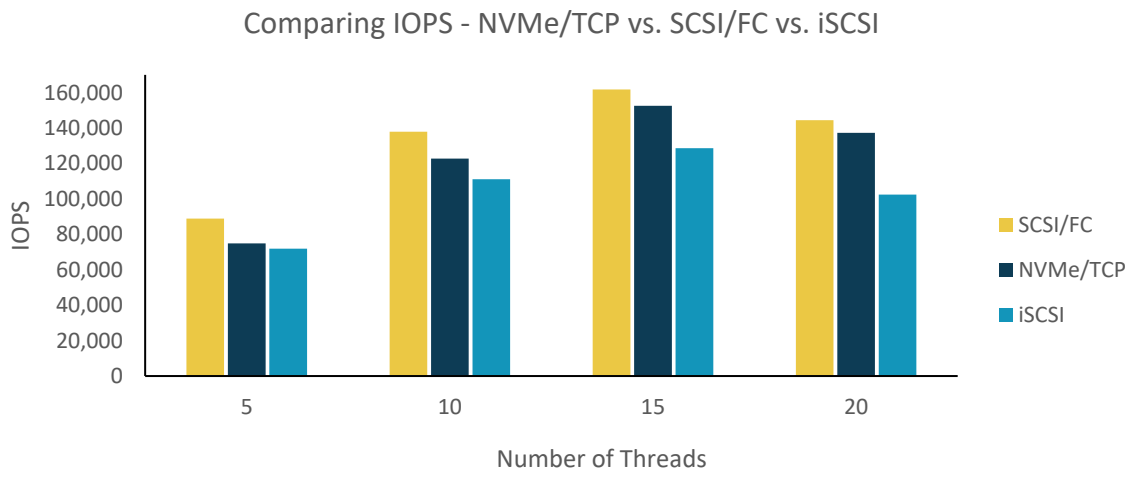

Source: Enterprise Strategy Group

Figure 2. Latency with NVMe/TCP Comparable to FC and iSCSI

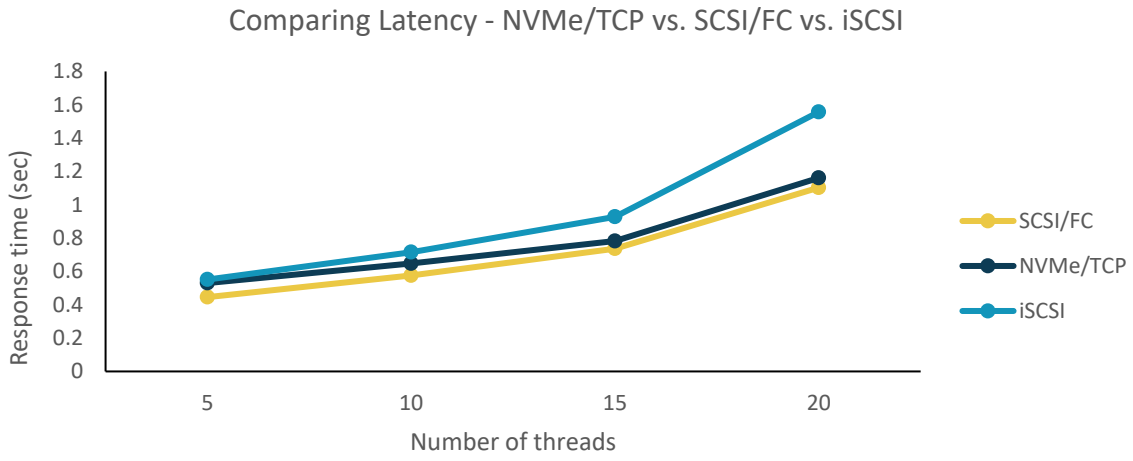

Source: Enterprise Strategy Group

Table 1. Results of Performance and Latency Tests

	SCSI/FC		NVMe/TCP		iSCSI	
# of threads	IOPS	Response time (sec)	IOPS	Response time (sec)	IOPS	Response time (sec)
5	88,825	0.4458	74,872	0.5302	71,872	0.5514
10	137,862	0.5758	122,729	0.6477	111,059	0.7155
15	161,737	0.7368	152,550	0.782	128,596	0.9278
20	144,373	1.1029	137,265	1.161	102,406	1.557

Source: Enterprise Strategy Group

[3] At the time of publication, the Fibre Channel Industry Association has roadmapped market availability of 128 Gb/sec for 2024.

Configuration of NVMe/TCP SANs with Dell EMC SmartFabric Storage Software

- ESG observed how easily a storage subsystem could be configured for a VMware ESXi host leveraging Dell EMC SmartFabric Storage Software. We navigated to VMware vCenter to configure a host's NVMe/TCP adapter with the required controllers so that the host could discover and connect to storage targets automatically. We noted that the Dell EMC SmartFabric Storage Software enabled the auto-discovery of controllers to be loaded onto the adapter.

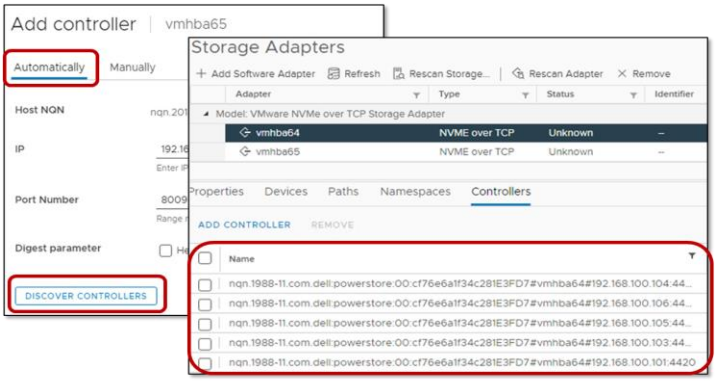

(auto-discovered controllers)

- We then navigated to the Dell EMC PowerStore management interface to add an ESXi host named "*VMworldDemo.*" Once we created this host, we simply specified an NVMe initiator to connect with an NVMe storage subsystem. We also created a 750GB storage volume named "*VMworldVolume*." Once the volume was created, the Dell EMC SmartFabric Storage Software enabled *VMworldDemo* to automatically connect to *VMworldVolume*, as the same volume appeared on vCenter associated with our ESXi host.

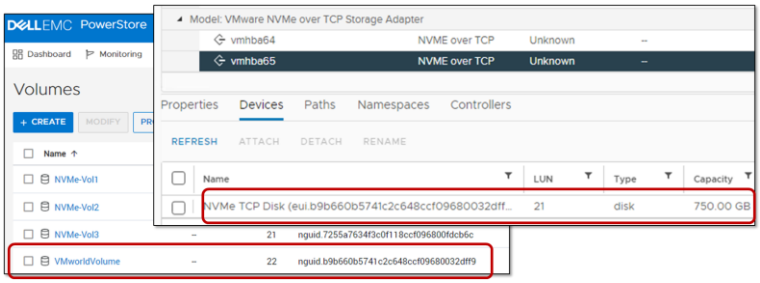

- Finally, ESG audited the results of an initial study comparing the capital expenses associated with implementing NVMe/TCP and SCSI/FC storage fabric technology. Costs of a 100GbE Ethernet switch and 32G FC switch, with similar switching capacities, were estimated. The Ethernet switch was configured with a dual port 25GbE NIC. A dual port 32G FC HBA was installed on the FC switch. Using publicly available information, estimates revealed that the cost per GB of available bandwidth on the Ethernet switch is at least 75% less than what was estimated for the FC switch.

First Impressions

Organizations using NVMe-oF technology can maximize storage performance and scalability, but implementations based on protocols such as FC and iSCSI actually lead to tradeoffs between incurring additional costs and maximizing performance. Dell Technologies end-to-end NVMe/TCP solution with VMware removes the need for that tradeoff so that organizations can minimize overall capital and operational expenses while achieving comparable performance typically seen with other SAN protocols, particularly FC.

ESG's first impressions of this end-to-end solution are that its implementation of the NVMe/TCP protocol can indeed achieve levels of performance comparable to other well-known protocols. We also observed how configuration of hosts and targets, using the Dell EMC SmartFabric Storage Software, does help in reducing overall time due to auto-discovery and automation of configuration tasks. Overall, ESG believes that Dell Technologies' end-to-end NVMe/TCP solution offers the best of both worlds: the cost-effectiveness of iSCSI with the higher performance that organizations expect with FC.