

# UNIVERSAL STORAGE FOR DEEP LEARNING



## CONSOLIDATE THE STACK. UNLEASH INSIGHTS.

Recent advances in machine and deep learning are moving GPU driven artificial intelligence applications from R&D labs and Silicon Valley startups into the computing mainstream. AI is presenting a new set of storage challenges that IT organizations have never had to deal with.

Customers need to keep GPU servers fed with more bandwidth than what a NAS can deliver, causing them to think that complex parallel file systems are the only answer. On the other hand, file accesses are small and random in nature, making AI the killer use case for flash storage – but parallel file systems and burst buffers were never designed for providing random read access to petabytes of flash. As AI becomes democratized, niche HPC storage products are difficult for enterprise IT organizations to adopt and operationalize on a broad basis.

VAST combines the single-host performance and scalability of a parallel file system with the simplicity of an all Flash NAS appliance. By marrying the performance of Flash with economics that are otherwise only found in HDD-based archive storage, VAST enables AI applications to consolidate infrastructure and accelerate the training and inference time.

## THE PYRAMID IS DEAD.

For decades, storage practitioners have been trained to tier their storage infrastructure in order to save money and to keep their largest datasets on slow, archival storage. Today, the rise in adoption of new machine and deep learning techniques require training on vast amounts of data, where data needs to be fed to farms of GPUs with maximum throughput. Training algorithms only get more effective as they are exposed to more and more data, thereby rendering the classic storage tiering model obsolete in the AI era.

### VAST'S KEY BENEFITS

#### **Parallel File System Speed & Scale**

Scale to TB/s, millions of IOPS;  
100GB/s+ per AI client over  
Ethernet or InfiniBand.

#### **NAS Simplicity**

Turnkey Universal Storage  
appliance – with no client SW  
dependency or complexity.

#### **Archive Economics**

Radical flash storage economics  
to make flash affordable for all AI  
datasets.



# BREAKING TRADEOFFS: VAST DATA UNIVERSAL STORAGE

By applying new thinking to decades old storage problems, VAST has broken the long-standing tradeoff between storage performance & the cost of capacity to make it possible to simplify and accelerate ML/DL pipelines.



### Exabyte-Scale All-Flash Namespace

Millions of IOPS from cost-effective QLC flash, enabled by 3D XPoint technology.



### RDMA Client Access

NFSoverRDMA provides over 10GB/s of mount performance over Ethernet or IB; scale multiple mounts in a single client to drive 100GB/s+ of single client bandwidth.



### Archive Economics

VAST has pioneered many innovations to democratize flash: QLC flash translation, low-overhead erasure codes and groundbreaking global data reduction.



### Multi-Tenant Training and Inference Infrastructure

VAST server pooling capability provides dedicated QoS for competing applications.



### 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.

**Scales linearly in performance and capacity.**  
**Universal Storage grows with your business needs.**

