

HPC solutions powered by VAST Data

Reimagining HPC for the AI Era



The changing HPC landscape

High-Performance Computing is rapidly shifting as AI and ML workloads take center stage. These workloads demand real-time data access, massive small-file ingestion, and hybrid-scale performance—far beyond what traditional storage systems were built for.

AI-driven HPC now delivers 10× higher data I/O, while global HPC data volumes are growing 2.5× every four years. From genomics to climate modeling, data intensity is exploding—making modern, high-speed, flexible storage essential for next-generation HPC.

Challenges with traditional **Parallel File Systems (PFS)**

Legacy PFS platforms are increasingly unable to keep pace with AI-driven, data-intensive HPC.



Performance limits: Slow tier access creates latency that stalls modern I/O-heavy workloads.

Complex tiering: Manual/policy-based data movement causes delays, mis-tiering and higher costs.

Version lock-in: Proprietary protocols require tight OS/kernel alignment, creating frequent compatibility issues.

Infrastructure overhead: Multi-tier, multi-OEM deployments and proprietary clients add complexity and management burden.

Traditional PFS can't deliver the agility, scalability or simplicity required for AI-enhanced HPC—demanding a modern, streamlined approach.

HCLTech + VAST Data: Modern HPC, simplified



HCLTech and VAST Data deliver a next-gen HPC platform built for AI-driven, data-intensive workloads. Powered by VAST's DASE all-flash architecture, it provides real-time data access, massive parallelism and a single global namespace that scales from TBs to exabytes—without tiering complexity.

The platform supports NFS, SMB, S3, NVMe/TCP and GPU Direct Storage, enabling in-place data processing for AI/ML pipelines. VAST's DataEngine automates AI and data workflows, cutting operational overhead and accelerating insights.

Combined with HCLTech's consulting and managed services, enterprises get a future-ready HPC foundation with the agility, simplicity and performance today's research and AI workloads demand.

Flexible deployment models

Appliance model

Pre integrated, turnkey platform for rapid deployment

OEM integrated Model

Hardware flexible option aligned with preferred vendors and existing procurement

U4X (Utility for X) model

Fully managed, pay per use HPC delivered by HCLTech—ideal for dynamic or project based workloads

All models are backed by HCLTech's end to end services—design, deployment, optimization and 24×7 support.

Customer benefits



Faster time-to-insight for analytics and simulations



Ransomware protection with indestructible snapshots



50× faster restores for backup and recovery



Global namespace and multi-tenancy for distributed teams



Simplified management with real-time observability and QoS



AI-optimized performance for deep learning pipelines

Together, HCLTech and VAST Data deliver a modern, AI ready HPC foundation that simplifies operations, accelerates insights and scales effortlessly—empowering enterprises to innovate faster and unlock the full value of their data.