Protecting AI Data: Why GPUs Change the Game -

Breakthrough parity RAID performance for massive parallel AI I/O with SupremeRAID™ Ultra and InnoGrit N3X

April 2026

Table of Content

EXECUTIVE SUMMARY 2

THE HARDWARE ADVANTAGE: SUPREMERAID™ 2.0, NVIDIA RTX 2000E ADA, AND INNOGRIT N3X 2

WHAT AI WORKLOADS DEMAND FROM STORAGE 3

PERFORMANCE COMPARISONS 4

Test Descriptions 4
Testing Environment 4
1. 4K Random Read Performance 5
2. 4K Random Write Performance 5
3. 1M Random Read Performance 6
4. 1M Random Write Performance 7
CPU Efficiency Analysis (RAID5 Optimal) 8
Read Efficiency: Small & Large Block 8
Write Efficiency: Small & Large Block 9

CONCLUSION 10

APPENDIX 11

Executive Summary

AI workloads are massive, highly parallel, and unforgiving. They generate intense bursts of small-block and mixed I/O patterns across thousands of threads, queues, and datasets. At scale, this pressure turns traditional storage controllers into a primary bottleneck for training, inference, and data preparation.

This joint whitepaper demonstrates how SupremeRAID™ 2.0 redefines parity RAID performance for the AI era. By pairing 24x InnoGrit N3X SLC NVMe drives with the SupremeRAID™ Ultra (powered by a 50W NVIDIA RTX 2000E Ada), we eliminate the traditional trade-off between data protection and performance.

The results focus on the metrics that define AI infrastructure efficiency. SupremeRAID™ delivers millions of random write operations per second with parity RAID in optimal mode and sustains unmatched throughput in degraded mode—scenarios where traditional software RAID often collapses. The result is faster data ingestion, higher metadata responsiveness, and resilient throughput when hardware failures occur at scale.

The Hardware Advantage: SupremeRAID™ 2.0, NVIDIA RTX 2000E Ada, and InnoGrit N3X

This test platform represents the ideal architecture for modern AI servers, combining next-generation GPU offload with ultra-low latency media to eliminate bottlenecks.

The RAID engine is the SupremeRAID™ Ultra, powered by the NVIDIA RTX 2000E Ada Generation GPU. This single-slot, low-profile accelerator operates within a highly efficient 50W power envelope. Crucially, it is driven by the SupremeRAID™ 2.0 driver, which features a completely redesigned I/O offload engine. This new architecture optimizes the data path to deliver a generational leap in RAID5/6 random write efficiency and ensures that performance remains consistent even during degraded states.

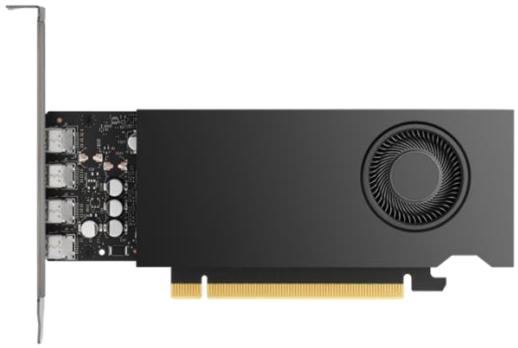

The storage media consists of **24x InnoGrit N3X NVMe SSDs,** powered by **KIOXIA XL-FLASH™** technology. These drives bridge the gap between DRAM and standard NAND, utilizing **SLC (Single-Level Cell)** architecture to ensure **extreme low latency** and deterministic performance. By pairing the SupremeRAID™ 2.0 engine with the N3X's raw speed, we achieve a storage subsystem capable of saturating the PCIe bus with parity protection enabled, preventing the "tail latency" spikes that can stall parallel AI training workloads.

Note: Official support for InnoGrit N3X drives will be included in the upcoming SupremeRAID™ 2.0 driver minor release.

What AI Workloads Demand From Storage

AI data platforms stress storage systems in ways that traditional enterprise applications do not. The workflow combines:

1. **High parallel reads** during training and dataset shuffling.
2. **Write-heavy bursts** during checkpointing and logging.
3. **Continuous metadata activity** from distributed data services.

Parity RAID (RAID5/6) is attractive because it offers capacity efficiency at scale. However, the challenge has always been keeping parity RAID fast under heavy random writes and maintaining performance during **degraded mode** (when a drive fails and reconstruction enters the data path). SupremeRAID™ 2.0 is specifically engineered to solve these two bottlenecks

---

Performance Comparisons

Test Descriptions

The results below compare **Linux MD (mdadm)** and the **SupremeRAID™ 2.0 Linux driver** on an identical 24-drive NVMe configuration. Each subsection presents the data followed by an analysis of its impact on AI workloads.

Testing Environment

- Hardware
  - **CPU:** AMD EPYC 9755 128-Core Processor × 2
  - **Memory:** 32 GB DDR5-6400 RDIMM × 24
  - **GPU RAID Accelerator:** SupremeRAID™ Ultra (NVIDIA RTX 2000E Ada), single slot, low profile, 50W
  - **NVMe Drives:** InnoGrit N3X SLC NVMe × 24
- Software
  - **OS:** Ubuntu 24.04.2 LTS
  - **Kernel:** 6.8.0-62-generic
  - **RAID Implementations:**
    - Linux MD (mdadm) v4.3
    - SupremeRAID™ 2.0 (2.0.0-uad-76-71)
  - **Benchmark Tool:** fio-3.40
- Configuration
  - One RAID group with 24 physical drives (RAID5 and RAID6)
  - **Optimal:** All drives healthy
  - **Degraded:** One drive failed

1. 4K Random Read Performance

Observation: Random read performance is essential for AI data loading and shuffling. While Linux MD performs well in optimal states, its performance plummets by over 99% in degraded mode due to CPU bottlenecks during reconstruction. SupremeRAID™ 2.0, utilizing the GPU to handle I/O, sustains massive throughput in optimal mode and maintains 12.6 million IOPS even during a drive failure.

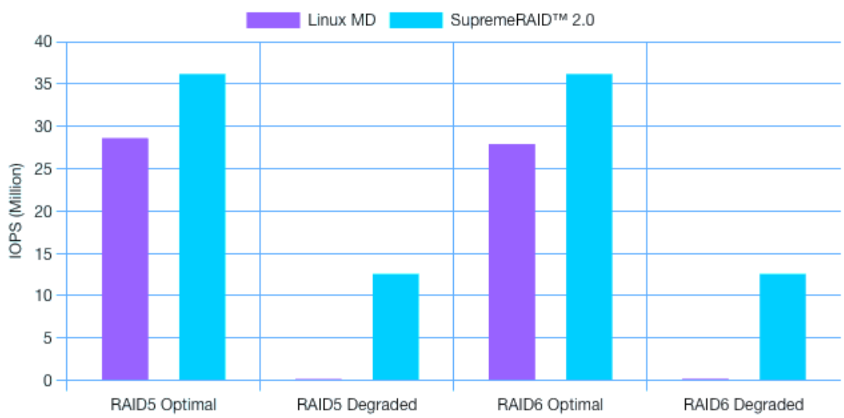

4K Random Read Data (Million IOPS)

Scenario	Linux MD	SupremeRAID™ 2.0	Improvement
RAID5 Optimal	28.6	**36.2**	+26%
RAID5 Degraded	0.163	**12.6**	**77x**
RAID6 Optimal	27.9	**36.2**	+29%
RAID6 Degraded	0.186	**12.6**	**67x**

2. 4K Random Write Performance

Observation: This is the most critical metric for metadata updates and checkpointing. Traditional RAID collapses under the "write hole" penalty, delivering sub-million IOPS. SupremeRAID™ 2.0 revolutionizes this workload, delivering over 6.4 million IOPS in optimal mode and sustaining nearly identical performance in degraded mode.

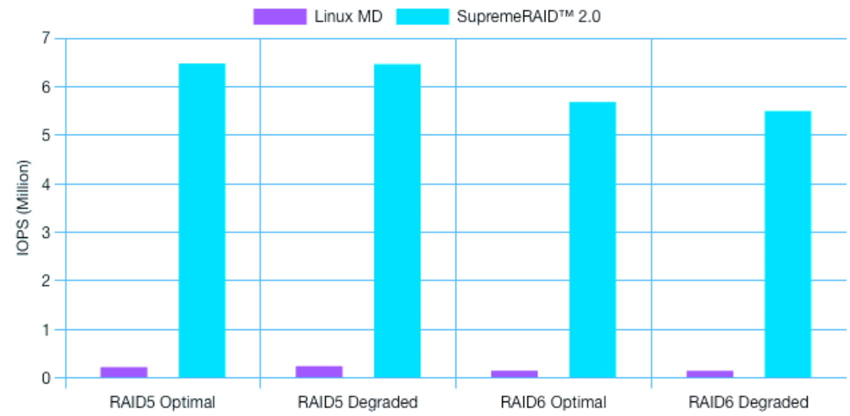

4K Random Write Data (Million IOPS)

Scenario	Linux MD	SupremeRAID™ 2.0	Improvement
RAID5 Optimal	0.223	**6.477**	**29x**
RAID5 Degraded	0.246	**6.466**	**26x**
RAID6 Optimal	0.149	**5.687**	**38x**
RAID6 Degraded	0.147	**5.499**	**37**

---

3. 1M Random Read Performance

Observation: Large block reads correspond to high-throughput ingestion and training data feeds. While optimal performance is similar (limited by drive/bus caps), degraded mode reveals the true difference. SupremeRAID™ sustains over 200 GB/s, whereas Linux MD drops to ~13 GB/s, a bottleneck that would stall an entire GPU cluster.

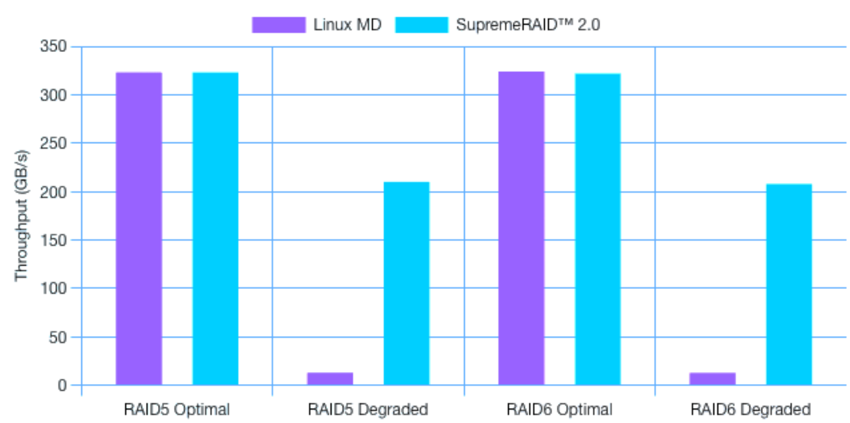

1M Random Read Data (GB/s)

Scenario	Linux MD	SupremeRAID™ 2.0	Improvement
RAID5 Optimal	323	**323**	
RAID5 Degraded	12.9	**210**	**16x**
RAID6 Optimal	324	**322**	
RAID6 Degraded	12.7	**208**	**16x**

---

4. 1M Random Write Performance

Observation: Large block writes are critical for checkpointing AI models. SupremeRAID™ 2.0 delivers 242 GB/s in optimal mode, outperforming Linux MD by over 16x. Crucially, it maintains >200 GB/s even in degraded mode, ensuring that a single drive failure does not interrupt the training process.

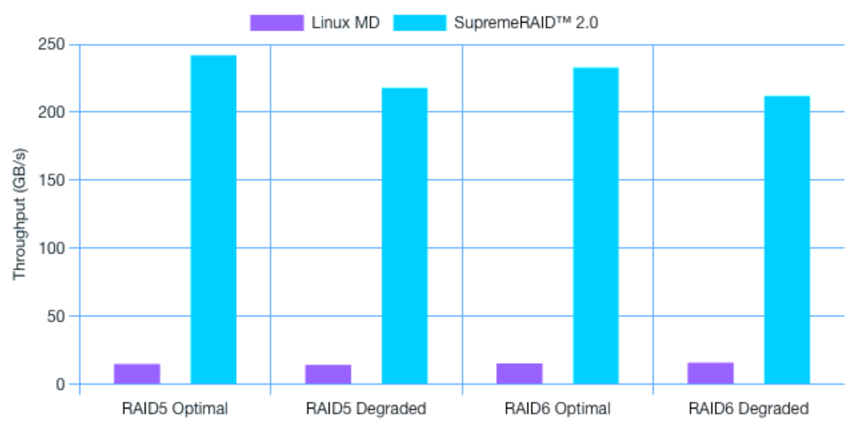

1M Random Write Data (GB/s)

Scenario	Linux MD	SupremeRAID™ 2.0	Improvement
RAID5 Optimal	14.8	**242**	**16x**
RAID5 Degraded	14.2	**218**	**15x**
RAID6 Optimal	15.4	**233**	**15x**
RAID6 Degraded	15.8	**212**	**13x**

CPU Efficiency Analysis (RAID5 Optimal)

This section evaluates the CPU cost of storage operations. By normalizing CPU utilization against performance output, we isolate the true "cost" of the storage software. Lower values indicate higher efficiency, meaning more CPU cycles are left available for AI and compute workloads.

Formula:
The "Normalized Cost" represents the percentage of total system CPU capacity required to deliver one specific unit of performance. It is calculated as:
Normalized Cost = Total CPU Utilization / Performance Achieved (in target units)

- Total CPU Utilization: Calculated as 100% - Idle
- Target Units: 1 Million IOPS (for 4K) or 10 GB/s (for 1M).

Read Efficiency: Small & Large Block

Observation: In read operations, SupremeRAID™ eliminates the CPU bottlenecks inherent in software RAID. For 4K random reads, it is 4.7x more efficient per IOPS. For large-block 1M transfers, the efficiency gap widens to 26x, demonstrating SupremeRAID™'s ability to act as a high-speed DMA engine that bypasses host CPU limitations.

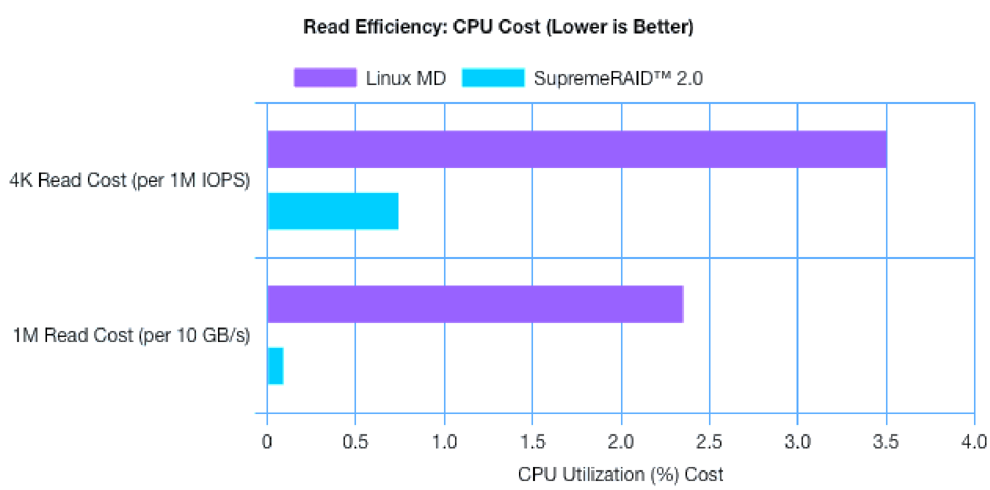

Read Efficiency Data

Workload	Metric Unit	Solution	Performance (Units)	CPU Used	Normalized Cost	Efficiency Gain
4K Random Read	per 1M IOPS	Linux MD	28.6	100.00%	**3.50%**	—
		SupremeRAID™	36.2	26.88%	**0.74%**	**4.7x**
1M Random Read	per 10 GB/s	Linux MD	32.3	76.06%	**2.35%**	—
		SupremeRAID™	32.3	3.04%	**0.09%**	**26x**

Write Efficiency: Small & Large Block

Observation: The contrast is most severe in write workloads. Linux MD suffers from heavy parity calculation and interrupt overhead, consuming over 77% of a CPU's capacity just to deliver a fraction of the performance. SupremeRAID™ offloads this burden to the GPU, resulting in a 41x efficiency gain for random IOPS and a massive 107x gain for throughput, effectively making high-performance RAID parity transparent to the host CPU.

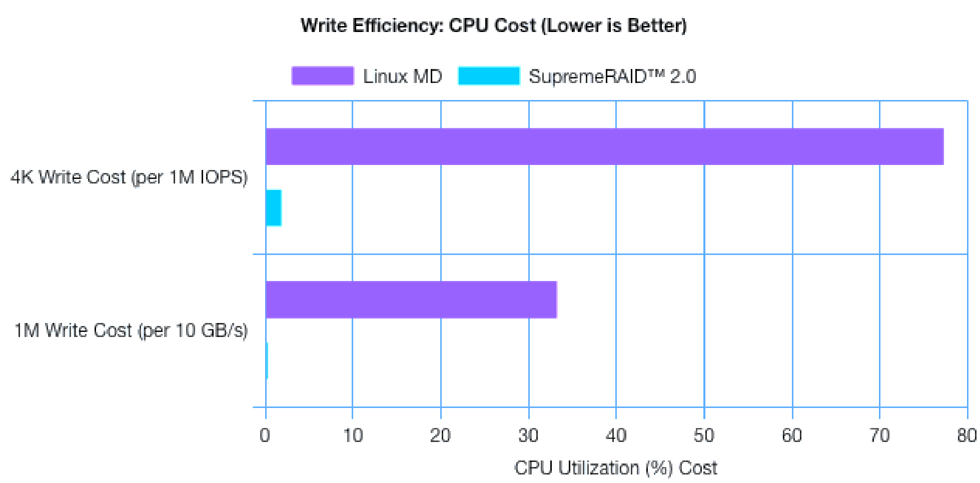

Write Efficiency Data

Workload	Metric Unit	Solution	Performance (Units)	CPU Used	Normalized Cost	Efficiency Gain
4K Random Write	per 1M IOPS	Linux MD	0.22	17.23%	**77.26%**	—
		SupremeRAID™	6.48	12.09%	**1.87%**	**41x**
1M Random Write	per 10 GB/s	Linux MD	1.48	49.20%	**33.24%**	—
		SupremeRAID™	24.2	7.60%	**0.31%**	**107x**

Conclusion

AI workloads demand a storage architecture that is massive, parallel, and resilient. The results in this evaluation confirm that the SupremeRAID™ 2.0 driver, with its redesigned I/O offload engine, delivers performance perfectly aligned with these requirements. By unlocking the full potential of InnoGrit N3X SLC NVMe drives powered by KIOXIA XL-FLASH™, the solution achieves multi-million IOPS random writes in optimal mode and robust, record-class performance even when failures occur.

Furthermore, SupremeRAID™ achieves this throughput with significantly superior CPU efficiency by offloading I/O handling to the GPU, freeing up host compute resources for core AI tasks. Powered by the compact 50W NVIDIA RTX 2000E Ada, this solution provides a dense, power-efficient foundation for next-generation AI storage nodes, ensuring that the extreme low latency and consistency of the N3X media are fully realized.

Appendix

Benchmarking Instructions

SupremeRAID™ RAID5

1. Create physical drives:
```sh
sudo graidctl create pd /dev/nvme0-23
```
2. Create a RAID5 drive group:
```sh
sudo graidctl create dg raid5 0-23
```
3. Create a virtual drive using all available space:
```sh
sudo graidctl create vd 0
```
4. Run fio with the parameters defined in the Detailed fio Parameters to measure optimal performance.
5. Mark the first physical drive offline to force degraded state:
```sh
sudo graidctl edit pd 0 marker offline
```
6. Run fio again to measure degraded performance.

SupremeRAID™ RAID6

1. Create physical drives:
```sh
sudo graidctl create pd /dev/nvme0-23
```
2. Create a RAID6 drive group:
```sh
sudo graidctl create dg raid6 0-23
```
3. Create a virtual drive using all available space:
```sh
sudo graidctl create vd 0
```
4. Run fio (optimal).
5. Mark the first physical drive offline:
```sh
sudo graidctl edit pd 0 marker offline
```
6. Run fio (degraded).

MD RAID5

1. Create the MD RAID5 array (/dev/md5):

```sh
sudo bash -c '
NVME_LIST=($(nvme list | grep INNOGRIT | awk "{print \$1}"))
mdadm --create /dev/md5 \
  --verbose \
  --level=5 \
  --raid-devices=24 \
  --chunk=16K \
  --consistency-policy=resync \
  --force \
  "${NVME_LIST[@]:0:24}"
'
```

2. Increase parity processing threads:

```sh
echo 32 | sudo tee /sys/block/md5/md/group_thread_cnt
```

3. Run fio (optimal).
4. Fail one member device:

```sh
sudo mdadm --manage /dev/md5 --fail /dev/nvme0n1
```

5. Run fio (degraded).

MD RAID6

1. Create the MD RAID6 array (/dev/md6):

```sh
sudo bash -c '
NVME_LIST=($(nvme list | grep INNOGRIT | awk "{print \$1}"))
mdadm --create /dev/md6 \
  --verbose \
  --level=6 \
  --raid-devices=24 \
  --chunk=16K \
  --consistency-policy=resync \
  --force \
  "${NVME_LIST[@]:0:24}"
'
```

2. Increase parity processing threads:

```sh
echo 32 | sudo tee /sys/block/md6/md/group_thread_cnt
```

3. Run fio (optimal).
4. Fail one member device:

```sh
sudo mdadm --manage /dev/md6 --fail /dev/nvme0n1
```

5. Run fio (degraded).

Fio Parameters

```sh
[global]
filename=/dev/gdg0n1
# Note: Use /dev/gdg0n1 for SupremeRAID, /dev/md5 or /dev/md6 for Linux MD
randrepeat=0
ioengine=libaio
direct=1
random_generator=tausworthe64
cpus_allowed_policy=split
group_reporting=1
norandommap=1
time_based=1
runtime=60
wait_for_previous=1

[precondition]
rw=write
bs=1m
numjobs=100
iodepth=32
size=1%
offset_increment=1%
time_based=0

[4k_random_read]
rw=randread
bs=4k
numjobs=512
iodepth=16
cpus_allowed=0-511

[1m_random_read]
rw=randread
bs=1m
numjobs=512
iodepth=8
cpus_allowed=0-511

[1m_random_write]
rw=randwrite
bs=1m
numjobs=256
iodepth=16
cpus_allowed=0-255

[4k_random_write]
rw=randwrite
bs=4k
numjobs=256
iodepth=16
cpus_allowed=0-255
```

Detailed Benchmark Results

RAID5 Optimal (SupremeRAID™ vs Linux MD)

Workload	SupremeRAID™ RAID5	CPU (user/system/idle)	Linux MD RAID5	CPU (user/system/idle)
4K Rand Read	36.2M IOPS	4.79 / 22.08 / 73.12	28.6M IOPS	8.87 / 91.13 / 0.00
4K Rand Write	6.477M IOPS	1.35 / 10.73 / 87.91	223k IOPS	0.11 / 17.13 / 82.77
1M Rand Read	301GiB/s(323GB/s)	0.11 / 2.93 / 96.96	301GiB/s(323GB/s)	0.12 / 75.94 / 23.94
1M Rand Write	225GiB/s(242GB/s)	1.07 / 6.53 / 92.40	13.8GiB/s(14.8GB/s)	0.05 / 49.15 / 50.80

RAID6 Optimal (SupremeRAID™ vs Linux MD)

Workload	SupremeRAID™ RAID6	CPU (user/system/idle)	Linux MD RAID6	CPU (user/system/idle)
4K Rand Read	36.2M IOPS	4.79 / 22.08 / 73.12	27.9M IOPS	7.09 / 92.91 / 0.00
4K Rand Write	5.687M IOPS	1.86 / 16.89 / 81.26	149k IOPS	0.03 / 50.31 / 49.66
1M Rand Read	300GiB/s(322GB/s)	0.10 / 2.80 / 97.10	302GiB/s (324GB/s)	0.13 / 78.46 / 21.42
1M Rand Write	217GiB/s (233GB/s)	0.84 / 7.11 / 92.05	14.4GiB/s (15.4GB/s)	0.05 / 52.67 / 47.28

RAID5 Degraded (SupremeRAID™ vs Linux MD)

Workload	SupremeRAID™ RAID5	CPU (user/system/idle)	Linux MD RAID5	CPU (user/system/idle)
4K Rand Read	12.6M IOPS	3.04 / 21.03 / 75.93	163k IOPS	0.06 / 89.37 / 10.56
4K Rand Write	6.466M IOPS	0.74 / 5.59 / 93.67	246k IOPS	0.11 / 14.16 / 85.73
1M Rand Read	195GiB/s (210GB/s)	0.11 / 7.30 / 92.59	12.0GiB/s (12.9GB/s)	0.01 / 98.47 / 1.51
1M Rand Write	203GiB/s (218GB/s)	0.75 / 5.07 / 94.18	13.3GiB/s (14.2GB/s)	0.05 / 48.89 / 51.06

RAID6 Degraded (SupremeRAID™ vs Linux MD)

Workload	SupremeRAID™ RAID6	CPU (user/system/idle)	Linux MD RAID6	CPU (user/system/idle)
4K Rand Read	12.6M IOPS	2.96 / 20.74 / 76.31	186k IOPS	0.06 / 91.28 / 8.66
4K Rand Write	5.499M IOPS	0.77 / 8.35 / 90.88	147k IOPS	0.03 / 51.03 / 48.94
1M Rand Read	194GiB/s (208GB/s)	0.10 / 6.97 / 92.93	11.8GiB/s (12.7GB/s)	0.01 / 98.42 / 1.57
1M Rand Write	197GiB/s (212GB/s)	0.66 / 6.19 / 93.15	14.7GiB/s (15.8GB/s)	0.05 / 51.69 / 48.26

Disclaimer: Performance results may vary depending on system configuration, workload, and other factors. The tests described in this whitepaper reflect performance under specific conditions.

About Graid Technology

Graid Technology is led by a dedicated team of experts with decades of experience in the SDS, ASIC, and storage industries, and continues to push boundaries in data storage innovation by protecting NVMe-based data from the desktop to the cloud. Cutting-edge SupremeRAID™ GPU-based RAID removes the traditional RAID bottleneck to deliver maximum SSD performance without consuming CPU cycles or creating throughput bottlenecks, delivering unmatched flexibility, performance, and value. With headquarters in Silicon Valley supported by a robust R&D center in Taiwan, we are globally committed to spearheading advancements in storage solutions. For detailed product information, visit our website, or connect with us on LinkedIn.

About InnoGrit

InnoGrit is focused on advancing storage technology to solve the data storage and data transport problem in artificial intelligence and other big data applications through innovative integrated circuit (IC) and system solutions. Our innovation in system performance, efficiency, reliability, and security is the foundation that supports our mission to unleash the potential limited by traditional data processing architectures and enable a new class of products to consumers, data centers, and enterprises. With a reputation for technology, quality and reliability, InnoGrit is a trusted partner to leading global OEMs, semiconductor companies, and flash vendors. For detailed product information, visit our website, or connect with us on LinkedIn.