

ADOPTING THE RIGHT NVME SOLUTION FOR YOUR APPLICATION

INTRODUCTION

There is an insatiable need for faster and higher capacity solid-state storage in high-performance environments. The non-volatile memory express (NVMe) is a communication protocol developed specifically for systems that utilize PCIe-based solid-state drives. It provides a much faster and more efficient alternative to legacy standards for interfacing between host systems and peripheral devices. These characteristics combined with the emerging M.2 small form factor is the ideal solution for next generation computing and storage systems. The performance benefits of PCIe along with the decrease in cost per bit continues to drive adoption of NVMe solutions.

Some of today's high-end systems can accommodate M.2 NVMe SSDs right on the motherboard. However, for computing environments that require high speed storage, system architects are considering expansion adapters to utilize four M.2 NVMe SSDs per PCIe slot. This would provide up to 8TB of high-performance storage, with future developments enabling even higher capacity.

DESIGN FOR QUALITY

Aplicata's family of M.2 NVMe SSD [adapters](#) are optimized to meet the needs of high-performance applications and systems. These adapters offer flexible configurations and choosing the right adapter will provide optimal results.

Since Aplicata adapters support any NVMe drive from any manufacturer, they work out-of-the-box in the major modern operating systems and hypervisors. The standard disk management utilities that come with the operating system enable users to view and manage the drives installed on the adapter and the partitions associated with those drives. These utilities are also used to create different RAID configurations. The RAID functionality is handled by software on the processor, hence, it is called software RAID. The most common configuration is RAID 0 (or striping) that is used to provide higher read and write operations with a negligible impact on processor performance.

Aplicata products are built using a rigorous design process, advanced material selection, and an emphasis on thermal management which impacts product reliability. The big caveat to M.2 NVMe drives, especially datacenter versions, is that they consume high power even when idle. By design, NVMe controllers throttle to prevent an overheat from causing damage to the drives but that defeats the purpose of using NVMe technology in the first place. Reliable heat dissipation is an absolute must for M.2 NVMe SSDs to give the best performance.

Aplicata adapters feature a passive heatsink that carries away the heat generated by NVMe controllers into moving air inside the server enclosure. For better heat dissipation, the adapter is composed of high thermal conductivity materials along with multiple heavy copper planes, which allow the heat to be carried away from the bottom of the drive through the adapter. The heatsink is integrated tightly with the adapter and the drives, using high conductivity thermal pads. These pads come in a variety of thicknesses, so it is possible to accommodate various drives. A 2TB drive, for example, has components on the bottom and that requires the use of thinner pads between the drive and the adapter.

Overall, advanced design techniques and materials employed in the making of Aplicata adapters produce more reliable and higher quality products.

	Quattro 400	Quattro 410	Quattro 200
Form Factor	FHHL PCIe	FHHL PCIe	Single XMC
Maximum Drives	4	4	2
NVMe Protocol	Yes	Yes	Yes
Host Interface	PCIe Gen 3	PCIe Gen 3	PCIe Gen 2 or Gen 3
Bus Width	x8	x16	x4 or x8
M.2 Format	2280 or 22110	2280 or 22110	2280
Bifurcation Required	No	Yes	No
PCIe Switch	Yes	No	Yes
Thermal Management	Passive	Passive	Passive

TYPICAL APPLICATIONS

Aplicata makes M.2 NVMe SSD [adapters](#) for high-performance computing and storage systems. Below, two key applications are highlighted to illustrate the versatility of these adapters. By extension, other applications that require very fast access to data or flexible storage configurations would benefit from deploying Aplicata adapters.

100GE Data Transfer Nodes

A data transfer node (DTN) is dedicated server built with high-quality components and configured specifically for wide area data transfer usually within a scientific infrastructure. It provides faculty and researchers a convenient and effective way to transfer data files between HPC resources. Achieving 100 Gbps (or 12.5 GB/s) of disk IO to local storage is challenging and will require several PCIe Gen 3 x4 drives. It is recommended to use at least two Aplicata adapters per DTN, where each adapter is configured in RAID 0 to achieve higher sequential read and write operations.

ZFS Caching

ZFS offers various advantages in comparison to regular filesystems. It ensures data will be safe and uses very effective write and read caching techniques. For write operations, ZFS uses a filesystem journal, called ZFS intent log (ZIL), for storing synchronous writes until they are committed to disk. Maintaining this journal in a dedicated SSD, keeps data safe in case of a system crash. For read operations, ZFS uses a very fast cache located in the server's memory. This is called adaptive replacement cache (ARC). When the option of adding more memory becomes either physically or economically impossible, a level 2 ARC (L2ARC) implemented in SSD can increase the cache size dramatically. Ultimately, Aplicata's quad SSD adapters are ideal for ZFS caching. In a perfect configuration, ZIL is implemented by mirroring two drives to protect the journal integrity, while L2ARC is implemented by stripping the other two drives for fast read access. This configuration would only use a single PCIe expansion slot, so the remaining slots can be available for other functions.

COMMITMENT TO EXCELLENCE

At Aplicata, we continue to strive for excellence in innovation, quality, and customer service. The right NVMe solution must meet customer expectations and needs. Our knowledgeable team is committed to supporting customers and helping them benefit the most from choosing Aplicata. Please contact us to learn more about our products and services.