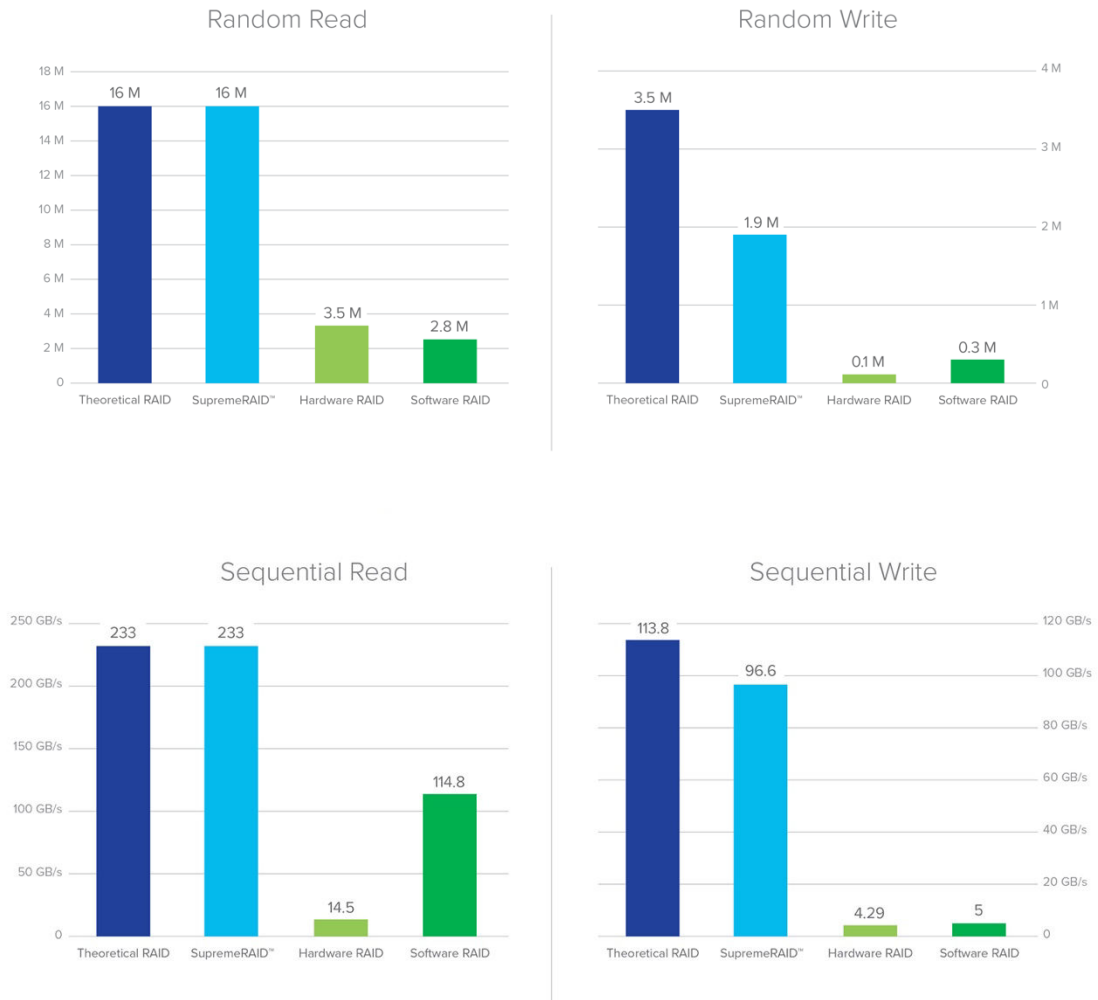


Performance Comparison of RAID 5 Solutions for PCIe Gen 5 NVMe SSDs

SupremeRAID™ vs. Software RAID vs. Hardware RAID

Executive Summary

SupremeRAID™ by Graid Technology offers an alternative option for adding RAID data protection to high-performance SSDs installed in servers. A primary benefit of choosing SupremeRAID™ over hardware and software RAID is preserving and using more of the available SSD performance. Extremely high efficiency allows RAID 5 performance to remain at or near theoretical levels (shown as dark blue bars in the chart below) when using SupremeRAID™. However, up to 97% of total SSD performance is wasted and unavailable when using hardware or software RAID 5.



Background

The fifth generation of the PCIe (Peripheral Component Interconnect Express) interface is known as PCIe Gen 5 (Generation 5). It supports maximum data transfer rates that are 2X faster than PCIe Gen 4 and 4X faster than PCI Gen 3. There is high interest in leveraging this higher performance because it enables servers to handle more data simultaneously, leading to improved system performance and faster data processing.

Other benefits of PCI Gen 5 include capabilities that improve the reliability and stability of the PCIe interface, reduce downtime, and ensure servers remain operational even in high-demand environments. Examples include new features like Forward Error Correction and Lane Margining.

PCIe Gen 5 Bandwidth

PCIe Gen 5 significantly improves server and solid-state drive (SSD) performance compared to previous generations, especially when combined with NVMe (non-volatile memory express) protocols. Together, they can achieve read and write speeds of up to 14 gigabytes per second (GB/s) or more. PCIe Gen 5 significantly improves over PCIe Gen 4 at up to 7 GB/s and PCIe Gen 3 at up to 3.5 GB/s.

	PCIe Gen 5	PCIe Gen 4	PCIe Gen 3
Maximum Bandwidth *	14 GB/s	7 GB/s	3.5 GB/s

* Typical performance with NVMe SSDs.

The increased performance of PCIe Gen 5 can be particularly beneficial for applications that require high bandwidth data access, such as high-performance computing, big data analytics, and artificial intelligence.

RAID Options for PCIe Gen 5 NVMe SSDs

Historically, there have been two options for adding RAID data protection to SSDs installed in servers. One option adds RAID using software running on the system's CPU (central processing unit), and the other adds RAID using hardware running on custom add-in cards with dedicated processors. A third option, SupremeRAID™, adds RAID using software running on standard GPU (graphical processing unit) add-in cards.

SupremeRAID™ features a unique off-data path architecture and utilizes GPU computing power for RAID data processing. It unlocks the full potential of PCI Gen 5 NVMe, especially when using RAID 5 data protection. Choosing SupremeRAID™ allows customers purchasing PCIe Gen 5 systems to experience the return on investment (ROI) they expected when investing in matching NVMe SSDs.

The following shows the results of SupremeRAID™ performance testing with a PCIe Gen 5 server and multiple PCIe Gen 5 NVMe SSDs and software RAID 5, hardware RAID 5, and SupremeRAID™ RAID 5.

Benchmark Testing

Hardware Components

The hardware selected for benchmarking consists of commercially available items, including one ASUS RS520A-E12-RS24U server, one AMD EPYC 9334 32-Core processor, twelve Micron MTC40F2046S1RC48BA1 DRAM, and sixteen Kioxia CM7 SSDs.

System Baseline Performance

Performance testing establishes the performance of a system with 16 PCIe Gen 5 NVMe SSDs. The results for sequential reads and writes match the aggregate performance of sixteen SSDs. However, performance for random reads and writes is lower than the aggregate performance of sixteen SSDs due to the impact of latency when using small IO sizes.

IO Type (8 jobs, 64 depth)	Tested System Performance*
Sequential Read 128KB	231 GB/s
Sequential Write 128KB	112 GB/s
Random Read 4KB	16 M IOPS
Random Write 4KB	7 M IOPS

* Performance without RAID data protection.

Performance Results

Hardware RAID 5 Performance Results

Testing suggests hardware RAID is too slow to manage two or more PCIe Gen 5 SSDs without significant performance losses. Choosing hardware RAID introduces unavoidable performance bottlenecks because SSDs connect to an add-in card. It generally involves connecting multiple SSDs with 4-lane PCIe interfaces to one add-in card with an 8- or 16-lane PCIe interface. Typical hardware RAID add-in cards support PCIe Gen 4, further bottlenecking performance.

IO Type (8 jobs, 64 depth)	Theoretical RAID Performance	Hardware RAID Performance	Hardware RAID Efficiency*
Sequential Read 128KB	233 GB/s	14.5 GB/s	6%
Sequential Write 128KB	113.8 GB/s	4.29 GB/s	4%
Random Read 4KB	16 M IOPS	3.5 M IOPs	22%
Random Write 4KB	3.5 M IOPS	108 k IOPs	3%

* Higher is better, and 100% is the maximum.

As shown in Figure 1, storage performance is negatively impacted by hardware RAID because reads and writes flow through the processor on the add-in card. Bandwidth is limited to the connection speed between the add-in card and the server. Compromised bandwidth can be particularly problematic in high-demand environments, such as data centers, where even minor delays can significantly impact overall system performance.

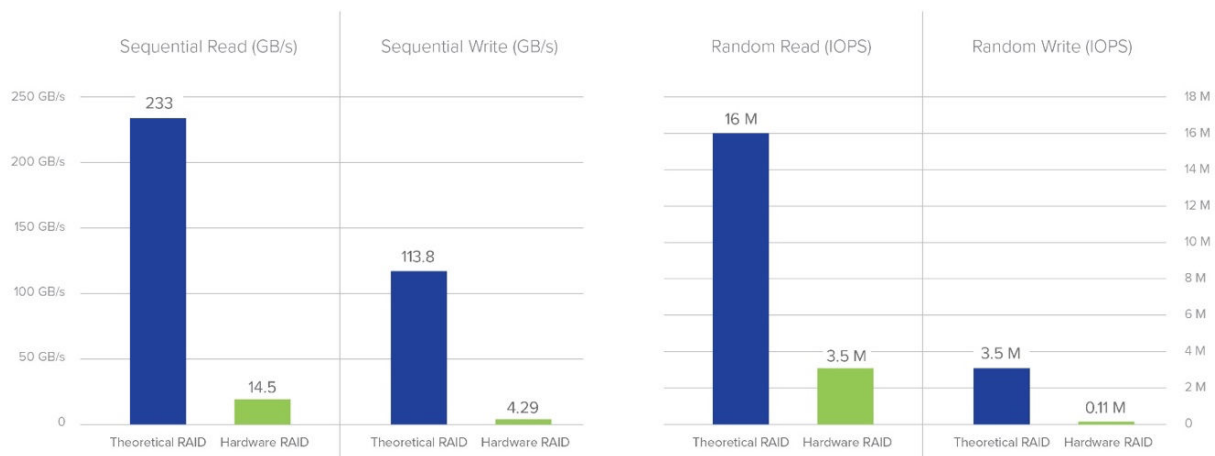


Figure 1: Hardware RAID Performance vs. Theoretical RAID

Software RAID 5 Performance Results

Testing suggests software RAID (e.g., Linux mdadm/mdraid) is too slow to manage two or more PCIe Gen 5 SSDs without significant performance losses. One exception is sequential reads that bottleneck at a higher number of SSDs, even when using servers with many CPU cores.

Choosing software RAID avoids using add-in cards but results in running RAID software and workloads alongside hypervisors, operating systems, databases, applications, and services. In other words, the performance of server workloads like databases and applications may suffer during periods of heavy storage activity.

IO Type (8 jobs, 64 depth)	Theoretical RAID Performance	Software RAID Performance	Software RAID Efficiency*
Sequential Read 128KB	233 GB/s	114.8 GB/s	49%
Sequential Write 128KB	113.8 GB/s	5 GB/s	3%
Random Read 4KB	16 M IOPS	2.8 M IOPS	17%
Random Write 4KB	3.5 M IOPS	355 k IOPS	10%

* Higher is better, and 100% is the maximum.

As shown in Figure 2, storage performance is negatively impacted by software RAID because reads and writes flow through the processor cores of the server. Bandwidth and IOPS are much slower than the aggregate speed of the SSDs, even though the SSDs remain connected to the server.




Figure 2: Software RAID Performance vs. Theoretical RAID

SupremeRAID™ RAID 5 Performance Results

Now let's benchmark the latest version of SupremeRAID™ (software v1.3.1) with SSDs remaining connected to the server. Sequential read performance matches the system baseline due to the SupremeRAID™ unique off-data path architecture. Sequential write performance,

however, requires extra parity data to be written, resulting in a speed of 90GiB/s. Random read performance can match the system baseline by leveraging the computing power of GPUs. Meanwhile, 4k random write performance reaches 1.9M IOPS due to the read-copy-on-write mechanism.

IO Type (8 jobs, 64 depth)	Theoretical RAID Performance	SupremeRAID™ Performance	SupremeRAID™ Efficiency*
Sequential Read 128KB	233 GB/s	233 GB/s	100%
Sequential Write 128KB	113.8 GB/s	96.6 GB/s	85%
Random Read 4KB	16 M IOPS	16 M IOPS	100%
Random Write 4KB	3.5 M IOPS	1.9 M IOPS	54%

* Higher is better, and 100% is the maximum.

As shown in Figure 3, SupremeRAID™ is highly effective at preserving the performance of PCIe Gen 5 NVMe SSDs. Theoretical and actual performance is similar for large sequential reads, large sequential writes and small random reads. While performance differs more broadly for small random writes, SupremeRAID™ still achieves about half of the theoretical performance.

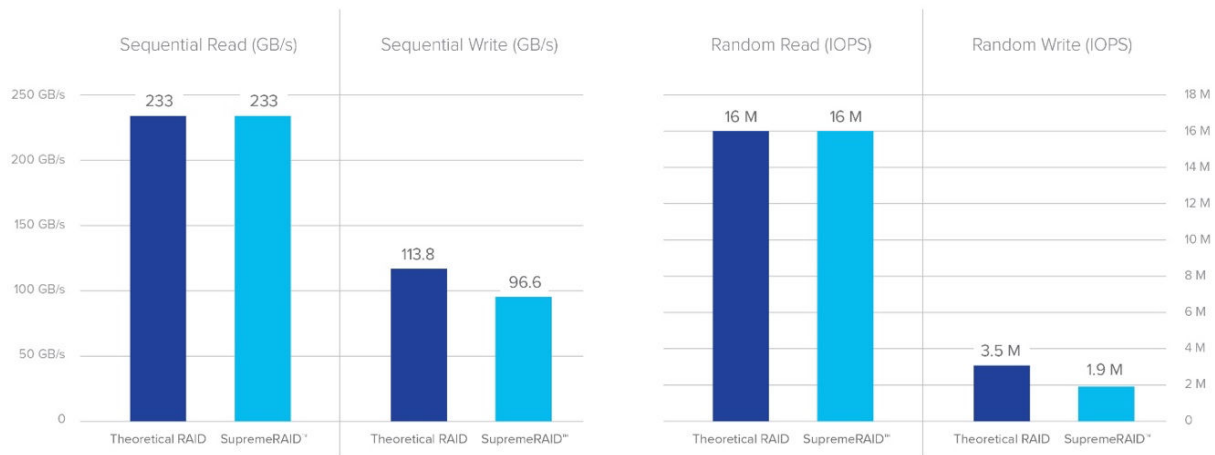


Figure 3: SupremeRAID™ Performance Results vs. Theoretical RAID

Conclusion

PCIe Gen 5 promises significant performance benefits for servers by enabling faster data transfer rates, improved I/O performance, increased lane density, better power efficiency, and improved reliability. Additionally, PCIe Gen 5 represents a significant advancement in SSD technology by offering improved performance, reliability, and efficiency. These benefits combine to help improve overall system performance, reduce operating costs, and ensure that data centers can handle the demands of modern computing applications.

Yet, adding data protection to PCIe Gen 5 NVMe SSDs requires an alternative to the unavoidable performance bottlenecks of using hardware or software RAID. With SupremeRAID™, customers can enjoy the benefits of PCIe Gen 5 and enterprise-grade data protection without sacrificing SSD performance.

About Graid Technology

Graid Technology, creator of SupremeRAID™ next-generation GPU-based RAID, is led by a team of experts in the storage industry and is headquartered in Silicon Valley, California with an R&D center in Taipei, Taiwan. Designed for performance-demanding workloads, SupremeRAID™ is the world's fastest NVMe and NVMeoF RAID solution for PCIe Gen 3, 4, and 5. A single SupremeRAID™ card delivers up to 19M IOPS and 110GB/s and supports up to 32 native NVMe drives, delivering superior NVMe/NVMeoF performance while increasing scalability, improving flexibility, and lowering TCO. For more information on Graid Technology, visit graidtech.com or connect with us on [Twitter](#) or [LinkedIn](#).