


VAST Data Platform for NVIDIA DGX SuperPOD



Start Fast, Deploy With Confidence, And Operate With Ease

Quick Start AI For The Enterprise

The era of artificial intelligence (AI) is well underway, and organizations across every discipline recognize the need to adopt AI or fall behind. However, many organizations struggle with the complexities of building and deploying AI solutions, and need a turnkey solution that can be easily supported by their existing IT teams.

Enter the NVIDIA DGX SuperPOD™ with the VAST Data Platform, engineered to eliminate the guesswork for innovative organizations and accelerate their AI initiatives. As the first enterprise NAS solution certified for DGX SuperPOD, the VAST Data Platform delivers virtually limitless levels of scale and performance via a system architecture that eliminates data tiers and silos. Now enterprise IT teams can deploy large-scale AI that is simpler, faster, and easier to manage.

All-Flash Performance, Archive Economics

The DGX SuperPOD with the VAST Data Platform is the culmination of years of partnership and success with joint customers. VAST's Disaggregated, Shared-Everything (DASE) architecture defies the conventional wisdom that only classic HPC storage based on parallel file systems has the scale and performance required for AI workloads. VAST's inherently parallel architecture provides the scale, performance, and resilience that make standard NAS protocols suitable for the data storage demands of DGX SuperPOD. Furthermore, the DASE architecture enables a set of compounding efficiencies: similarity-based data reduction, low-overhead erasure codes, and support for hyper-scale flash that deliver archive-tier economics on all-flash infrastructure.

Why VAST?

AI Infrastructure Simplified

Start fast with a turnkey, validated architecture; Streamline operations with enterprise NAS simplicity; Scale to exabytes in a single namespace; Optimize costs with all-flash at archive economics

NAS Simplicity

Native NFS, SMB, and S3 cross-protocol support; Infiniband and ethernet support in the same namespace

Manageability

Low touch administration; Multi-tenant isolation with QoS policies; Secure Multi-Tenancy; Ransomware-Proof Snapshots

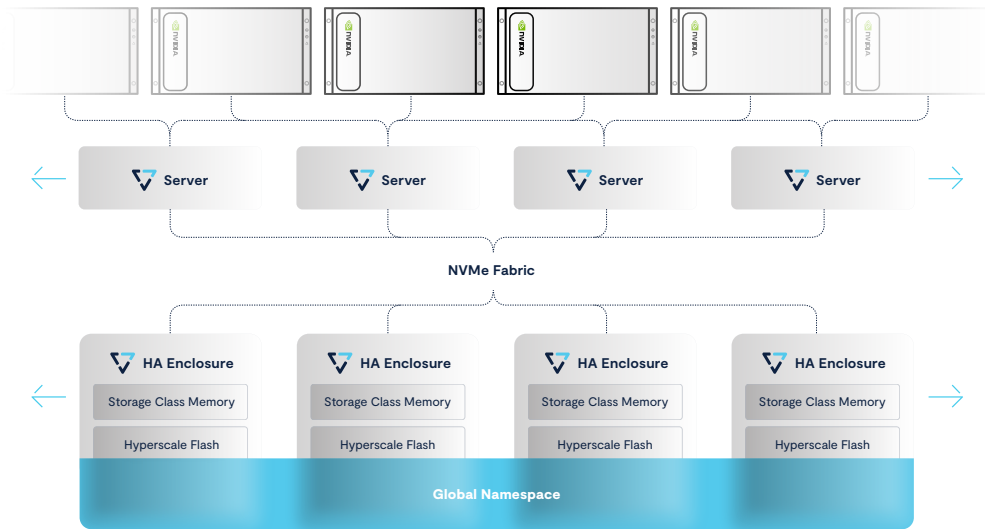
About Us

VAST is the data platform for simplicity at scale. VAST engineered the first new scale-out architecture in 20 years to take you beyond the conventional limits of scale, resilience and cost.

↳ Looking to simplify your data center and unlock insights from all of your data? Contact us at hello@vastdata.com

Parallel File System Performance, NAS Simplicity

VAST Disaggregated, Shared-Everything Architecture



Simplicity For Non-Stop Data Pipelines

VAST is truly differentiated as an enterprise NAS datastore that is simple for enterprise IT teams to manage with operational efficiency that scales with the system. Built for non-disruptive upgrades and expansions, VAST customers enjoy unparalleled uptime and reliability along with support for enterprise features like snapshots, replication, secure multi-tenancy, encryption and more. DGX SuperPOD with VAST Data lets customers focus on AI-driven outcomes while eliminating the need to tune for different workloads.

Trusted by Industry Leaders

VAST is trusted by leading companies around the world to leverage a single radically affordable all-flash innovative platform.



Exabyte-Scale All-Flash Namespace

Millions of IOPs, TB/s of performance from ultra-dense cost-optimized NVMe infrastructure



Non-stop Operations

Upgrades and expansions are always on line and non-disruptive



Archive Economics

Radical data reduction combined with low-overhead erasure codes and support for 10 year endurance on hyperscale flash



Multi-Tenant Training and Inference Infrastructure

Secure tenant and work isolation with QoS policies



NAS Simplicity

One simple to manage scale out file system appliance, no need for complex client-side file system client software



Embarrassingly-Parallel Scalability

With no east-west cluster traffic, VAST's DASE architecture sets a new standard in scaling to the needs of massive ML/DL server farms.