
CUSTOMER WIN

# Large Federal System Integrator

## V5000 Deployment

### Use cases

- Computational Fluid Dynamics (CFD) simulations
- AI model training, inference, and analytics
- Mixed HPC+AI pipelines requiring continuous checkpointing and ultra-low-latency I/O

### Win summary

A Tier 1 U.S. federal system integrator selected the VDURA Data Platform V5000 to power a mission-critical AI and CFD cluster that is the backbone of one of the world's largest defense GPU installations. The Phase 1 rollout combines 4 PB of ultra-fast NVMe flash with 16 PB of HDD capacity, with transfer rates above 800 GB/s today, scaling to ~200 PB and 2.5 TB/s as the program expands. VDURA's unified global namespace delivers consisten low latency, automated flash-to-disk tiering, and multi-level erasure coding, slashing cost-per-TB by greater than 60 percent compared to flash-only solutions, boosting TB-per-watt energy efficiency by 44 percent, and trimming operations to manage the system to one-half an FTE.

### Problem

The integrator needed storage that could simultaneously deliver:

1. Consistent low latency to keep thousands of GPUs saturated.
2. A single, unified namespace that marries NVMe flash performance with petabyte-scale capacity.
3. Automated data placement with no external movers or manual migrations.
4. Dramatically lower cost-per-TB and more efficient TB-per-watt than all-flash alternatives.
5. Enterprise-grade encryption and a minimal operational footprint.
6. Freedom from the reliability issues experienced with the incumbent solution and the higher TCO quoted by other competitors.

VDURA

CUSTOMER WIN

## Solution

| Phase | Configuration | Performance |
|---|---|---|
| Phase 1 (2025) | 20 PB total – 4 PB NVMe flash + 16 PB HDD capacity extensions | Transfer rates above 800 GB/s |
| Future | Incremental growth to ~200 PB usable | Transfer rates above 2.5 TB/s |

## Platform architecture

- Director Nodes for metadata and orchestration.
- All-NVMe Flash Nodes for extreme AI performance.
- HDD Capacity Expansion Nodes for cost-efficient bulk storage.
- Unified global namespace provides one data plane and one control plane.

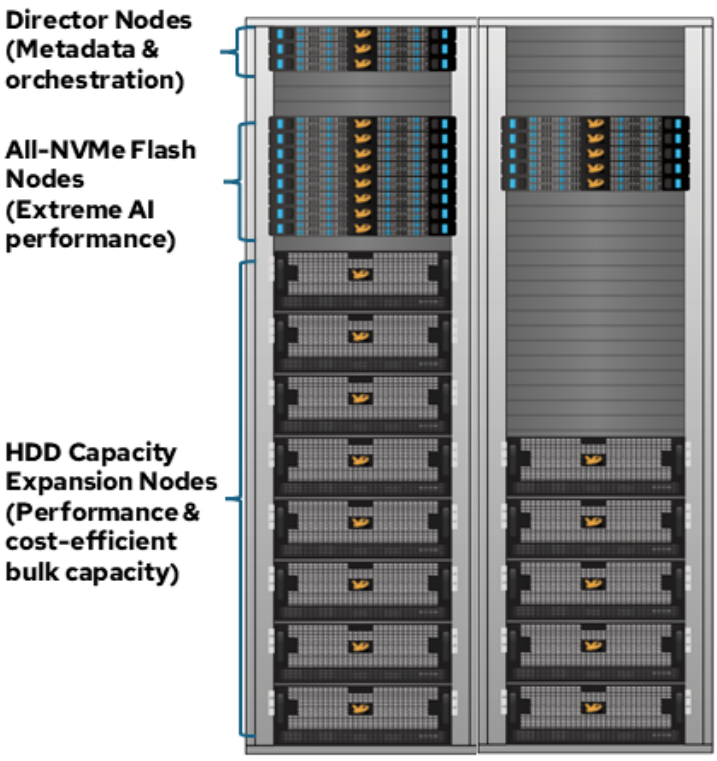


## Key selection factors

- True parallel-file-system throughput with consistent low latency.
- Single namespace spanning flash and capacity tiers.
- No external data movers required for flash-to-disk tiering.
- Greater than 60 percent lower cost per TB versus all-flash designs.
- Built-in AES-256 encryption.
- Only one-half FTE needed to manage system.
- 44 percent better energy efficiency (TB/W) than other all-flash offerings, cutting power and carbon footprint.

CUSTOMER WIN

## Integrator benefits

| Benefit | Impact |
|---|---|
| GPU utilization maximized | Transfer rates above 800 GB/s keep accelerators busy; latency stays consistently low |
| Lower capital cost | Design reduces spend by greater than 60 percent per TB vs. all-flash alternatives |
| Energy and space savings | 44 percent more efficient TB/W lowers OPEX and rack footprint |
| Operational simplicity | Unified namespace and policy-driven tiering requires only 1/2 FTE |
| Security and compliance | Built in end-to-end encryption and automated key management |

## Scalability

- Linear capacity growth to ~200 PB usable without namespace splits.
- Performance scales incrementally to 2.5 TB/s sustained as additional All-Flash and Capacity Expansion Storage Nodes are added.
- Architecture supports future flash-heavy or disk-dense expansions with no forklift upgrades.


## Competitors

- VDURA replaced a commodity-based TLC all-flash solution using a third-party object store that was cited for poor reliability and flash/object storage integration challenges.
- VDURA beat the other top competitor, who offered a QLC-based scale-out file system, by delivering superior total cost of ownership while meeting every performance and efficiency requirement.

## Conclusion

By combining flash-class responsiveness with object-storage economics under a single unified global namespace, VDURA enabled this federal integrator to launch a mission-critical HPC+AI environment that is both budget-aligned today and ready to scale an order of magnitude in the years ahead.