

2023

ATTO Technology ExpressNVM™ NVMe Smart Switch Host Adapters and Xinnor Software RAID (xiRAID)

Data integrity, flexibility in system design, and management of storage resources are key requirements when designing a scalable high-performance storage solution.

Overview


In compute intense and high-performance ecosystems, RAID has been widely accepted for critical applications that include protection from drive failures, enabling data redundancy and fault tolerance to improve storage reliability and performance. These characteristics are fundamental to any enterprise and client system backbones to ensure a robust and resilient infrastructure.

NVMe SSD devices are increasingly becoming the standard for providing extremely low latency and high-performance block storage in the enterprise data center.

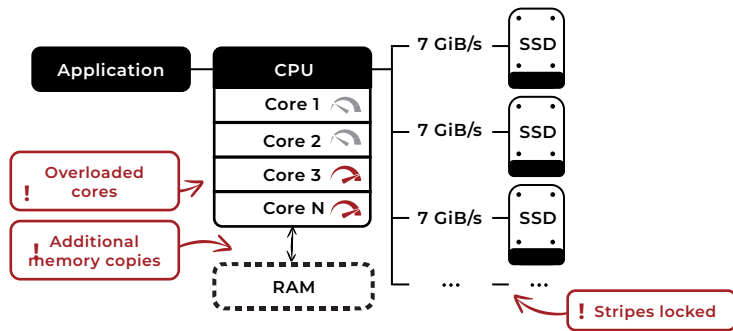
For mission-critical applications, low latency is important as customers expect to interact with technology in real-time with no delays, and access to fast storage in large capacities with reduced downtimes is extremely valuable to such IOPS intensive scenarios. Enterprises and data centers today, demand robust NVMe storage connectivity to deliver best-in-class end user experience.

Challenges

NVMe devices are not protected in any way and in that they bring with them new challenges when it comes to data integrity, management, and scalability. Addressing these challenges that are a hindrance to deploying RAID on NVMe SSDs is of great importance.



Software RAID options are designed to work with slower devices with less parallelism and I/O intensity, like HDDs or SATA and SAS solid-state drives. When used with NVMe, most software RAID products struggle with CPU load, memory usage and overall performance. This is magnified during RAID rebuilds in the case of drive failures.



These drawbacks are one of the main reasons why in many installations customers have to rely on simple mirroring or RAID 1+0, schemes to keep high levels of performance and data protection. However, this is costly due to the necessary doubling of needed drive capacity.

Server platforms today offer NVMe backplanes, but these are usually attached to simple PCIe retimers to boost signal strength. It solves signal quality issues but little else. In the storage world, a fitting analogy would be SATA with 1:1 connections and simple S.M.A.R.T. management.

Data intensive and mission critical applications require advanced scalability options, similar to SAS expanders and extensive device health monitoring and performance analytics for efficient management of storage resources.

Solution

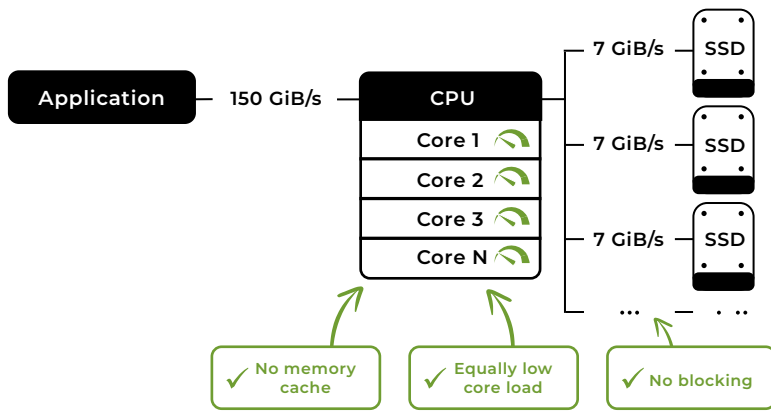
Xinnor xiRAID is a software RAID solution designed from the ground up for NVMe performance. It provides a high level of flexibility while keeping host resource utilization to an unprecedented minimum.

When dealing with hundreds of thousands or millions of IOPs and tens of gigabytes per second of throughput, software RAID starts using significant amounts of host resources, which can starve business applications of CPU cycles.

“ xiRAID relies on a relatively rarely used feature of modern x86 CPUs called AVX (Advanced Vector Extensions) and has a lockless architecture that helps spread computation evenly across CPU Cores. ”



This is a highly effective approach that needed years of research but allows RAID parity to be computed “on the fly” for tens of millions of IOPs with negligible effect on host resources. It also enables faster reconstruction of RAID in case of a drive failure, minimizing the window of increased risk of data loss and degraded performance.



ATTO ExpressNVM™ NVMe Smart Switch Host Adapters are an innovative storage connectivity solution that adds scalability and flexibility to dense NVMe storage ecosystems while also including advanced management capabilities not available on other NVMe adapters.

ExpressNVM is a PCIe Gen4.0 host adapter with dedicated connections for up to 16 NVMe SSDs scaled to add high-performance storage capacity in a host server. The intelligent hardware architecture onboard the ExpressNVM is tightly coupled with highly optimized algorithms to help deliver maximum throughput in the fastest, most reliable, and smartest way possible.

ExpressNVM comes with a host of management tools and administrative features that allow data centers and hyperscalers to design smart and flexible enterprise-grade systems while retaining the full performance of NVMe SSDs. The FPGA-based management engine on the ExpressNVM, monitors key metrics via SMBUS and translates them into highly useful information on the software stack, ATTOView, that collects and analyzes metric data. IT managers and system admins can now gain insights into the NVMe ecosystem at a granular level.

ExpressNVM also features UBM and Virtual SES capabilities that extend device management via the backplane.

In addition to keeping track of temperature and firmware revisions, metrics like drive health, faulty device isolation, and containment, drive read/write statistics, add incredible value to the NVMe ecosystem – helping reduce downtime and lower overall TCO.

“When used together, xiRAID’s exceptional software RAID performance and ExpressNVM’s ability to scale, manage, and monitor NVMe drive arrays combine to create an economical and ultra high-performance storage solution unlike any other on the market.”

It is a powerful yet highly flexible solution to the challenge of affordable, efficient data protection, and drive management for NVMe in the data center.

Testing

We installed 8 NVMe SSDs via PCIe Gen4 on a Dell PowerEdge server using the ATTO ExpressNVM Adapter and ran sequential I/O across the RAID group.

Sequential workloads were measured using FIO with a 128K IO size. xiRAID recorded an impressive RAID5 performance with reads clocking at 27.5 GB/s in a x4 PCIe configuration across the RAID group, and writes averaging at 9 GB/s.

Sequential IO (M&E)				
	S4FF		S48F	
	Read	Write	Read	Write
RAW	25.2 GB/s	10.8 GB/s	28.1 GB/s	11.1 GB/s
RAID	25.0 GB/s	9.081 GB/s	27.5 GB/s	9.3 GB/s

Test Setup:

Server: Dell Power Edge R750

CPU: Intel® Xeon® Gold 6330 CPU @2.00GHz

Memory: 128GB DRAM

Operating System: Oracle Linux 8.7

Kernel Version: Linux 5.4.17-2136.313.6.el8uek.x86_64

NVMe Adapter: ExpressNVM S48F/S4FF x16PCIe Gen4

NVMe SSDs: 8x Western Digital Ultrastar DC SN650 NVMe drives

RAID Level: RAID5

RAID stripe size: 128K – with Xinnor recommended optimization

Performance Benchmarking: FIO

Conclusion

Xinnor xiRAID and ATTO ExpressNVM bring the cost-effectiveness and familiarity of traditional RAID to the NVMe storage ecosystem while providing exceptional performance, reliability and scalability. xiRAID enables full-featured NVMe software RAID with minimal resource penalties while ExpressNVM provides the flexibility and manageability that storage infrastructure architects demand for their RAID designs.

Get In Touch



+1.716.691.1999



www.ATTO.com



ATTO Technology, Inc.

155 CrossPoint Parkway

Amherst, New York 14068

XINNOR

+972 43 740 203

request@xinnor.io

XINNOR

Bar Jehuda road 300

«Lev ha'Zomet», Check Post, Haifa, Israel