

SupremeRAID™ AE (AI Edition)

**AI-Driven RAID Purpose Built
for Enterprise NVMe Storage**



SupremeRAID™ AE (AI Edition) is a next-generation, software-defined **RAID solution** optimized for high-performance NVMe SSD arrays. Designed for modern data centers, AI/ML workloads, and mission-critical enterprise environments, SupremeRAID™ AE combines advanced redundancy, intelligent recovery, and seamless management to deliver unmatched data protection and operational efficiency.



Intelligent Data Offload Engine

By offloading data tasks to optimize GPU utilization, SupremeRAID™ AE enhances system efficiency, allowing for faster, more effective processing of complex AI workloads.



NVMe-oF Support

With the ability to scale storage across multiple hosts, SupremeRAID™ AE supports the massive datasets required for modern AI models while preserving performance, even as data volumes grow.



Enterprise-Grade Data Protection

Designed to maintain uninterrupted access to critical datasets, SupremeRAID™ AE prevents workflow disruptions and safeguards data integrity during intensive AI training and inference processes.



GPUDirect Storage (GDS) Support

SupremeRAID™ AE enables direct data transfers from NVMe drives to GPU memory, bypassing host memory. This seamless integration accelerates workflows with near-zero latency, ensuring faster model training and inference.



Flexible GPU Deployment

SupremeRAID™ AE provides the flexibility to begin with shared GPUs and easily transition to dedicated configurations as workloads expand. This adaptability ensures cost-effective scalability and prepares enterprises for future growth.



Clustering AI Storage Compatibility

SupremeRAID™ AE seamlessly integrates with BeeGFS, Lustre, and Ceph distributed file systems, simplifying large-scale dataset management and reducing storage costs without the need for data migration.

SupremeRAID™ AE (AI Edition)

Benchmark Performance Numbers

✓ **Prevents GPU idling** caused by drive failures, ensuring uninterrupted workloads.

✓ **Achieves over 95% of raw NVMe performance** from GPU with RAID protection.

Hardware


Server Model: Supermicro SYS-821GE-TNHR, CPU: Intel(R) Xeon(R) Platinum 8462Y+ *2, Memory: M321R8GA0BB0-CQKZJ 64GB DDR5 4400 MT/s *32, GPU: NVIDIA H100 *8, NVMe Drive: SAMSUNG PM1743 3.84T MZWL03T8HCLS-00A07 *9

Software

OS: Ubuntu 22.04.5 LTS, Kernel: 5.15.0-131-generic, SupremeRAID™ AE Driver: 1.6.1-38, Benchmark Tool: fio-3.30, gdsio-1.11 on H100*8 in GDS compatible mode

RAID Configuration

RAID5 with SAMSUNG PM1743 3.84T *9



	128K Random Write	128K Random Read
RAID5 Theoretical (9 NVMe)	45 GiB/s	110 GiB/s
fio	44 GiB/s	107 GiB/s
gdsio	43 GiB/s	104 GiB/s

Competitive Differentiators

- **AI-Optimized Parity & Recovery:** Enhanced algorithms for faster rebuilds and minimal performance impact during failures.
- **Enterprise-Grade Resilience:** Double failure protection and proactive recovery exceed traditional RAID implementations.
- **Automation-Ready:** RESTful API and CLI enable full automation and integration into DevOps and SRE workflows.
- **NVMe-Centric Design:** Built from the ground up for NVMe SSDs, not retrofitted from legacy RAID architectures.

Use Cases

- **AI/ML Data Pipelines:** High-throughput, low-latency storage for training and inference workloads.
- **Enterprise Databases:** Maximum data protection and uptime for mission-critical transactional systems.
- **Virtualization & Cloud:** Scalable, resilient storage for private clouds and virtualized environments.
- **Media & Content Delivery:** Fast, reliable storage for high-resolution video and content distribution.

Key Features & Technical Benefits

Double Failure Protection with Distributed Journaling

Distributed journaling architecture provides double failure protection across multiple drives and controllers.

SupremeRAID™ AE maintains data integrity and consistency even if two drives or a drive and controller fail simultaneously, ensuring business continuity during unexpected outages or power loss.

Automated Bad Block Recovery

Real-time detection of uncorrectable read errors (UREs) with automated recovery using RAID parity.

SupremeRAID™ AE minimizes risk of data loss and maximizes array reliability by automatically reconstructing and repairing corrupted data blocks without manual intervention.

Dataset Management (DSM) Deallocate Command (TRIM) Support

Advanced support for TRIM commands on virtual drives, including partial stripe and parity management.

SupremeRAID™ AE enhances SSD performance and lifespan by efficiently reclaiming unused storage, reducing write amplification, and optimizing garbage collection.

RESTful API for Automation & Integration

Comprehensive RESTful API supporting full CRUD operations, HTTP response codes, and detailed status reporting.

SupremeRAID™ AE simplifies integration with orchestration platforms, monitoring tools, and custom automation scripts, enabling seamless management and scaling in modern data center environments.

Intuitive Management GUI

Web-based graphical interface for system monitoring and management, including dashboards, event logs, and real-time statistics.

SupremeRAID™ AE empowers administrators with clear visibility and control, streamlining day-to-day operations and enabling proactive issue resolution.

Detailed Technical Specifications

For questions about product specs, email us at info@graidtech.com

Supported RAID Levels	RAID 0/1/5/6/10 and advanced proprietary levels (AI-optimized parity schemes)
Supported Enterprise GPUs	Compatible with NVIDIA Datacenter and RTX Pro GPUs in Ampere, Ada Lovelace, Hopper, and Blackwell architectures
Maximum Drives per Array	Up to 32 NVMe SSDs per array (scalable based on hardware platform)
Supported NVMe SSDs	Dapustor, FADU, Hagiwara, Intel/Solidigm, Kingston Technologies, KIOXIA, Memblaze, Micron, Petaio, Phison, QUANXING, Samsung, Scaleflux, Seagate, Western Digital/SanDisk
Maximum Array Size	Up to 1PB per array (subject to hardware and OS limits)
Host OS Compatibility	RHEL 9, Ubuntu 20.04 / 22.04 / 24.04
CPU/Memory Requirements	Minimum 8-core CPU, 32GB RAM (higher specs recommended for large arrays or AI workloads)
Network Support	Up to 800GbE, RDMA-capable (RoCE, iWARP)
Data Protection	Double failure protection, distributed journaling, proactive parity verification
Error Handling	Automated URE recovery, SMART monitoring, predictive failure analytics
Performance Optimization	Read/write performance optimized for AI workloads, NVMe multipath support, parallel rebuilds
Management Interfaces	Web-based GUI, CLI (Command Line Interface), RESTful API
Monitoring & Alerts	Real-time statistics, event logging, SNMP, email notifications
Security	Role-based access control (RBAC), audit logging, secure API endpoints (HTTPS)
Integration	API for orchestration platforms, plug-ins for monitoring tools
Licensing	Perpetual

SupremeRAID™ AE (AI Edition)

Powering AI with Unmatched Performance & Resiliency

Learn more: <https://graidtech.com/products/supremeraid-ae>



Built for the Demands of Modern AI.

SupremeRAID™ AE (AI Edition) is purpose-built to meet the data-intensive demands of modern AI workloads, empowering enterprises to eliminate transfer bottlenecks and fully optimize GPU server performance. It delivers breakthrough speed, reliability, and efficiency—transforming how AI data is stored, protected, and accessed at scale. **Contact us to unlock the full potential of your AI infrastructure: email info@graidtech.com.**



Accelerate Results

GPUDirect Storage (GDS) support ensures faster data throughput and near-zero latency, accelerating model training and inference for AI workloads.



Maximize GPU Performance

The Intelligent Data Offload Engine optimizes your workloads, enabling full utilization of your GPU resources for peak performance.



Protect Critical Data

Our resilient enterprise RAID protection safeguards your datasets, ensuring uninterrupted access during intensive AI processes.



Scale Seamlessly

Flexible deployment options and GDS support allow for smooth scalability to meet your current and future performance demands.

