SOLUTION BRIEF

# MAXIMIZING AI/ML PERFORMANCE AND DATA RESILIENCE WITH SUPREMERAID™ AE ON SUPERMICRO GPU SERVERS

TABLE OF CONTENTS



## Executive Summary


The era of Artificial Intelligence (AI) and Machine Learning (ML) demands unprecedented data processing capabilities and unwavering data integrity. Data resilience is essential for both AI training and inference, as these processes rely on large, valuable datasets that must be protected from loss, corruption, and hardware failures. During training, interruptions or data loss can significantly set back progress and increase costs, while for inference, resilient storage ensures uninterrupted service delivery and reliable model outputs, which are critical in production environments supporting real-time decision-making.

Supermicro GPU-accelerated servers are at the forefront of this revolution. This whitepaper explores how Graid Technology's SupremeRAID AE (AI Edition), tested on a Supermicro SYS-821GE-TNHR GPU server, delivers exceptional RAID 5 data protection for NVMe SSDs, ensuring both high performance and critical data resilience for AI/ML pipelines —from training to inference. We demonstrate that SupremeRAID AE not only protects data but also sustains near-native NVMe bandwidth, even under simulated AI I/O patterns using NVIDIA GPUDirect® Storage (GDS) compatible tools.


**GPU SUPERSERVER**
**SYS-821GE-TNHR**

## The AI Data Challenge: Performance Meets Unyielding Resilience

The proliferation of AI and ML is transforming industries. Training complex models and deploying them for real-time inference requires processing colossal datasets. Supermicro GPU servers, equipped with powerful NVIDIA and AMD GPUs and high-speed NVMe storage, provide the computational backbone for these demanding tasks.

However, the sheer volume of data and the intensity of I/O operations in AI environments amplify the risks associated with data integrity:

- **Drive Failures:** While NVMe SSDs are reliable, failures are inevitable in large-scale deployments. A single drive failure in an unprotected array can halt multi-day training jobs, leading to significant financial and time losses.

- **Uncorrectable Bit Errors (UBER):** SSDs can experience read errors that cannot be corrected by their internal ECC. The UBER for enterprise SSDs might be low (e.g., 1 in 10^17 bits read), but in AI workloads that continuously read petabytes of data, the probability of encountering such an error becomes non-negligible. A single UBER can corrupt a dataset, invalidate model training progress, or bring down an inference service until the affected data is identified and reconstructed.

For AI/ML, data isn't just stored; it's the lifeblood. Any interruption or corruption can have catastrophic consequences. Therefore, data resilience is not a luxury but a fundamental requirement.

## The Imperative for NVMe RAID in GPU Servers

The need for robust data protection for NVMe drives in GPU servers is paramount. Key considerations include:

- **Protection Against Drive Failure:** RAID 5 or RAID 6 can sustain one or two drive failures, respectively, allowing workloads to continue uninterrupted while a failed drive is replaced and the array rebuilt.
- **Mitigation of Read Errors (UBER):** RAID's inherent redundancy can be leveraged to recover data affected by UBER, preventing silent data corruption and a forced, lengthy data reconstruction process.

The challenge lies in implementing RAID in a manner that does not compromise the extreme performance demanded by AI workloads.

## Introducing SupremeRAID AE: Resilience Without Compromise

Graid Technology's SupremeRAID AE is a revolutionary NVMe/NVMeoF RAID solution that leverages the power of the GPU to eliminate I/O bottlenecks and deliver exceptional performance with robust data protection. By offloading RAID calculations to the GPU, SupremeRAID AE frees up CPU resources and enables NVMe SSDs to operate closer to their full potential.

**Key Features relevant to AI/ML:**

- **GPU-Powered Performance:** Utilizes existing GPU resources for RAID processing, ensuring minimal CPU overhead.
- **RAID Levels:** Supports RAID 0, 1, 5, 6, 10 for flexible data protection strategies.
- **Automated Bad Block Recovery:** Proactively handles drive errors and UBER.

## Automated Bad Block Recovery with SupremeRAID AE

NAND flash, the core of NVMe SSDs, is susceptible to wear and data degradation over time, potentially leading to Uncorrectable Bit Errors (UBER). While SSD controllers have sophisticated error correction, UBERs can still occur.

SupremeRAID AE's Automated Bad Block Recovery mechanism enhances data integrity:

1. **Error Detection:** When a read operation encounters a URE on a drive within a redundant RAID array (e.g., RAID 5), SupremeRAID AE detects this.
2. **Data Reconstruction:** Using the parity data (for RAID 5) or dual parity (for RAID 6) from the other drives in the array, SupremeRAID AE reconstructs the data from the affected block in real-time.
3. **Transparent Overwrite:** The reconstructed, correct data is then written back to the original location on the problematic drive. This action can often resolve the issue by remapping the bad block at the drive level or by refreshing the data.

4. **Seamless Operation:** This entire process is transparent to the application, ensuring data consistency and preventing the URE from propagating or causing application-level errors.

Automated Bad Block Recovery leverages the RAID redundancy to mitigate the impact of UBERs, preventing data loss and minimizing disruptions to critical AI workloads without requiring manual intervention or full drive rebuilds for isolated bit errors.

## Performance Validation: SupremeRAID AE on Supermicro GPU Server

To quantify the benefits of SupremeRAID AE, Supermicro conducted a series of performance tests on a Supermicro GPU server.

### 5.1. The Supermicro SYS-821GE-TNHR: An AI Powerhouse

The Supermicro SYS-821GE-TNHR is an 8U GPU server meticulously engineered for the most demanding AI/ML, HPC, and big data analytics workloads. Its key features include:

- **Processor Support:** Dual Socket E (LGA 4677) supporting 4th and 5th Gen Intel® Xeon® Scalable processors, providing substantial CPU horsepower for complex tasks.
- **GPU Capacity:** Optimized for up to 8x double-width, full-height, full-length PCIe 5.0 x16 GPUs, enabling massive parallel processing capabilities crucial for AI model training and inference.
- **High-Speed Memory:** Supports up to 8TB of DDR5-4800MHz ECC RDIMM memory across 32 DIMM slots, ensuring ample bandwidth for large datasets.
- **Flexible, High-Performance Storage:** Features up to 24x 2.5" hot-swap drive bays with support for NVMe (PCIe 5.0), SAS, and SATA drives, offering versatile and scalable storage options. This test configuration utilizes 9 of these NVMe bays.
- **Advanced Networking:** Includes AIOM/OCP 3.0 support for flexible, high-speed networking options (e.g., 200Gb/s InfiniBand or Ethernet).
- **Robust Design:** Built with redundant power supplies (Titanium Level) for high availability and system reliability.

This server platform provides an ideal environment for showcasing the synergy between high-performance NVMe storage and SupremeRAID AE's GPU-accelerated data protection.

### 5.2. Test Bed Configuration

- **Server:** Supermicro SYS-821GE-TNHR
- **CPU:** 2 x Intel® Xeon® Platinum 8462Y+
- **GPU (for AI & RAID):** 8 x NVIDIA H100 80GB HBM3
- **NVMe SSDs:** 9 x Samsung PM1743 3.84TB (Model: MZWL03T8HCLS-00A07)
- **Operating System:** Ubuntu 22.04.5 LTS (Kernel: 5.15.0-131-generic)
- **RAID Solution:** SupremeRAID AE (Version: 1.6.1-38.g5c902892.hp90)
- **Benchmarking Tools:**
    - fio (Version: 3.28) for direct storage performance.
    - gdsio (NVIDIA GDS tools, nvidia_fs v2.19, libcufile v2.12) for GDS-compatible AI I/O simulation.

## Testing Methodology

The performance evaluation focused on throughput (GiB/s) for large-block sequential and random I/O patterns, which are representative of data loading, checkpointing, and large-scale data processing phases common in AI/ML workloads. All tests were conducted with a significant dataset size to ensure steady-state performance and a runtime of 300 seconds for fio tests (or

60 seconds for gdsio examples, potentially averaged over multiple runs) to ensure stable results. Throughout this evaluation, baseline measurements were obtained using standard Linux MD/RAID configurations, allowing for direct comparison between the tested system and widely adopted RAID solutions.

A key aspect of this methodology is the understanding that a RAID array, by its nature of striping data across multiple drives, can transform larger logical I/O requests from the application into multiple, smaller physical I/O requests to its member drives. Therefore, to estimate the theoretical potential of the RAID 5 array:

- The **128k random read/write throughput on RAID 5** was theoretically estimated based on the **16k random read/write performance of a single NVMe baseline drive**. This reflects the scenario where a 128k RAID I/O might be serviced by multiple ~16k I/Os distributed across the data (and parity for writes) drives.
- The **1M sequential read/write throughput on RAID 5** was theoretically estimated based on **the 128k random read/write performance of a single NVMe baseline drive**. This simulates a scenario in which large sequential streams are split into 128k chunks for distribution across the drives in the RAID set.

1. **Baseline Performance (Single NVMe Drive & Aggregated 9 NVMe Drives - No RAID):**

- **Objective:** Establish the raw performance capabilities of individual Samsung PM1743 drives and the true aggregate performance of all 9 drives working in parallel.
- **Tool:** fio (Version: 3.28)
- **Target (Single Drive):** Individual /dev/nvmeXnY devices (e.g., /dev/nvme0n1).
- **fio Parameters for Single Drive (Example for 16k Random Read):**

```sh
fio --name=baseline_16k_randread_single_nvme0n1 --filename=/dev/nvme0n1 --direct=1 \
--rw=randread --bs=16k --ioengine=libaio --iodepth=128 --numjobs=8 --group_reporting \
--time_based --runtime=300 --norandommap
```

- Repeat for rw=randwrite, bs=128k, rw=read (sequential), rw=write (sequential), bs=1M.
- The values for "16k random read/write" and "128k random read/write" in the "Single NVMe Disk - fio" column of the document (page 1) are derived from such single-drive tests.

**2. SupremeRAID AE RAID 5 (9 Drives - fio Tests):**

- **Objective:** Measure the performance of the SupremeRAID AE RAID 5 array using standard block I/O tools.
- **Tool:** fio (Version: 3.28)
- **Target:** SupremeRAID AE RAID 5 volume (/dev/gdg0n1)
- **fio Command (Example for 1M Sequential Read, matching FIO config on page 4):**

```sh
fio --name=srae_raid5_test --filename=/dev/gdg0n1 --ioengine=libaio \
--direct=1 --randrepeat=0 --random_generator=tausworthe64 --cpus_allowed_policy=split \
--group_reporting --norandommap=1 --time_based --runtime=300 --rw=read --bs=1M \
--iodepth=64 --numjobs=128
```

- Adjust rw (e.g., randread, write, randwrite) and bs (e.g., 128k) for other tests.

**3. SupremeRAID AE RAID 5 (9 Drives - gdsio Compatibility Mode Tests):**

A. **Objective:** Simulate AI workloads accessing the SupremeRAID AE RAID 5 volume using NVIDIA GPUDirect® Storage (GDS) in compatibility mode.
B. **Tool:** gdsio (NVIDIA GDS tools)
C. **Target:** SupremeRAID AE RAID 5 volume (/dev/gdg0n1)
D. **Environment Variable:** For "gdsio (compatible) on 8 GPUs" tests: export CUDA_VISIBLE_DEVICES=0,1,2,3,4,5,6,7.
E. **gdsio Command (Example for 1M Sequential Read on 8 GPUs):**
The command on page 3 appears to be the most direct reference for testing /dev/gdg0n1 with gdsio. It uses multiple -f, -w, and -d flags. This structure launches multiple parallel I/O streams/workers, each assigned to a GPU and targeting the same file, with its own I/O depth.

```sh
# For testing with 8 GPUs, using CUDA_VISIBLE_DEVICES=0,1,2,3,4,5,6,7
# The -w value comes from the table note e.g., (-w 128) for 1M Seq Read on 8 GPUs
/usr/local/cuda-12.4/gds/tools/gdsio -s 1024G -T 300 -i 1M -x 0 \
-f /dev/gdg0n1 -w 128 -d 0 \
-f /dev/gdg0n1 -w 128 -d 1 \
-f /dev/gdg0n1 -w 128 -d 2 \
-f /dev/gdg0n1 -w 128 -d 3 \
-f /dev/gdg0n1 -w 128 -d 4 \
-f /dev/gdg0n1 -w 128 -d 5 \
-f /dev/gdg0n1 -w 128 -d 6 \
-f /dev/gdg0n1 -w 128 -d 7
```

F. Adjust -x (0 for read, 1 for write) and -i (e.g., 128k) for other block sizes/operations.
G. The -w value (e.g., 128) is repeated for each GPU stream as per the example and should correspond to the (-w XXX) value in the "gdsio (compatible) on 8 GPUs" column of the results table.

**4. Linux MD RAID 0 (9 Drives - gdsio Full Path Tests):**

- **Objective:** Establish a high-performance, unprotected baseline using Linux software RAID 0 with gdsio utilizing full GDS paths.
- **Tool:** gdsio (NVIDIA GDS tools)
- **Target:** Linux MD RAID 0 volume (/dev/md0)
- **Environment Variable:** export CUDA_VISIBLE_DEVICES=0,1,2,3,4,5,6,7
- **gdsio Command (Example for 1M Sequential Read):**
For the "md RAID 0 9 disks - gdsio on 8 GPUs" column, the document doesn't explicitly show the (-w XXX) value. We'll use a structure similar to the SupremeRAID gdsio test but assume a potentially higher individual -w value to fully saturate MD RAID 0 with direct GDS. If a specific -w was used, that should be substituted. Page 10 shows gdsio with MD RAID 0 using multiple -w 256 for a 4k test; for 1M, we might infer a similar structure or a higher single effective -w.

```sh
# Example for 1M Sequential Read on MD RAID0 with 8 GPUs
# Assuming -w 128 per stream as per other 1M read tests, adjust if known otherwise for this column
/usr/local/cuda-12.4/gds/tools/gdsio -s 1024G -T 300 -i 1M -x 0 \
-f /dev/md0 -w 128 -d 0 \
-f /dev/md0 -w 128 -d 1 \
-f /dev/md0 -w 128 -d 2 \
-f /dev/md0 -w 128 -d 3 \
-f /dev/md0 -w 128 -d 4 \
-f /dev/md0 -w 128 -d 5 \
-f /dev/md0 -w 128 -d 6 \
-f /dev/md0 -w 128 -d 7
```

- Adjust -x and -i as needed. The -w value for the "md RAID 0 ... gdsio on 8 GPUs" (non-compat) column should ideally be determined from specific test logs if it differs from the compat mode tests. If it's the same, the command structure would be similar, just targeting /dev/md0.

## Results & Analysis

Before presenting the measured performance of the SupremeRAID AE RAID 5 array, it's essential to establish the baseline performance of the individual NVMe drives and the derived theoretical maximums for a 9-disk RAID 5 configuration. As outlined in the methodology, the theoretical RAID 5 performance is estimated based on single-drive performance, considering that larger RAID I/Os are often broken down into smaller I/Os to member disks. Specifically:

- Theoretical 128k RAID 5 throughput is derived from the single-drive 16k performance.
- Theoretical 1M RAID 5 throughput is derived from the single-drive 128k performance.
- RAID 5 read throughput is calculated as Single_Drive_Performance * N (where N=9 drives), reflecting the document's estimation method for reads.
- RAID 5 write throughput is calculated as Single_Drive_Performance * (N-1) (where N=9 drives, so N-1=8), accounting for one drive's capacity being dedicated to parity.

The following table shows these baseline figures and the calculated theoretical RAID 5 throughputs:

**Table 5.4.1: Single NVMe Drive Baseline and Theoretical 9-Disk RAID 5 Performance**

| Metric | Block Size | Single NVMe Drive (fio Baseline) | Calculation for 9-Disk RAID 5 | Theoretical 9-Disk RAID 5 Throughput |
|---|---|---|---|---|
| For 1M Sequential Read | 128k | 13.9 GiB/s | 13.9 GiB/s * 9 | 125.1 GiB/s |
| For 1M Sequential Write | 128k | 5.67 GiB/s | 5.67 GiB/s * 8 | 45.36 GiB/s |
| For 128k Random Read | 16k | 12.3 GiB/s | 12.3 GiB/s * 9 | 110.7 GiB/s |
| For 128k Random Write | 16k | 5.63 GiB/s | 5.63 GiB/s * 8 | 45.04 GiB/s |

With these theoretical values as a reference, the measured performance of SupremeRAID AE RAID 5 and MD RAID 0 is presented below. All results are in GiB/s.

**Table 5.4.2: Measured Performance Comparison**

| Workload | Block Size | Theoretical SupremeRAID AE RAID 5 (9 Drives) | SupremeRAID AE RAID 5 (fio on 9 Drives) | SupremeRAID AE RAID 5 (gdsio compatible on 8 GPUs) | MD RAID 0 (gdsio on 8 GPUs, No Protection) |
|---|---|---|---|---|---|
| Sequential Read | 1M | 125.1 GiB/s | 117 GiB/s | 113.03 GiB/s | 119.9 GiB/s |
| Sequential Write | 1M | 45.36 GiB/s | 43.3 GiB/s | 43.17 GiB/s | 50.1 GiB/s |
| Random Read | 128k | 110.7 GiB/s | 111 GiB/s | 102.61 GiB/s | 116.4 GiB/s |
| Random Write | 128k | 45.04 GiB/s | 42.0 GiB/s | 39.92 GiB/s | 39.9 GiB/s |

**Visual Performance Comparison:**

**A. Read & Write Performance of SupremeRAID RAID 5 Array - Absolute**

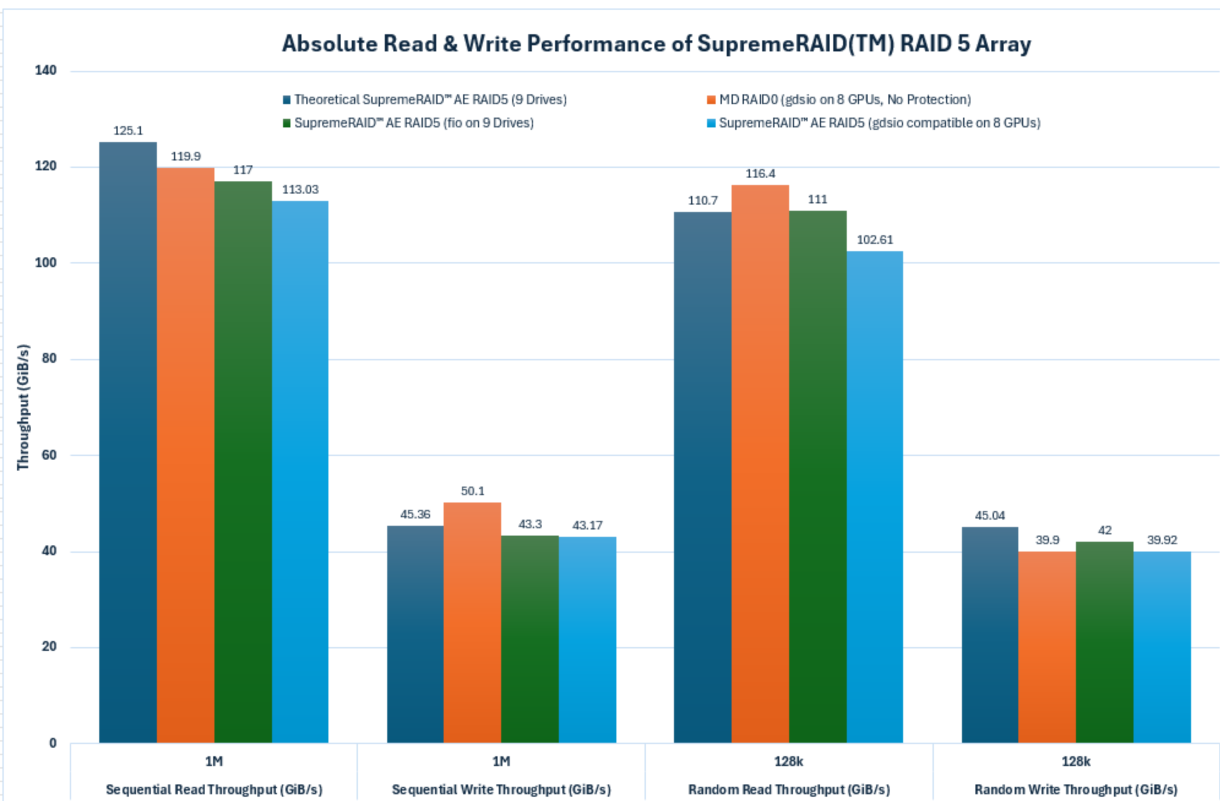

## B. Read & Write Performance of SupremeRAID RAID 5 Array – Relative

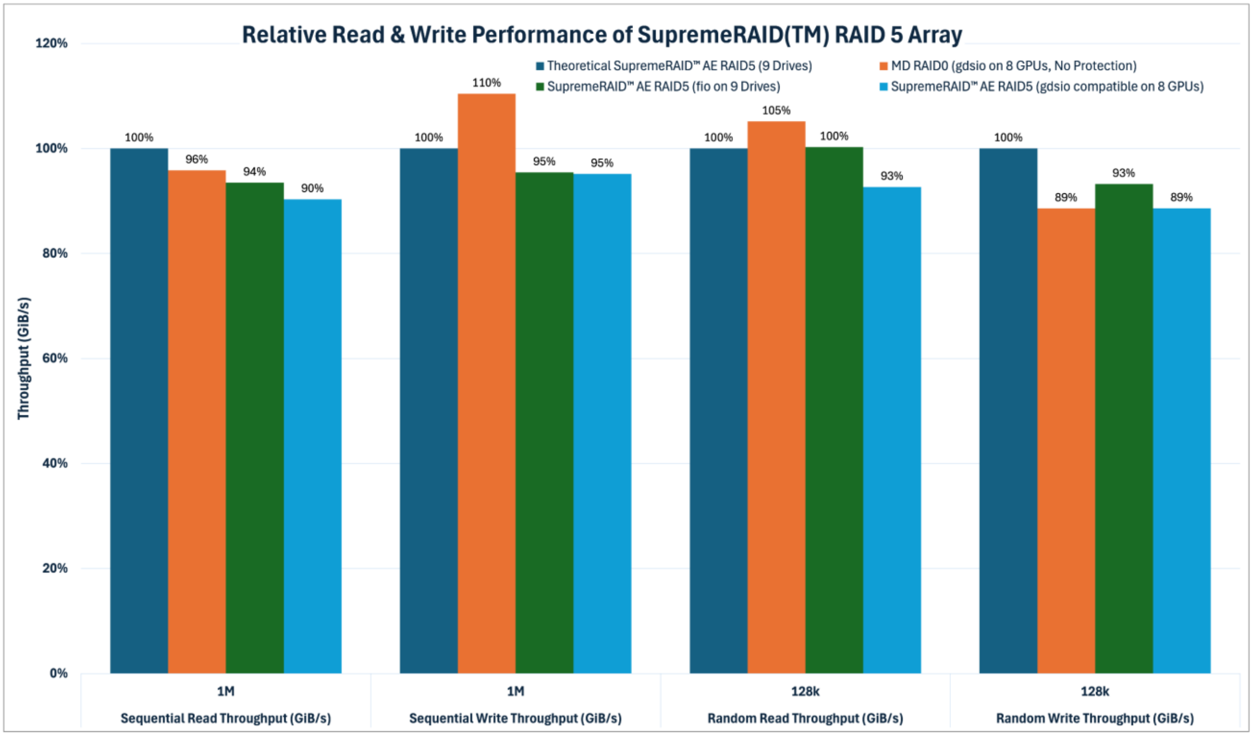


**Analysis of Results:**

- **Near-Theoretical Performance with Standard I/O (fio):**
  - For 1M sequential reads, SupremeRAID AE RAID 5 with fio achieved 117 GiB/s, which is ~93.5% of the calculated theoretical 125.1 GiB/s. This demonstrates excellent efficiency even without GDS-specific data paths.
  - For 128k random reads, the fio result of 111 GiB/s is remarkably close to (and slightly exceeds) the theoretical 110.7 GiB/s, indicating minimal overhead for random read operations on the protected array.
  - Sequential writes (43.3 GiB/s vs. 45.36 GiB/s theoretical) and random writes (42.0 GiB/s vs. 45.04 GiB/s theoretical) also show SupremeRAID AE performing very close to its theoretical potential under standard fiotesting.
- **Strong Performance with GDS Compatibility Mode:**
  - When using gdsio in compatibility mode, SupremeRAID AE RAID 5 continues to deliver robust throughput: 113.03 GiB/s for 1M sequential reads and 102.61 GiB/s for 128k random reads. While slightly lower than direct fio (which doesn't involve the GDS stack complexity), these figures are still very high and represent a significant portion of the theoretical maximum.
  - This performance in GDS compatibility mode is crucial, as it showcases SupremeRAID AE's ability to integrate into AI/ML workflows that leverage GPUDirect Storage, providing data protection without becoming a significant bottleneck.
- **Comparison with Unprotected MD RAID 0 (gdsio):**
  - The MD RAID 0 configuration (using full gdsio paths, no protection) serves as a high-water mark for raw GDS throughput.

        - For 1M sequential reads, SupremeRAID AE RAID 5 (fio: 117 GiB/s, gdsio compat: 113.03 GiB/s) is competitive with MD RAID 0 (gdsio: 119.9 GiB/s), especially considering RAID 5 protection.
        - For 128k random reads, SupremeRAID AE RAID 5 (fio: 111 GiB/s, gdsio compat: 102.61 GiB/s) compares favorably to MD RAID 0 (gdsio: 116.4 GiB/s).
    - The key takeaway is that SupremeRAID AE, even with standard fio (simulating non-GDS optimized applications), achieves performance levels that are very close to, or even exceed in some direct I/O cases, what a fully GDS-enabled unprotected MD RAID 0 can achieve, all while providing essential RAID 5 data protection. For write workloads, the parity overhead of RAID 5 naturally results in lower numbers compared to RAID 0, but SupremeRAID AE still delivers strong protected write performance.

The results clearly demonstrate that SupremeRAID AE provides excellent data protection with minimal performance overhead. Its ability to achieve near-theoretical performance with standard I/O tools like fio is particularly noteworthy, indicating its efficiency even for applications not fully optimized for GDS. When GDS compatibility is engaged, it continues to provide substantial throughput, making it a versatile solution for diverse AI/ML workloads requiring both speed and data safety. Because SupremeRAID AE leverages a small portion of the GPU to deliver high-performance, software-defined RAID, introducing only minimal overhead—approximately 4%—to GPU-intensive AI inference workloads. This allows organizations to maintain strong overall GPU performance for training and inference, while benefiting from the massive throughput and resiliency provided by SupremeRAID AE without needing additional hardware or impacting primary compute tasks.

## Conclusion: Supermicro and SupremeRAID AE – A Resilient Foundation for AI

The demands of modern AI/ML workloads necessitate a storage infrastructure that is not only high-performing but also exceptionally resilient. The test results on the Supermicro SYS-821GE-TNHR server clearly demonstrate that Graid Technology's SupremeRAID AE provides a compelling solution.

By leveraging GPU acceleration, SupremeRAID AE delivers robust RAID 5 data protection for high-speed NVMe SSDs with minimal performance impact. It effectively saturates the available NVMe bandwidth for read operations and provides strong write performance, even when simulating AI I/O patterns with GDS-compatible tools. The added benefit of Automated Bad Block Recovery further solidifies its role in maintaining data integrity against insidious UBERs.

For organizations deploying AI/ML workloads on Supermicro GPU servers, SupremeRAID AE offers an optimal balance of blazing-fast performance and critical data resilience, ensuring that valuable AI models and datasets are protected without compromising the speed and efficiency required to stay competitive in the rapidly evolving AI landscape.

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

## For More Information

Supermicro site: https://www.supermicro.com/en/solutions/graidtechnology

Graid Technology site: https://graidtech.com/partners/supermicro

## SUPERMICRO

As a global leader in high performance, high efficiency server technology and innovation, we develop and provide end-to-end green computing solutions to the data center, cloud computing, enterprise IT, big data, HPC, and embedded markets. Our Building Block Solutions® approach allows us to provide a broad range of SKUs, and enables us to build and deliver application-optimized solutions based upon your requirements.

Visit: www.supermicro.com

## 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.

Visit: www.graidtech.com