



NVMe Storage Evolution

Tsiyon Sadiky



Tsiyon Sadiky
Product Manager
Infinidat

Biography

Tsiyon Sadiky is the Product Manager, Infinidat (<https://www.infinidat.com>) with nearly 20 years' experience in storage product software development, pre-sales, customer support, and product management.

As Product Manager, Tsiyon works closely with Infinidat's enterprise customers and the technical development team, to ensure that product innovations continue to exceed customer expectations and provide the very best storage performance.

Tsiyon holds a bachelor's degree in Computer Science.

Tsiyon blogs at <https://www.infinidat.com/en/blog>

Keywords Storage, Infinidat, Non-volatile Memory Express (NVMe), Solid-state drives (SSDs), Data speed
Paper type Opinion

Abstract

As the evolution in storage continues, Non-volatile Memory Express (NVMe), the new storage access and transport protocol for flash and next-generation solid-state drives (SSDs) is delivering the highest throughput and fastest response times for all types of enterprise workloads. Input/Output tasks performed using NVMe drivers begin faster, transfer more data, and finish faster than older storage models using older drivers, such as AHCI (Advanced Host Controller Interface) a feature of SATA SSDs. In this article, the author explains how NVMe is becoming the new industry standard for both servers in the data centre and in client devices, and how Infinidat is supporting the technology.

Introduction

Evolution in storage continues to deliver means computing power is delivering highest throughput and fastest response times. At Infinidat we recently announced the extension of our support for Non-volatile Memory Express/Transmission Control Protocol (NVMe/TCP) to VMware environments with the release of InfiniBox 7.1. As part of the NVMe over fabrics family, our extension of NVMe/TCP for VMware provides customers with an easy and cost-effective way to get the most out of their existing investments in network infrastructure and our industry acclaimed enterprise storage platforms.

VMware customers using vSphere 7.0 update 3 can now benefit from NVMe/TCP and will continue to enjoy improvements to the protocol support with the upcoming vSphere 8 release. Together with the release of InfiniBox 7.1 and the



Data Centre and Virtualization

accompanying Host PowerTools for VMware (HPT-VM), VMware customers using InfiniBox get the same simple, intuitive and self-configuring experience for NVMe/TCP that they are familiar with when using iSCSI and Fibre Channel (FC).

Together with the newly added NVMe/TCP support for VMware, InfiniBox 7.1 offers support for all major Linux distributions, as NVMe/TCP is gaining momentum in that space as well. The same benefits exist with Linux and NVMe/TCP, from the cost of infrastructure to the performance gains, and we want all our enterprise customers to benefit from all our solution developments.

InfiniBox support for NVMe/TCP on VMware comes at no additional cost via a simple and non-disruptive software upgrade, and is available for all supported InfiniBox models, regardless of HW generation or capacity.



Protocol (R)Evolution

NVMe/TCP is sometimes referred to as an evolution, mostly of Internet Small Computer System Interface (iSCSI), and sometimes as a revolution. Both are true.

You can say it's an evolution because, at its essence, iSCSI and NVMe/TCP serve the same purpose: high speed and low latency access to remote centralized storage using block semantics. Both can use flexible provisioning based on application needs, while similarly enjoying data services on the storage array, such as snapshots, clones, and data reduction.

The revolution part comes from the fact the protocols had a fresh start. The clean slate starts with NVMe, the protocol used for direct attached storage, which aims to replace SCSI-based protocols, such as SAS. NVMe, compared to iSCSI, has more parallelism built into it to take advantage of the advancements in solid state drives. Those advancements were carried over to NVMe over Fabrics (NVMe-oF), and its specific derivatives such as NVMe/TCP which runs on standard Ethernet networks.



NVMe/TCP is particularly interesting when compared to its siblings running over other transports (FC and RoCEv2). Each variant has its advantages and disadvantages, but NVMe/TCP offers the lowest admission fee to NVMe-oF, particularly if you're replacing iSCSI:

- If you're looking to replace iSCSI with NVMe/TCP, your existing network infrastructure should suffice, your investment in new hardware should be minimal, and the upgrade should deliver clear value over iSCSI.
- If you're looking to replace FC with NVMe/FC, your existing network infrastructure should also suffice, but it's an infrastructure which is considered more expensive per machine.
- If you're looking to deploy a RoCE v2 network, you have the cost of RDMA enabled switches and RNICs.
- If you're looking to deploy an FC network, you have the cost of FC switches and HBAs.





NVMe/TCP on InfiniBox

NVMe/TCP for VMware extends the capabilities of the existing block and file protocols already supported by InfiniBox, while leveraging the existing InfiniBox architecture. These capabilities include:

- Triple redundant architecture for high availability and business continuity.
- Ease of management with fully functional APIs for autonomous automation, and clear GUI and CLI for simple manual tasks.
- Integration with VMware vCenter through Host PowerTools.
- All the enterprise features and data services from data reduction to replication. (Synchronous and asynchronous replication is already supported, but Active-Active replication will be available for volumes served by the NVMe/TCP protocol in a future InfiniBox release.)

Looking back at the inception of NVMe and the NVMe-oF family of protocols, performance was a major driver in their design, and we're happy to announce we delivered on that front as well. When compared to iSCSI with the same equipment, we've seen workloads with IOPS results 30% higher, while delivering 35% lower latency.

In conclusion

NVMe/TCP overall has very clear advantages in terms of simplicity, cost, and performance, and it is not surprising to see VMware enhancing their support for it. VMware with an InfiniBox enterprise storage platform leverages all the existing advantages of using InfiniBox today and provides protocol modernization as well as what should be a nice (and free!) performance bump.