

# Goethe University

## V5000 Deployment

### Use cases

- AI-accelerated physics research (e.g., lattice-QCD, plasma and astrophysics models)
- Traditional HPC simulation and data-intensive analytics
- Continuous training and inference pipelines that demand high sustained throughput across an NVIDIA NDR200 (400 Gb/s) fabric

### Win summary

Goethe University's Center for Scientific Computing, a leading research institution, achieved substantial measurable results by implementing the VDURA Data Platform V5000 for their AI-driven physics and HPC cluster. With one of Europe's largest AMD GPU clusters combined with VDURA's V5000 system, the solution delivers exceptional real-world performance, including 8.4 million IOPs and 670,000 metadata operations per second. As the system grows, performance will scale linearly, effectively meeting future demanding throughput requirements. VDURA's innovative Multi-Level Erasure Coding™ (MLEC) significantly enhances data durability and availability, ensuring robust 24x7 operation.

Professor Volker Lindenstruth, Director of the Center for Scientific Computing at Goethe University, emphasized that VDURA's unique architecture provided an optimal balance of superior price/performance, unparalleled durability, and minimal management overhead:

*"The new VDURA system which is going to be our core storage facility will grow to become larger and larger. Already today it is a very high performing system which is what we need and it's also an extremely reliable system which is of very high importance to us because losing data is not something we are looking forward to."*

See Goethe University discuss VDURA at ISC 2025 [here](#).

## Problem

Goethe University faced the challenge of finding a storage platform that delivered both high performance and long-term durability at an economical price, something all-flash systems could not provide.

## Solution

After a competitive RFP which included the primary incumbent bidder, Goethe University awarded the project to VDURA.

Phase	Configuration	Performance
Phase 1	20 PB VDURA V5000 with ~2 PB NVMe flash + 18 PB HDD; delivers 90 GB/s aggregate bandwidth into the GPU fabric	90 GB/s sustained
Phase 2 (Future)	Expansion path to >100 PB hybrid capacity	2.5 TB/s sustained

Key selection factors cited by Professor Lindenstruth during onsite acceptance testing:

- Multi-Level Erasure Coding provides superior data protection and ensures end-to-end data integrity.
- VDURA architecture that matches flash-tier speed with disk-tier economics.
- True parallel file system (PFS) performance paired with simple, policy-free management.
- VDURA's proven durability of up to 12 nines after experiencing prior data loss events with the legacy system.

Benefit	Impact
Superior price/performance	NVMe performance front end with HDD capacity expansion design avoids the cost premium of all-flash bids while still meeting every GPU throughput SLA
Data durability and reliability	MLEC and rapid rebuilds eliminate the outage risk that had plagued the previous platform
Low operational overhead	Single global namespace, low management resources (one-half FTE to manage system), and ability to scale to future EB capacity while increasing reliability and performance
Research acceleration	Physics AI training jobs that previously stalled on I/O now run uninterrupted, maximizing utilization on Europe's largest AMD-GPU cluster

Professor Lindenstruth summarized the outcome:

*"After a thorough evaluation, we selected VDURA's Data Platform and V5000 for its superior price/performance and durability...which all-flash solutions couldn't meet. VDURA's hybrid architecture delivered the ideal balance of price/performance and low management overhead, making it the perfect choice."*



## Scalability

- Linear capacity growth: The V5000 architecture allows Goethe to scale beyond 100 PB while preserving a single namespace.
- Performance at scale: All-Flash NVMe Storage Nodes ensure small-file IOPS and high throughput that increases as the system scales.
- Future-proof roadmap: VDURA's platform supports both disk-dense and flash-heavy nodes, enabling optimization for future workload mixes without forklift migrations.

## Competitors

The legacy environment suffered from cumbersome upgrade cycles, lacked durability, demanded significant expertise to administer, and consumed excessive resources to stay online. Competing all-flash systems carried a significantly higher cost per petabyte.

VDURA's offering beat competitors due to its best price/performance ratio for AI/HPC workloads, proven durability, and minimal management overhead.