

NVMe SSDs: How virtualized platforms can benefit from innovative setup of high performance and large capacity NVMe storage arrays

Industry trends

Over recent years, the IT industry has continued its long-term shift from mechanical hard disk drives (HDDs) to solid-state drives (SSDs) - taking advantage of the speed and performance that the latter offers. According to IDC^[1], worldwide SSD industry revenues will continue to increase strongly between now and 2026, experiencing a compound annual growth rate (CAGR) of around 9.6% during this period.

The speed and performance requirements of today's applications continue to rise, whether it is a new game, a new application or a more demanding business service. Many customers are looking for faster storage and others are focused on lower power consumption, while latency needs to get shorter, and there is a continual struggle for hardware to keep up with software developments.

To meet increasing demands, data centres and enterprise operations are now also looking to move away from SSDs based on older SATA and SAS technologies, which were originally developed for HDDs. Instead, they want to take advantage of the newer NVMe (Non-Volatile Memory Express) technology, which is specifically optimised for SSDs usage and delivers lower latency than other legacy storage technologies.

While NVMe SSDs provide substantial operational advantages, there are many companies that have not yet made the transition from AHCI. This is partly due to a perception that changing from SATA-based servers and storage to NVMe involves substantial extra costs in changing and adapting backplanes, cables and other peripherals. However, recent price drops for NVMe SSDs mean that they are now competitively positioned against their SATA SSD equivalents - and this is encouraging migration to the newer technology.

In fact, enterprise and data centre NVMe SSDs are projected to represent the majority of data storage devices across data centres globally, when compared to SATA and SAS SSDs. The combined segment (in units) is forecast to grow from 42.5% in 2019 to 91% by the end of 2023^[2].

NVMe technology in a system

The key to unleashing the performance of NAND Flash memory in server and storage systems is to free it from any limitations and burdens. Other than legacy protocols such as the previously discussed AHCI, the NVMe protocol is tailored to the needs of Flash memory-based storage components. The protocol was designed from the ground up as a light and efficient standard to support the high performance requirements of forthcoming IT deployments in cloud datacentres and enterprises.

As SATA and SAS controllers introduce latency on the path to storage components, directly connecting SSDs to the CPU via the PCIe interface is the logical alternative.

The PCIe interface offers a future-proof specification roadmap in line with upcoming data transfer standards. In this combination, SSDs ensure a fast, low-latency and non-volatile interconnection of relevant data to the applications.

NVMe and RAID

As well as considering individual disks, we also need to think about the employment of redundant array of independent disks (RAID) technology. This is commonly used across the IT industry, especially in enterprise applications, with multiple drives working together to ensure the resilience of critical data should one drive fail. However, with a conventional hardware RAID set-up there is a performance restriction in terms of increased latency, because of the RAID controller that handles the communication between the host and the SSDs.

There is an alternative to this, which is known as 'software RAID'. Here the CPU handles the controller task, but that uses a lot of CPU resources. By using software RAID the CPU may be unable to handle other workloads at the level required. In practice, as the CPU is one of the most costly components in a server, this may mean that a more expensive CPU may be required because of software RAID. There may be additional cost penalties in relation to enterprise software which is licensed on a per-core basis.

Instead, an increasingly popular approach is to use a dedicated card to handle the RAID control tasks, while eliminating the traditional bottleneck and taking advantage of the speed and capacity of NVMe. The SupremeRAID™ card by Graid Technology is an excellent fit for this task, and can support up to 32 NVMe SSDs - bringing the high performance and increased capacities of NVMe to a server.




Figure 1: SupremeRAID™ by Graid Technology enables high performance with low CPU load.

There is no bottleneck or compromise in performance, as the NVMe SSDs will connect directly to a PCIe bus (as shown in Figure 1). A single SupremeRAID™ card is able to support 19 M IOPS and 110 GB/s operation.

The NVMe RAID card can take over the responsibility of managing the storage, and release the CPU for other important tasks. The RAID card also enables a less complicated server setup - the card can simply be plugged into any one of the PCIe slots without extra cables to connect with NVMe SSDs, thus reducing the total cost of ownership (TCO) when building a server.

Testing activities

To demonstrate the performance of NVMe and the SupremeRAID™ card, the KIOXIA team has cooperated with primeLine Solutions and implemented a test system that would simulate real-world conditions of RAID storage.

The hardware system for the test was as follows:

- 1x ASUS RS520A-E11-RS24U | BIOS 0901-C
- 1x CPU AMD EPYC | SP3 | 7543P | 2.8GHz | 32-Core
- 8x MEMS DDR4 3200 ECC Reg. 64GB
- 1x Raid GRAID SupremeRAID™ SR-1010
- 1x NVMe Standard | M.2 | KIOXIA XG6 | 1024GB
- 20x Data Center NVMe | 2.5" | KIOXIA CD6-R | 7.68TB FW 0106

The operating system was Ubuntu 20.04, running the FIO 3.16 performance benchmarking tool. The RAID modes tested were RAID 10, RAID 5, and RAID 6 - representing the typical modes that are most commonly used in real-world applications.

The SSD used was the CD6-R from KIOXIA. This NVMe SSD is optimised for read-intensive data centre scale-out and cloud applications, including big data/IoT, online transaction processing and virtualization. Storage capacities of up to 15.36 TB are available in a 2.5-inch form factor, with active power consumption of 13 to 19 W. The series is compliant with both the PCIe 4.0 and NVMe 1.4 specifications. SSDs in the series deliver consistent performance up to 1 M IOPS (random read) and 85 K IOPS (random write).

The SupremeRAID™ SR-1010 by Graid Technology supports RAID levels 0/1/10/5/6/JBOD while the core software license supports up to 32 native NVMe drives. Designed for both Linux and Windows operating systems, it delivers superior performance while increasing read/write in all flash array and HPC applications.



Figure 2 and 3: CD6-R Data center NVMe SSD and SupremeRAID™ SR-1010 by Graid Technology

The test results obtained

The results of the testing work conducted are detailed in Figure 4 (showing sequential reads and writes) and Figure 5 (showing random reads and writes). The tests were run for a period of 30 minutes.

Figure 4 shows that the testing achieved very high performance for both reading and writing, which can be considered a maximum level that might be reached in a real-world scenario using the SupremeRAID™ cards together with SSDs.

Note that the KIOXIA tests used only 20 SSD drives with the SupremeRAID™ card, while it has a maximum capacity of 32 SSDs. This means that higher performance could be obtained by adding extra drives (up to the limit of 32).

As shown in Figure 5, when we test with random reads and writes, the performance figures are lower than for the optimal use cases detailed in Figure 4, but still achieve excellent results. The performance varies depending on the RAID mode used, and the various combinations of reads and writes.

In practice, many applications have highly predictable workloads, so the right SSD and RAID architecture can be chosen to match their demands. For example, database applications benefit from low latency, dual ports for multi-path and high availability, plus the highest available transactions per second. Alternatively, data analytics searches benefit from high bandwidth for sequential and random reads. Similarly, in data centres, cloud compute applications need high performance and low latency, while content delivery networks are best equipped with SSDs that excel at read-intensive workloads (typically 90% read and 10% write). Finally, media streaming benefits from very high read bandwidth.

The system set-up described in this whitepaper is suitable for applications including (but not limited to) machine learning, cache servers, streaming, simulation, virtualization platforms and cloud solutions.




Figure 4: RAID performance, sequential reads and writes




Figure 5: RAID performance, random reads and writes

Conclusions

As we have discussed in this white paper, NVMe SSDs provide a cost-effective, scalable and flexible way for enterprises and data centres to achieve the highest levels of performance. They present a marked operational boost over legacy designs, while being competitively priced. Furthermore, they are a future-proof technology that will continue to adapt to the latest application requirements.

With KIOXIA's latest generation of SSDs, working in conjunction with SupremeRAID™ cards by Graid Technology, customers can take data storage into a whole new era by exploiting the potential performance of their NVMe drives to the maximum – maintaining excellent throughput, while freeing-up CPU resources in a system that is simple to install.

For businesses that are considering a migration to NVMe SSDs, there has never been a better time to make the switch. NVMe SSDs have become the default standard in new platform and system designs, and SupremeRAID™ by Graid Technology has increased the appeal of NVMe SSDs even further.

Sources:

[1] Source: IDC. "Worldwide Solid State Drive Forecast Update, 2022-2026", Jeff Janukowicz, September 2022, <https://www.idc.com/getdoc.jsp?containerId=US47831722>

[2] Source: IDC. "Worldwide Solid State Drive Forecast Update, 2019-2023", Market Forecast Table 12, Jeff Janukowicz, December 2019, IDC #44492119.

Notes:

*The following trademarks, service and/or company names - IDC, SATA, SAS, NVMe, PCIe, AHCI, RAID, ASUS, AMD EPYC, Graid Technology, primeLine Solutions, SupremeRAID, Ubuntu, PCIe 4.0, HPC - are not applied, registered, created and/or owned by KIOXIA Europe GmbH or by affiliated KIOXIA group companies. However, they may be applied, registered, created and/or owned by third parties in various jurisdictions and, therefore, protected against unauthorised use. All other company names, product names and service names may be trademarks of their respective companies. All company names, product names and service names may be trademarks of their respective companies.

About KIOXIA Europe GmbH

KIOXIA Europe GmbH (formerly Toshiba Memory Europe) is the European-based subsidiary of KIOXIA Corporation, a leading worldwide supplier of flash memory and solid state drives (SSDs). From the invention of NAND flash memory in 1987 to today's renowned BiCS FLASH™ 3D flash memory technology, KIOXIA continues to pioneer innovative memory solutions with high quality and reliability. The company's innovative BiCS FLASH™ memory is an important component in almost all electronic devices where data need to be stored. By evolving "memory", KIOXIA creates uplifting experiences and changes the world. For more information, please visit [KIOXIA Website](#)

About primeLine Solutions GmbH

Headquartered in Bad Oeynhausen NRW, primeLine Solutions GmbH focuses on solution-based manufacturing and sales of server, storage, and industrial systems. In addition to established product lines, primeLine also manufactures individually according to customer requirements. With many years of experience in the IT market, the company sees itself as a reliable, flexible manufacturer and partner with an emphasis on uncomplicated service and support. Customers are data centers, research institutions, public institutions, SMEs, system houses and the retail trade. Learn more at www.primeline-solutions.com/de/.

About Graid Technology

Graid Technology is headquartered in Silicon Valley, California with an office in Ontario, CA, and an R&D center in Taipei, Taiwan. Graid technology's SupremeRAID™ performance is breaking world records as the first NVMe and NVMeoF RAID card to unlock the full potential of your SSD performance: a single SupremeRAID™ card delivers 19 million IOPS and 110GB/s of throughput. For more information on Graid Technology, visit graidtech.com or connect with us on [Twitter](#) or [LinkedIn](#).