

Hammerspace Data Platform

Unify, automate, and provide high-performance standards-based access to unstructured data across any site, any cloud, any storage

DATA SHEET

Data Silos are Holding Back Your Business

The next data cycle is here. AI, GPU computing, data analytics, and other digital transformation initiatives are forcing a long overdue reckoning to fundamentally change how data is used and preserved. Data is no longer consumed primarily by humans, it is used by machines at a pace legacy systems weren't designed for.


Data-driven organizations are being held back by a **storage-centric** infrastructure that leaves valuable unstructured data trapped in storage silos.

Getting data to high-performance local compute clusters, GPU clouds, remote applications, and decentralized global users with the speed and efficiency needed, is a challenge. Data copy sprawl impacts data governance and security. Historically, storage systems have not provided the agility to use cloud resources seamlessly, and existing IT architectures can not feed GPU clusters for AI training, inferencing, and analytics.

Benefits

- Unify Unstructured Data to Eliminate Data Silos
- Keep GPU and Compute Resources Fully Utilized
- Move Data Between Sites and Clouds Quickly and Non-Disruptively
- Simplify Data Governance, Security, and Protection

Hammerspace Solves with a Fundamentally Different, Data-Centric Architecture



Eliminate Data Silos and Deliver Extreme Parallel Performance to Compute, Applications, and Users Across Any Site, Any Cloud, Any Storage


The Hammerspace Data Platform unifies unstructured data across sites, clouds, and any storage into a single, parallel global filesystem with a global namespace. It provides HPC-class performance for high speed data processing, and automates data placement in the appropriate storage tier, cloud region, or cloud service nondisruptively with automated Data Orchestration.

This eliminates data silos and makes data an instantly accessible resource for compute clusters, applications, and users, no matter where they are located.



A Truly Global, Standards-Based Parallel File System

Spans multiple sites and multiple clouds to present a single global namespace. Store data on both new and existing storage from all leading vendors.



HPC Performance with Enterprise NAS Simplicity

Parallel NFSv4.2 combines the benefits of HPC file systems with standards-based NFS to accelerate AI data pipelines and GPU computing.



Unify Existing Unstructured Data Into a Single Data Set

Make existing data instantly accessible to users with Data-in-Place Assimilation. Files are visible to users in minutes.



Non-Disruptive Multi-Site Data Orchestration

Automate data placement and protection to give users, applications, and compute clusters fast access to global data.



Objectives-Based Data Policies and Programmable Metadata

Define plain-language policies that control data placement and protection, via GUI, API, or CLI.



Advanced Cross-Platform Data Services

Data storage and protection services for simplifying data security and governance, across sites, clouds and storage systems.



Hammerspace Architecture

Architectural Attributes	Description
Parallel Global File System	Spans multiple sites and clouds, presents as a single global namespace, and provides ability to read/write files simultaneously from multiple sites.
Separate Metadata and Data Paths	Metadata layer, or control plane, separated from the data layer or data plane to provide for faster performance, data placement flexibility and independence.
Deploy on Bare Metal, VM, or Cloud	Deploy as software on bare metal, in VMs, and on cloud instances in AWS, Azure, and Google Cloud.
Multi-Protocol File Access	Supported protocols include NFSv4.2, NFSv3, SMB 2, SMB 3, and S3. In addition, Hammerspace is NVIDIA GPUDirect Storage certified and a CSI driver is available for use with containers. All data in the file system is accessible from any protocol.
Parallel NFS with Flex Files	Uses NFSv4.2 client with Parallel NFS and NFS Flex Files - all part of the standard Linux kernel - to deliver a standards-based parallel file system architecture using NFS.
No Custom File System Clients or Agents	Delivers parallel file system access to underlying storage volumes without the use of proprietary client software, agents on the storage, or custom NFS implementations.
Direct Data Path Between Client and Storage	With NFSv4.2, create a direct data path between client and storage, with metadata out of band, to reduce latencies and network transmissions.
Linear Scalability to Thousands of Nodes	Scale from one storage node to thousands of storage nodes without experiencing performance degradation.
Live Data Mobility	Move, copy, and manipulate files in the file system concurrent with file read/writes. File movement is transparent to users and applications, even on files that are open and in use.
Store Data on Any Storage Type	Store and tier data on NVMe flash, SSDs, HDDs, and tape, all within a single file system.
Support for Third Party Storage	Supports any storage type from any vendor, including off-the-shelf servers, NAS, block, object, cloud, and tape.
Ability to Use Existing Storage	Use existing storage systems from any vendor, and assimilate metadata from those systems non-disruptively and with near zero downtime.
Supports Wide Range of Networking Technologies	Supports both TCP and RDMA transport for data, with either Ethernet or Infiniband.



Hammerspace Enterprise Data Services

Data Orchestration Services	Description
Data-in-Place Assimilation	Assimilate metadata from third-party storage systems, while keeping data in place. Files are visible to users in minutes, non-disruptively and with minimal downtime.
Objectives-Based Data Policies	Define plain-language policies that automate data placement and protection. Policies can be set via API, CLI, or via GUI using a simple policy builder.
File-Granular Data Services	Any policy or file motion can be applied at a share, directory, and even a file-granular level.
Custom Metadata Tagging	Administrators may apply custom metadata tags and labels.
Automated Metadata Inheritance	Authorized users may assign custom metadata tags to any folder within the file system hierarchy so new files and folders below that point in the hierarchy will automatically inherit those tags.
Metadata Triggers	Policies can be triggered by any combination of metadata, including file system metadata and custom metadata.
Automated Data Services	Data services can be automated globally at a file-granular level across all storage types and locations, including data protection services (snaps, clones, replication, etc), and data mobility actions (tiering, migration, etc.)
Cross Platform Tiering	Multi-site replication can be automated using Hammerspace policy or on-demand via user/application activity. Supports one-to-one, and one-to-many schemes.
Data Storage and Protection Services	Description
Multi-Site Replication	Multi-site replication can be automated using Hammerspace policy or on-demand via user/application activity. Supports one-to-one, and one-to-many schemes.
Erasur Coding	High-performance erasure coding that is built on the Mojette Transform erasure code. Highly resilient and efficient, and up to 2x faster than traditional erasure coding schemes.
Snapshots and Clones	Share-level snapshots across multiple storage types, which may be scheduled or immediate. Snapshots are consistent across all sites, and may be stored anywhere, including in the cloud. Where snapshots apply at the share-level, clones are at the file and directory level. Snapshots are read-only, where clones are writable and can be moved.
Deduplication and Compression	Data stored in public cloud storage is automatically deduplicated and compressed for faster replication and lower bandwidth and capacity usage.
Encryption	Supports encryption utilizing 3rd-party key management servers (KMS) as well as passphrase encryption.
File Versioning	File versioning can be enabled to occur automatically as a declarative Objective at a file-granular level. Versioning may be triggered anytime when a file or dataset that has such a versioning policy assigned to it is changed.
WORM	Block I/O to specific tiers and lock instances of files for data protection and to comply with immutability requirements.
Undelete	Undelete may be enabled for files as an additional layer of protection, with time-stamped versions kept for a specific time range in the snapshot.
Anti-Virus	Scanning files for anti-virus is supported for on-access and background scanning. Hammerspace will scan files on access and prevent file opens if a virus is detected.

