

VDURA for Artificial Intelligence

Fortify your AI infrastructure with enhanced protection, unmatched durability, and optimized application performance.

Introduction

The race to AI leadership is not just about algorithms – it's also about infrastructure. As organizations push the boundaries of innovation with machine learning (ML) and generative AI workloads, they face relentless demands for faster insights, seamless scalability, and uncompromising reliability. In this era of exponential data growth and compute complexity, ordinary storage systems simply cannot keep up.

Downtime is more than an inconvenience; it's a multimillion-dollar setback. Data corruption doesn't just waste resources; it stalls breakthroughs. What AI leaders need is a solution designed from the ground up to match the pace and precision of their ambitions.

Enter the VDURA® Data Platform – a transformative force in AI infrastructure. By combining industry-leading durability, unparalleled scalability, and intelligent data management, VDURA empowers organizations to overcome bottlenecks, unlock new capabilities, and accelerate their path to groundbreaking discoveries. With VDURA, AI workloads move beyond limitations, delivering on the promise of true innovation.

AI data solutions – evolved

The VDURA Data Platform empowers AI-driven organizations by excelling in five key areas.

1. The gold standard in reliability and durability

- **Multi-level erasure coding (MLEC):** Protects against data loss and ensures data integrity end-to-end.
- **Uninterrupted performance:** Guarantees system availability and rapid recovery, maintaining mission-critical AI operations.
- **Enterprise-grade durability:** Prevents data corruption and reruns, optimizing resources for uninterrupted workflows.

2. Infinitely scalable performance and capacity

- **Parallel file system:** Distributes data across multiple nodes for simultaneous access, eliminating I/O bottlenecks.
- **Dynamic Data Acceleration (DDA):** Enables low-latency, high-throughput performance for generative AI models and other data-intensive workloads.
- **Future-proof scaling:** Seamlessly accommodates growing data needs and supports thousands of GPUs.

3. Optimized GPU Throughput

- **Hybrid storage architecture:** Balances the best attributes from storage technologies to deliver optimal performance and cost efficiencies.
- **Data striping:** Distributes data across nodes, reducing latency and preventing GPU starvation.
- **Flash-optimized metadata engines:** Streamlines data access for complex AI and HPC workloads.

4. Intelligent data management

- **Smarter decision-making:** Enables organizations to search, analyze, and mobilize data precisely where and when needed.
- **Efficient checkpointing:** Saves application states quickly, minimizing disruptions and ensuring fast recovery.

5. Best-in-class TCO

- **Hybrid architecture:** Combines the performance of flash with the cost-efficiency of HDD.
- **Unified management plane:** Reduces IT overhead, allowing teams to focus on strategic objectives.
- **Superior price-to-performance:** Delivers unmatched value compared to all-flash competitors, supporting budget-conscious scalability. A global leader in academic research chose VDURA to support its AI and HPC workloads "for its superior price/performance and durability...which all-flash solutions couldn't meet."

Overcoming I/O bottlenecks and accelerating GPU performance

The unprecedented processing power of GPUs has revolutionized AI and HPC, but performance gains are often hindered by bottlenecks in data storage and retrieval. VDURA's parallel architecture addresses this challenge by:

- **Eliminating I/O bottlenecks:** Parallel read/write operations maintain uninterrupted data streams to GPUs.
- **Enhancing checkpointing efficiency:** Concurrent writes ensure minimal disruption to workflows.
- **Ensuring continuous GPU operation:** High-throughput data streams reduce idle time and maximize computational output.

AI use cases

VDURA partners with industry leaders to transform data into powerful AI outcomes.

- **Generative AI models:** Seamless scalability and low-latency storage support faster iterations and innovation.
- **Predictive analytics:** Reliable, scalable infrastructure processes and analyzes large datasets with precision.
- **Autonomous systems:** Real-time data processing and high availability ensure functionality and safety.

Advanced reliability for mission-critical applications

Checkpointing and data integrity

- **Checkpointing excellence:** Enables efficient recovery and reduces resource wastage during long-running workloads.
- **Multi-level erasure coding:** Ensures data integrity even during node or drive failures, safeguarding critical AI operations.

Scalability for growing workloads

- **Adaptive architecture:** Dynamically adjusts to accommodate evolving AI and HPC demands.
- **Seamless performance growth:** Supports expanding compute and storage requirements with ease.

About VDURA

VDURA elevates AI and HPC infrastructure with a powerful and flexible Data Platform that delivers an unparalleled blend of performance, durability, and reliability. From supporting GPU-accelerated workloads to enabling groundbreaking AI innovations, VDURA is the trusted partner for organizations aiming to stay ahead in a data-driven world.

