

Global Energy Leader Modernizes HPC and AI Infrastructure for Seismic and Renewable Research

Why Choose VDURA?

Scalable Expansion for Next-Gen AI Workloads: Seamless cluster expansion improves IOPS, capacity throughput, and metadata/file stats to accommodate unpredictable research growth.

Linear Scale-Out Performance: Every node adds increased throughput and capacity, eliminating scaling bottlenecks.

Hybrid Efficiency: Flash for active workloads and HDD for capacity are both managed as one namespace.

Proven Support: VDURA single-stack model comes with expert engagement and proactive tuning.

Overview

A global energy and renewables innovator selected the **VDURA Data Platform V5000** to power its next-generation AI and HPC environment for exploration analytics, seismic modeling, and renewable-energy simulation.

The company's prior storage environment based on a **long-standing Lustre-based incumbent** had reached its limits in price/performance, manageability, and efficiency. By moving to VDURA, the organization gained a unified data environment that delivers flash-class performance, petabyte-scale capacity, and linear scalability across all dimensions of growth.

Challenge

The existing infrastructure created operational bottlenecks that slowed data analysis and consumed valuable resources:

- Limited flexibility for mixed I/O workloads across AI and seismic pipelines
- Increasing management overhead and performance tuning complexity
- Inefficient scaling that required separate systems for performance and capacity tiers
- Rising power consumption and floorspace constraints as workloads expanded

The customer sought a modern data platform that could consolidate tiers, reduce complexity, and scale seamlessly for unpredictable, next-generation AI workloads.

Solution

The company deployed the VDURA Data Platform V5000, integrating NVMe flash and HDD capacity storage under a single control plane. This architecture provides high throughput for AI and seismic workloads while maintaining cost-efficient scaling for long-term data retention.

- **Cluster Type:** New HPC/AI built on x86 servers with NVIDIA GPUs
- **Usable Capacity:** ≈ 26 PiB
- **Performance:** >150 GiB/s sustained throughput on seismic workloads that scales linearly as the system grows
- **Architecture:** Tuned for both AI and HPC workloads with small, low-latency workload files on SSD and large high-throughput files on HDD
 - **Director Nodes:** Metadata and orchestration for high parallel throughput
 - **HDD Capacity Expansion Nodes:** Storage Nodes with both SSDs for small file workloads and high-efficiency HDD storage for bulk scientific data
 - **Unified Global Namespace:** One control plane for both parallel file and object-based workloads

Competitors

The customer evaluated several long-standing AI and HPC storage vendors. Traditional Lustre and scale-out file systems offered maturity but could not match VDURA's balance of performance, economics, and simplicity.

Results and Benefits

The VDURA solution has provided significant improvements in every aspect of the customer's data storage needs.

- **Performance:** >150 GiB/s sustained throughput
- **Efficiency:** Significantly lower TCO
- **Scalability:** Linear expansion across IOPS, throughput, capacity, and metadata without re-architecture
- **Flexibility:** Unified namespace supports both AI and HPC data tiers in one platform
- **Reliability:** Erasure coding and online rebuilds ensure continuous uptime

Conclusion

The VDURA Data Platform V5000 provides a future-ready foundation for AI, HPC, and renewable energy innovation that delivers the scalability, efficiency, and durability required to power the next decade of energy discovery.

About VDURA

VDURA delivers the world's most powerful data platform for AI and HPC. Combining breakthrough durability, economics, and performance, VDURA powers the most demanding research and industrial workloads, helping organizations turn data into discovery.