VAST

White paper

Choosing the Right Storage for Modern Workloads

Table of Contents

Successful AI Projects Need More Than Just Storage 3

How Enterprise File Storage Got Here 4

The Four Architectures 5

- Clustered Dual-Controller NAS 5
- Shared-Nothing Scale-Out NAS 5
- Parallel File Systems 5
- Disaggregated, Shared Everything (DASE) 6

Solution Comparison 7

Comparison Table 8

From autonomous agents to biomedical breakthroughs, today's most exciting technologies depend on mining intelligence from vast pools of unstructured data. As AI becomes democratized, high-performance file storage is no longer enough on its own. It remains essential, as AI workloads demand it, to keep GPUs utilized, but models and applications behave very differently from traditional workloads. Delivering at scale requires a complete, end-to-end platform. Some vendors simply stop at high-performance storage and leave gaps that organizations need to fill on their own. This report compares common approaches to storage and profiles a representative of each.

Successful AI Projects Need More Than Just Storage

Just five years ago, most enterprise storage roadmaps centered on consolidating block and file workloads, holding the line on cost, and tiering cold data to the cloud. AI was a workload, not the primary workload.

That has rapidly shifted. Training runs now span thousands of GPUs and ingest petabytes of unstructured data. Inference pipelines pull from vector stores, document repositories and live event feeds. Retrieval-augmented architectures push reads to wherever the model needs context. Agents loop through tools, write intermediate state back to a shared store, and call other agents that do the same.

These patterns share a common requirement: a single, high-throughput, low-latency data plane that every GPU, every container, and every analytics engine can hit at the same time. Legacy enterprise storage wasn't designed for this. Even the scale-out systems built for industries like life sciences, energy, and media are showing strain when asked to feed modern AI infrastructure.

This paper compares the architectures contending for that role and the products that represent each:

- **Clustered dual-controller NAS,** represented by NetApp ONTAP
- **Shared-nothing scale-out NAS,** represented by Dell PowerScale and Everpure FlashBlade
- **Parallel file systems,** represented by IBM Storage Scale
- **Disaggregated, shared everything (DASE),** the architecture beneath the VAST AI OS

How Enterprise File Storage Got Here

By the mid-1990s, enterprise IT had standardized on dual-controller NAS appliances. Two heads, shared disk, redundant paths. This was good enough for home directories, archives, and VM datastores. NetApp built one of the most successful businesses in storage on that pattern, and the architecture still anchors many corporate file environments today.

Dual-controller systems work when storage is a supporting actor, but AI is a completely different story. A single training run can consume more sustained read bandwidth than a Fortune 500 data center used to consume in total.

Scale-out emerged in response, and two approaches took hold. Shared-nothing NAS, popularized by Dell PowerScale (formerly Isilon) in the early 2000s, bundles compute and capacity into identical nodes and spreads data across them. Parallel file systems, developed in HPC and best represented by GPFS, layer a POSIX client over a flexible mix of media and protection schemes. Both extend dual-controller architectures, however both inherit limits that show up clearly under AI load.

The newest architecture, DASE, was developed at VAST and sits beneath the VAST AI OS. It treats compute and capacity as independent resources connected over an NVMe fabric, not a backplane.

The Four Architectures

Clustered Dual-Controller NAS

NetApp ONTAP is the canonical example. ONTAP clusters are built from HA pairs (two controllers sharing storage), and clusters scale by adding more pairs. FlexGroup volumes let a single namespace span multiple constituent FlexVols across nodes. That is how NetApp delivers scale-out for unstructured data.

The architecture inherits the strengths of dual-controller NAS: hardened code, deep protocol support, broad ecosystem integration. It also inherits the boundaries. Each constituent volume is owned by one node. Any single file lives on a single node. For AI training, where a checkpoint write can saturate a controller and where dataset loaders fan out reads across millions of small files, those node boundaries become contention points. NAS-only ONTAP clusters also cap out well below the node count modern AI infrastructure asks for.

NetApp's AI positioning leans on the AIPod reference architecture with NVIDIA and the newer AFX and AIDE platforms. NetApp has innovated with these recent releases, but the underlying architecture, by design, remains the same.

Shared-Nothing Scale-Out NAS

Dell's PowerScale (formerly Isilon) is the longest-running shared-nothing NAS in the market. Everpure's FlashBlade is the all-flash version of the same idea in a chassis-and-blade form factor. Both bundle compute and storage media into one building block. Data is replicated or erasure-coded across nodes, and the cluster protects itself by spreading work and data widely.

Shared-nothing is operationally simple and scales linearly for many workloads. The trade-off is rigid. Compute and capacity scale together, even when only one is needed. Adding capacity means buying CPU you may not want. Adding performance means buying capacity you may not need. Inter-node traffic also grows with cluster size, which limits how large any single cluster can usefully get.

PowerScale's feature surface is mature after two decades, including cross-pool tiering and tier-to-cloud. FlashBlade is the more modern of the two and supports S3, NFS, and SMB, however files and objects sit in separate namespaces. Both architectures were set when "modern workload" meant a render farm, not a 10,000-GPU training cluster.

Parallel File Systems

IBM's Storage Scale (formerly GPFS, then Spectrum Scale) is a leading commercial parallel file system. It originated in HPC, where many compute nodes write to a single file at the same time, and where POSIX semantics matter more than NAS protocols.

Storage Scale supports nearly every data protection model in use: shared-nothing, SAN, software RAID over shared media. That flexibility is a strength when an in-house team has the expertise to design and tune the system. It is a liability when the goal is a turnkey AI platform that a single infrastructure engineer can run without years of parallel file system experience.

For AI shops with deep HPC heritage, Storage Scale still has its place. For the broader enterprise running mixed AI, analytics, and traditional workloads, the configuration surface is wider than most teams want to manage.

Disaggregated, Shared Everything (DASE)

DASE is the architecture foundation beneath the VAST AI OS. The protocol servers that handle NFS, SMB, S3, and the higher-level database and pipeline services run as stateless containers. The storage media (Storage Class Memory and Hyperscale Flash SSDs) sits in highly available enclosures. The two layers are connected over an NVMe fabric with single-digit-microsecond latency.

Because the front-end servers don't own data, they can scale independently of capacity. Add containers when you need more throughput. Add enclosures when you need more capacity. Every server can see every byte, so there is no node ownership, no cross-node fan-out for a single file, no FlexGroup-style constituent boundaries.

DASE is what lets VAST collapse all-flash file, object, database, and pipeline services onto one platform. The VAST AI OS extends that platform with a native database (DataBase), event-driven function execution (DataEngine), a global namespace (DataSpace), and agent and inference services (InsightEngine). Storage is the foundation. The AI OS is what sits on top.

Solution Comparison

- **NetApp ONTAP,** (clustered dual-controller NAS) is one of the most widely deployed enterprise NAS platforms in the world. ONTAP runs on the AFF and ASA all-flash arrays, in the major public clouds (FSx for ONTAP, Cloud Volumes ONTAP), and at the edge. For AI, NetApp leads with AIPod reference architectures co-designed with NVIDIA and more recently the AFX and AIDE platforms. The architectural ceiling for AI is the HA-pair boundary, which constrains single-file performance and the practical size of a single cluster.

- **Dell PowerScale (formerly Isilon),** is a 20-plus-year-old shared-nothing scale-out NAS. It supports heterogeneous clusters of all-flash, hybrid, and HDD nodes, and includes mature features like automated tiering between PowerScale pools and to the cloud. Its limits are the limits of shared-nothing: compute and capacity scale together, and inter-node traffic grows with cluster size.

- **Everpure FlashBlade,** packages the shared-nothing model into a chassis of all-flash blades. The newer FlashBlade//EXA utilizes FlashBlade as a dedicated metadata server which expands into higher-performance environments. FlashBlade products support S3, NFS, and SMB, however files and objects occupy independent namespaces, which complicates multi-protocol pipelines.

- **IBM Storage Scale (formerly GPFS and Spectrum Scale),** is a parallel file system first released in 1998. It gives storage architects extensive flexibility in protocol, media, and protection design. That flexibility creates a complex planning and operations surface as features like NFS or S3 access require gateways to scale, and teams typically need deep HPC expertise to run it well.

- **VAST AI Operating System (DASE),** is built on the DASE architecture. It delivers all-flash file, object, database, and pipeline services from a single namespace, with N+4 data protection at as little as 3% overhead and similarity-based data reduction that consistently exceeds competing systems. VAST positions the platform as an AI OS because storage, database, function execution, and global namespace services share the same data plane.

Comparison Table

	NetApp ONTAP	Dell PowerScale	Everpure FlashBlade	IBM Storage Scale	VAST AI OS
Class	Clustered HA-pair NAS	Scale-out NAS	Scale-out NAS / object	HPC parallel file system	AI data platform
Architecture	Clustered dual-controller	Shared-nothing	Shared-nothing	Shared-nothing, shared-media, SAN	DASE
Application focus	Enterprise NAS, AI	Media, life sciences, AI	AI, backup, media	HPC	AI, HPC, analytics, media, backup
POSIX client	No	No	No	Yes	No
Scale	Hundreds of TB per HA pair; ~24-node NAS cluster	11 TB – 185 PB raw	123 TB – 19.2 PB	4 TB – 8 EB	338 TB – EB+
HDD / Hybrid	Yes / Yes	Yes / Yes	No / No	Yes / Yes	No / No
All-flash	Yes (AFF / ASA)	Yes	Yes	Yes	Yes
Low-cost hyperscale flash (QLC)	Limited	Yes	Yes	Yes	Yes
Data protection	RAID-DP, RAID-TEC	N+1 to N+4 (most N+2)	N+2+	3- and 4-way mirror, N+2, N+3	N+4
Data protection overhead	12-25%	11-25%	35-50%	20-75%	3-20%
Storage pools	Aggregates per HA pair	One per node type minimum	One	Several	One
Asymmetric scaling (compute and capacity independent)	No	No	No	Partial	Yes
Symmetric scale-out	Yes	Yes	Yes	Yes	Yes
Internal tiering	Within ONTAP (FabricPool)	Scheduled	No	Scheduled	No

	NetApp ONTAP	Dell PowerScale	Everpure FlashBlade	IBM Storage Scale	VAST AI OS
Tier to external storage	Yes (FabricPool)	Yes	Yes	Yes	No
Deployment on-prem	Yes	Yes	Yes	Yes	Yes
Deployment in-cloud	Yes (FSx, CVO)	Yes	No	Yes	Yes
Snap to cloud	Yes	No	No	No	Yes
Tier to cloud	Yes	Yes	By file	Yes	No
NFS	Yes	Yes	Yes	Yes	Yes
SMB	Yes	Yes	Yes	Yes	Yes
S3	Via StorageGRID	Yes	Yes	Yes	Yes
Multi-protocol SMB / NFS	Yes	Yes	Yes	Yes	Yes
Multi-protocol file / object	Limited	Yes	Separate namespaces	Yes	Yes
Inline duplicate block elimination	Yes (efficiency varies)	Inline (managed jobs still required)	No	No	Inline
Single-block compression	Yes	Inline on current models	Yes	Yes	Inline

	NetApp ONTAP	Dell PowerScale	Everpure FlashBlade	IBM Storage Scale	VAST AI OS
Global, cross-block (similarity) compression	No	No	No	No	Inline
Native database service in platform	No	No	No	No	Yes (VAST DataBase)
Native function / pipeline service	No	No	No	No	Yes (VAST DataEngine)
Global namespace across regions	Limited (FlexCache, SnapMirror)	Limited	No	Yes (AFM)	Yes (DataSpace)

Figures are directional and reflect publicly available product information as of July 2026. Vendor documentation should be consulted for current limits and capabilities.

©2026 VAST Data, Inc. All rights reserved. All trademarks belong to their respective owners.
07/26 I 12364