Innovating Defense AI Infrastructure

USE CASE

Overview

A leading U.S. federal systems integrator is building one of the largest GPU-accelerated environments globally, focusing on defense and aerospace. The project is a mission-critical program requiring low-latency performance, massive scale, and resilience.

Challenges

The integrator faced the following challenges:

- Sustaining GPU training and inference pipelines without bottlenecks.
- Eliminating complex manual data migrations.
- Reducing total cost of ownership.
- Enhancing energy efficiency and space utilization.
- Ensuring enterprise-class encryption and security.

Solution

The VDURA Data Platform V5000 was deployed to meet these needs. Key features of the solution include:

- **Unified Global Namespace**: High performance NVMe and scale out capacity tiers.
- **Multi-level Erasure Coding**: Provides robust data protection.
- **Parallel File System Throughput**: Combines all-flash performance with hybrid storage economics.
- **Security**: Built in end-to-end encryption per file basis.

Results & Benefits

- **Accelerated time-to-insight**: GPUs remain fully saturated thanks to ultra-low-latency I/O.
- **Optimized TCO**: Hybrid architecture cuts cost per terabyte by more than half compared to all-flash competitors.
- **Sustainable efficiency**: Improved TB-per-watt reduces both operating costs and carbon footprint.
- **Streamlined operations**: One global namespace, automated tiering, and minimal staffing requirements.
- **Elastic scalability**: Capacity and throughput scale linearly to future-proof national defense programs.

About VDURA

VDURA offers a powerful, cost-efficient, and secure data platform for AI and HPC, supporting organizations in deploying resilient infrastructure at any scale. VDURA powers the most demanding workloads in government, science, and industry, helping organizations unlock the full potential of their data.

© 2025 VDURA, Inc. All Rights Reserved.