HAMMERSPACE

WHITEPAPER

Hammerspace Global Data Environment

Local access to global data everywhere

Table of Contents

What is a Global Data Environment?

A Global Data Environment (GDE) is defined as a means by which users and applications can get the experience of 'local' access to data that may be stored in a decentralized cloud across widely distributed storage types/locations, while at the same time providing global control for data services transparently across them all. The data environment could include multiple otherwise incompatible storage silos in a data center or across multiple data centers, and may include one or more Cloud vendors.

When data and storage are managed in a GDE users and applications no longer face the difficulties of finding and accessing their files across often incompatible storage silos and/or locations. Such problems become even more costly when IT managers must wrangle multiple storage types, point solutions and file copies globally, and must do so without interrupting access to users and applications.

The key ingredient to creating an effective GDE is in the ability to globally leverage file system metadata and other information into a common layer that bridges incompatible storage infrastructures. This metadata layer must be abstracted from the physical storage devices and locations, so it can provide a unified view and control of all data across any storage resource of any type, from any vendor, in any location. Such a cross-platform metadata control plane can then drive workflows across any storage type/vendor anywhere, and do so transparently to users and applications

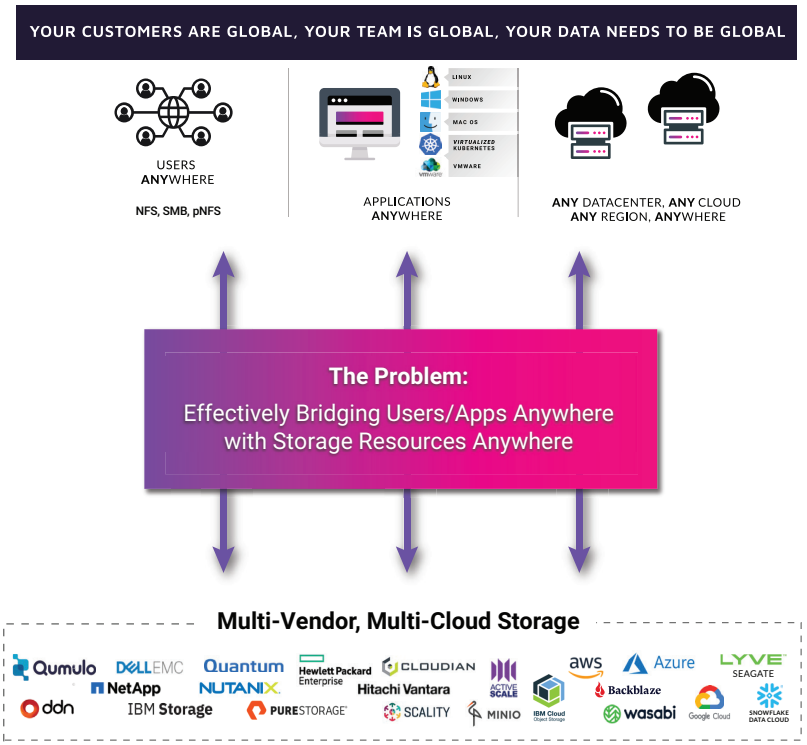

The Root of the Problem:

The problem a GDE addresses is despite the dramatic increase in performance, density, and capabilities in the storage industry over many decades, the paradigm users and applications still use today to access their data is in many ways the same as when their files were stored on 512k floppy disks and 10MB hard drives on individual workstations.

Users and applications interface with their data via a filesystem, which is an abstraction that organizes the raw bits on any storage medium via metadata into a file form that people can understand and use.

There have been multiple revolutions in how filesystems are managed over time. Initially the file system was part of a computer's operating system, with the disks being direct-attached to the PC. As such, the PC was an individual island. Data and its metadata had to be copied or sent outside of that system and then reloaded onto another computer for it to be shared.

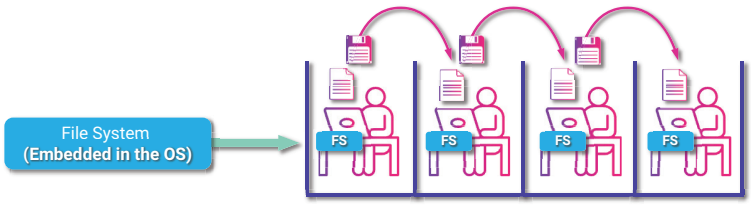

Figure 2: When the file system was trapped at the OS level, users had to make physical copies of their files, an early version of silos.

Then NetApp created a revolution by pulling the file system metadata out of the computer's operating system into a network attached storage device that could share file access to users throughout a local workplace. Now instead of users having to put copies of files onto floppy disks to physically hand them to others, multiple users on a local area network could all work off of the same file share.

This network effect was expanded when Isilon pioneered the next revolution by enabling the file system to emerge out of a single storage controller and to span multiple nodes and racks. That innovation simplified the expansion of data by making file access usable at scales that went far beyond what was previously possible.

The Silo Problem Emerges

This trend has obviously exploded over the last couple of decades, with an entire industry of many diverse storage types emerging at different performance and price bands, all building on the ability to present their file system across networks within a local environment.

The common thread across all of these technologies even until today, however, is that users and their applications still must interact with a file system that presents raw bits into usable information in a form they can consume.

The problem is that despite the innovation in other areas, this key file system metadata remains embedded in each vendor's proprietary storage platform. If that storage device fills up, or the files need to be moved to another storage type or location, both the data, or file essence, and its metadata must be copied and written to another file system on another device. To access this new file copy, users have to remount another file system, and find a second copy of the file metadata and the file essence.

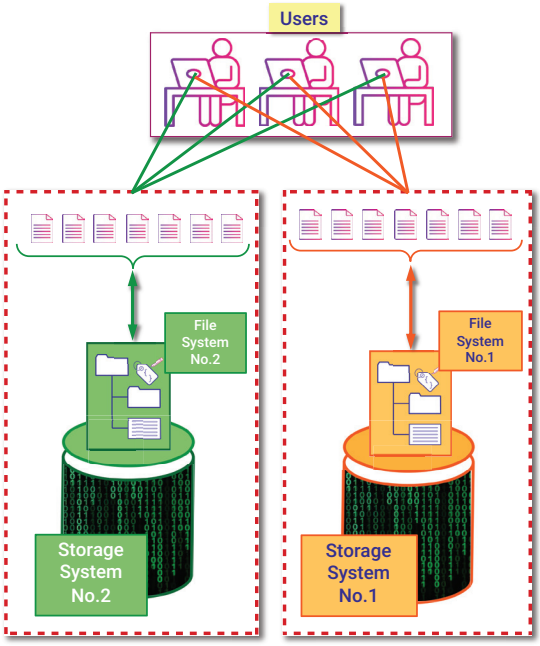

Figure 3: Although file systems expanded to the network, they are still embedded in the storage platform. So moving data to different storage creates copies. Although much faster, it still perpetuates the silo problems of when we were putting files on floppy disks.

When data volumes were small, it was an inconvenience, much like putting copies on a floppy disk and handing them to someone else. Inconvenient, but it was manageable. Indeed, an entire industry of data management point solutions emerged to help bridge the storage vendor silos that emerged.

But even until today, these point solutions are all attempts to overcome the fundamental problem that the file system is trapped in a vendor-locked storage layer, requiring users and their applications to jump from one to another to gain access to all of their digital assets.

Silos Go Global

As the cloud era emerged, enterprises realized they could move from managing their own data centers to outsourcing some of their IT needs to public cloud providers.

Here's the rub: The traditional paradigm of the file system trapped at the storage layer was inconvenient within a single data center. But migration to the cloud dramatically compounded the problem. And for enterprises with large volumes of unstructured data, it became difficult to move all of it entirely to the cloud. This led to more point solutions, gateways, and other complex ways to often manually orchestrate copies across ever greater silos.

Local access with standard SMB and NFS protocols is still needed for most use cases, especially for high performance

and large files. Cloud resources were beneficial as a means to scale businesses beyond fixed infrastructure, but the problem of distance created another more difficult silo barrier.

As data volumes continued to explode and users/applications and their compute/storage resources became more decentralized, the silo problem that had been an inconvenience has now become a severe business limiter. Far from reducing costs, distributed data silos typically add significant operational costs for managing data across both on-premises and cloud resources.

Tackling the Problem: Creating a Global Data Environment

While storage devices can orchestrate data internally within their own resources, they cannot manage data or enable file access outside of their proprietary vendor silo without the use of gateways, symbolic links or other point solutions. But in a GDE, by aggregating metadata from all storage types into a high performance global file system above the infrastructure layer, users, applications, and IT data services are now able to take advantage of a decentralized cloud to see and manage all resources globally, regardless of the which underlying storage type the data lives on today, or needs to move to tomorrow.

In this way, a GDE strategy enables the experience of 'local' access to datasets across any storage type for users, applications, and compute environments located anywhere in the world. This ability to aggregate metadata into a global file system that is independent of the data orchestration and services layer is critical, and distinguishes a well-implemented GDE from environments that simply copy files between silos, or leave symbolic links to files that have physically moved to other storage locations.

A 'global' solution that only works within a single vendor's storage ecosystem is just a silo by another name.

That brute-force copy methodology is often referred to as a 'storage-centric' approach to data management.

As an alternative to this, rather than working on forked copies of data, users and applications need direct access to their files via a true cross-platform global file system. In this way they can share access from anywhere to files that may live on any storage type anywhere. Importantly, such a solution needs to work within the realities of today's applications, network protocols, and across all storage vendor platforms.

The ability to bridge storage vendor silos via a shared metadata layer is called a 'data-centric' approach. And it is the data itself and this shared metadata that provides the intelligence to drive data policies across any storage type, from any vendor, globally.

For a GDE to be effective, it must also be based upon existing standards. Users and applications need to access files via industry-standard protocols (container environments, NFS, pNFS, and SMB), and the GDE implementation must not require proprietary software running on the client, or need agents to be installed on the storage. A well-designed GDE embraces open source throughout its entire architectural stack wherever possible to leverage standard network interfaces and kernel-level code built into standard operating systems for performance and broad compatibility.

In addition, a GDE must be able to run anywhere (aka software-defined) and have the flexibility to run on bare metal, in VMs, or in the cloud.

In the content to follow, we will outline Hammerspace's approach to creating an effective vendor-neutral GDE.

Using Hammerspace Software to Build Your Global Data Environment

Local Data Read/Write for Any Application, Any User, Any Compute Region, Anywhere

To be competitive in today's data landscape, organizations increasingly need to ensure that a distributed workforce can be productive with immediate, secure, and shared access to any data, regardless of their location. Whether teams are working from home offices in the same city, or from different countries around the world, they all need direct secure access to the company's data resources for research, development, and innovation.

The fact is, whether intentional or not most organizations have global data requirements to some degree, and have many of the attributes of a global data environment that they typically manage manually, often with multiple point solutions, and a lot of effort and complexity.

These increasingly decentralized data environments may be caused by the number of workers who must now access their files from remote locations. Or the fact that most data centers have more than one class of storage, with storage silos being created to accommodate different performance or cost bands. Or maybe backups or archives are being pushed to the cloud or object stores or other cool/cold tiers.

The problem is that today managing data across all these storage choices requires significant manual orchestration by IT administrators, and often interruptions or complexity for users or applications to locate and access those files. So enterprises are left to wrangle global data problems with a collection of manual processes and point solutions, and usually must do so with a bottoms-up storage-centric approach using tools designed to manage each individual storage platform locally.

Hammerspace enables customers to implement a Global Data Environment to solve these problems with a top-down data-centric approach. It does so with metadata-driven automation to achieve the accessibility and data protection requirements globally across all storage resources from any vendor, and in any location. It starts from the premise that since data is accessed and stored globally across a myriad of storage choices; so shouldn't those data and those storage resources be managed globally as well?

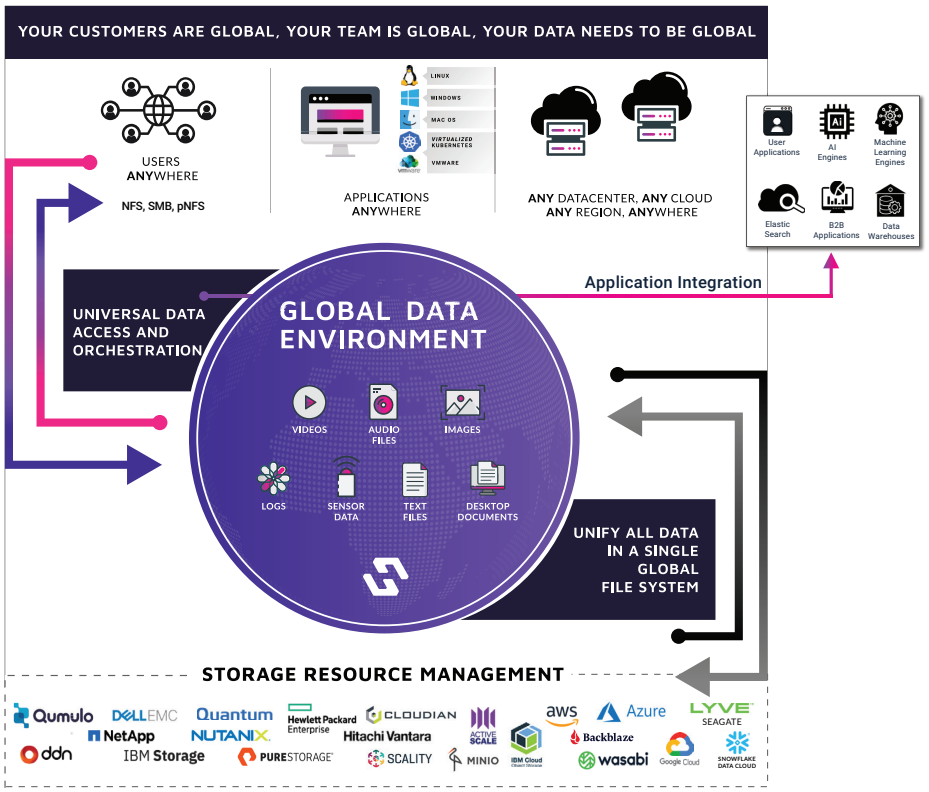

Hammerspace - The Missing Link

Hammerspace is a software solution based upon the principles noted above so that storage silos are effectively eliminated, and users and applications share a unified view and control of all data globally, exactly as if all those files were

on local storage exposed through a single local file system.

Hammerspace has spent years of development, and done the heavy lifting needed to completely reimagine from first principles the way file systems need to work in a decentralized environment. It has done so to enable the promise of the cloud to be fully realized, and in so doing solve these fundamental silo issues within the data center as well.

Just as network file system access in the 1990s revolutionized how businesses could use their data in a single location, Hammerspace's Parallel Global File System now enables access across multi-vendor storage silos, locations, and cloud resources to support use cases that were previously impractical or impossible to achieve.

In this way, all users and applications can achieve local read/write performance to all of the organization's data, which may in fact be stored across different vendors' storage platforms, such as data center DAS and SAN/NAS clusters, object stores, and even in multiple cloud storage platforms.

When a user or application needs to access files that may physically live on remote storage, Hammerspace presents them with unified access to the files via standard network shares, as though they were on a local NAS. This is not shuffling copies of data across storage types and locations, which adds confusion to users, and creates headaches for IT admins. This is globally accessing the same files, via a universal metadata control plane that intelligently bridges the underlying physical storage resources of any type and location.

This is an important distinction: A point solution may be able to copy a file from Site A to Site B, but that is not much different than putting a file on a floppy disk, and physically handing it to your colleague in the 1990s. The user has copied both the file metadata and the file essence, and has created a new redundant, forked copy that must be managed.

With the Hammerspace Parallel Global File System, the metadata layer is common across all users, anywhere. If a local instance of a file is needed for processing in one part of the world, the metadata is still the same. It is the same file, wherever it is being accessed.

Unprecedented Flexibility

Bridging the asynchronous distance gap between locations, or on-premises and cloud with a high-performance Parallel Global File System enables our customers to rapidly ramp up or down resources anywhere. And they can do so without the penalty associated with re-tooling fixed on-premises infrastructures or managing access and data orchestration by point solutions across the incompatible silos. They can even get more life out of their existing compute and storage resources because the boundaries between otherwise incompatible devices disappear.

Storage capacity and performance of any vendor's platform become a variable that can be factored into automated workflows, driven by objective-based policies. Even the cost profiles of different regions from the same cloud vendor can be factored in when workflow decisions are made.

Reduce Complexity and Costs

This capability also has the benefit of dramatically reducing complexity for IT administration in multi-siloed environments, since data orchestration and data services are now back-end functions and are completely transparent to users and applications. In Hammerspace, users and IT Admins can use business rules to establish objective-based policies to ensure file-granular control over file copy management as well as other data management and file protection services. Such tasks are typically managed by numerous point solutions in today's siloed environments.

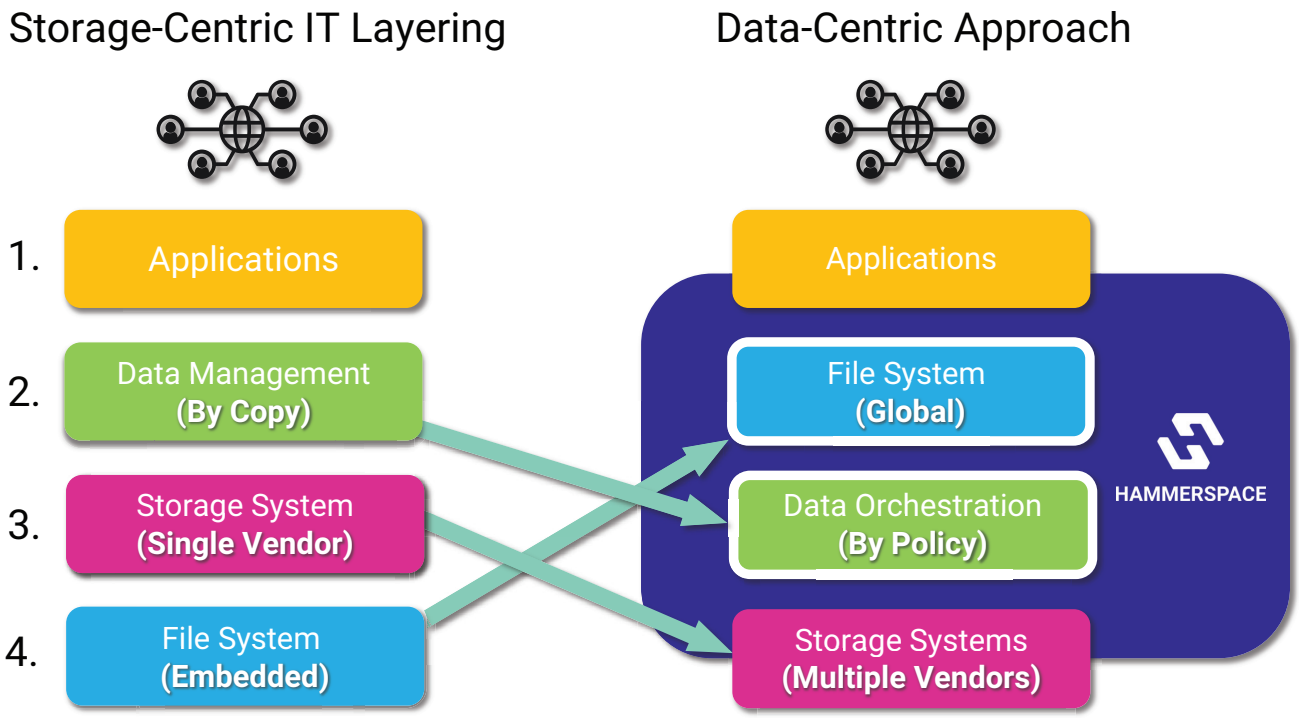

Figure 5: Hammerspace fixes the broken IT layering, elevating the file system out of the infrastructure layer to be close to users and applications. In this way, users can access their data globally across any vendor storage silo or location. And data orchestration and services may be automated across them all, without user interruption.

In this way, users or Admins can establish workflow-driven policies based upon these service-level objectives to ensure the data is in the right storage location at the right time to meet latency, performance, and cost requirements.

This is possible because of the power of global metadata, which enables users and applications to interact directly with all their data globally across all tiers, data centers, and multiple clouds. When data needs to be physically moved, Hammerspace stages only the subset of files that are being used between storage resources. From the user's perspective, the file is still accessible at the same share, in the same file system with no change, since they're interacting with the shared Parallel Global File System across all resources and locations.

This capability eliminates the problem of redundant copies, manual replication, or fragmented data protection strategies and other symptoms of data and storage sprawl. All data services are built into Hammerspace software to automate processes for IT Administrators. This reduces the number of copies of data, while also reducing the number of software applications and point solutions that are required to manage a multi-silo data environment.

Hammerspace Capabilities Overview

Hammerspace is a software solution that empowers customers to create their own Global Data Environment, leveraging any combination of their existing storage resources with new and/or cloud-based storage. In this way, organizations can solve today's challenges for distributed data and remote workers. It is a software-defined solution that may be deployed on commodity bare metal servers, or any virtual environment, and in the Cloud. It supports virtually all storage types from any vendor, including most public and private Cloud solutions.

- **Any Datacenter, Any Cloud, Any Region, Anywhere**. Hammerspace enables organizations to easily store, protect, and operate on data by automatically moving it to the best location by policy or on demand. In this way, customers can not only access data storage anywhere as though it were local, but also take advantage of remote cloud-based compute resources, dynamically taking advantage of the lowest cost infrastructure.
- **Applications Anywhere.** Applications can access data stored in remote locations leveraging automated file-granular data orchestration capabilities to provide high-performance local access when needed for processing.
- **Users Anywhere.** People are increasingly working from all parts of the world. Organizations seek to grow their talent pools with access to team members no matter where they reside. Hammerspace via its Parallel Global File system presents local access to users anywhere, and eliminates the need to replicate a full copy of data volumes to each site.

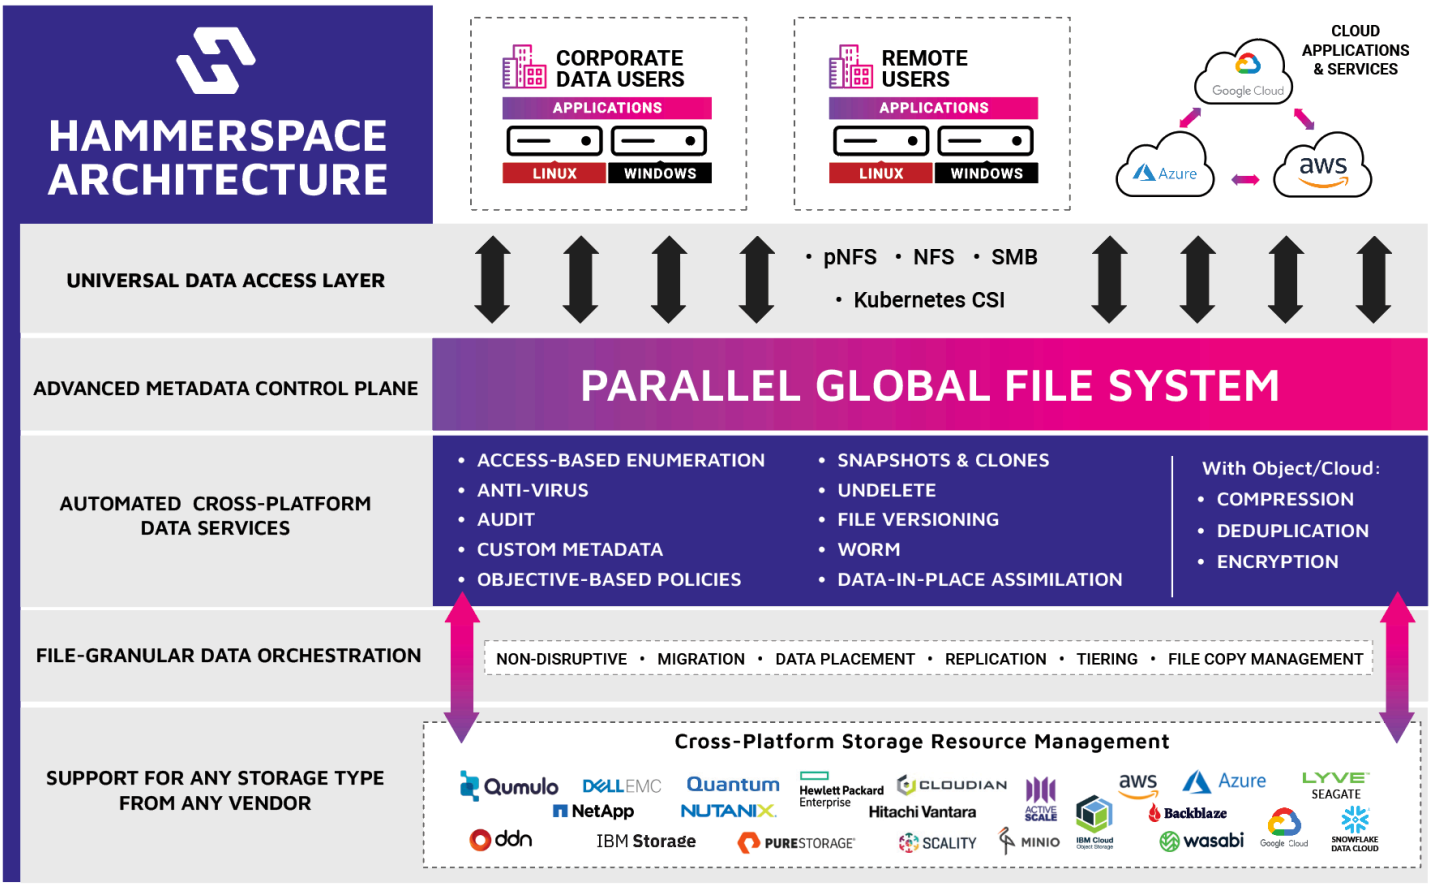

Figure 6: A logical view of the Hammerspace Architecture. Hammerspace software is deployed as a fully integrated software solution built upon open standards, with a single installer that includes all components necessary for deployment on bare metal servers, VMs or in cloud-based compute environments. No external software dependencies are required, including the Linux OS.

The Building Blocks: Hammerspace software is composed of 5 distinct capability groups

- Universal Data Access Layer
- Parallel Global File System
- Automated Cross-Platform Data Services
- File-Granular Data Orchestration
- Vendor-Neutral Storage Support

1. Universal Data Access Layer

The Universal Data Access Layer is the on-ramp to Hammerspace, presenting multi-protocol file access for users and applications. All users/apps see the same file system view, based upon their permissions, regardless of which protocol they are using. Authenticated SMB users see the same folder/file structure that authenticated NFS users do across all underlying storage, including NAS, object, and cloud storage.

Additionally, Hammerspace provides a CSI (container storage interface) driver supporting both block and file-based persistent volumes in Kubernetes and other container environments.

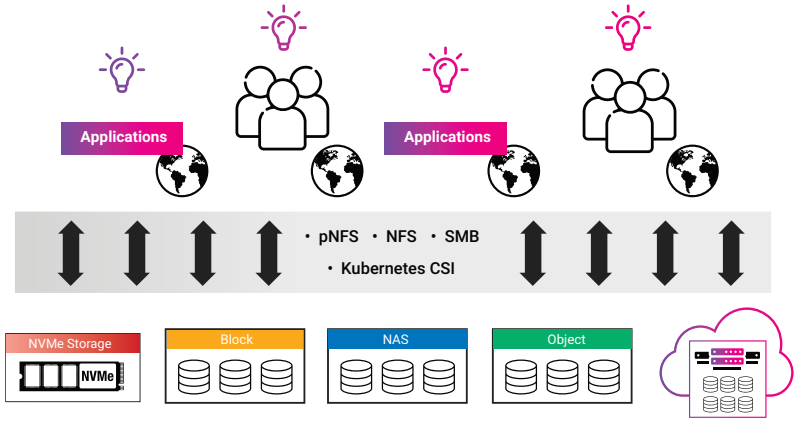

Figure 7: The Hammerspace Universal Access Layer presents a global view to all files on all storage types and locations via standard protocols, as though all were on a local NAS device.

Data is presented to users via file shares exactly as they are used to, with industry-standard protocols. From a user's perspective, it is a hierarchy of directories and files just like any local NAS or file server. But unlike conventional storage platforms, the file system metadata is elevated above the infrastructure layer and does not need to be kept in the same disk silo a the contents of the files.

2. Hammerspace Parallel Global File System

As noted above, unlike conventional storage platforms that embed the file system within the infrastructure layer, Hammerspace elevates the file system so it can span otherwise incompatible storage silos across one or more locations including the cloud.

In traditional storage architectures where the file system is embedded in the storage platform, if files need to be moved to another storage type or location, a copy of both the file metadata and the file essence must be sent. That action now creates a second, forked copy of the file that must be later reconciled.

Because the Hammerspace Parallel Global File System is independent of the storage layer, the need to wrangle such forked file copies is no longer necessary. With Hammerspace, all authorized users in all locations have read/write

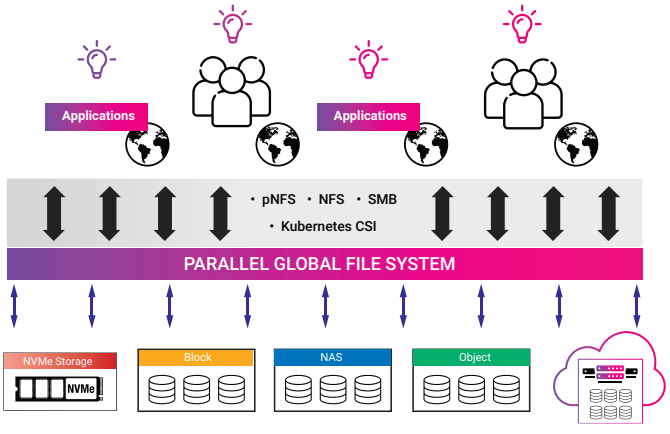

Figure 8: The Hammerspace Parallel Global File System elevates file system metadata above the infrastructure layer, spanning all storage types from any vendor, and across multiple locations, clouds, and cloud regions.

access to their data everywhere. Not to file copies, but to the same files via this unified global metadata control plane.

In addition, Hammerspace supports custom metadata tagging, which adds an additional level of control to globally manage digital assets and to automate workflows. This makes it easy to better describe, classify, and manage the orchestration of file data based on your organization's business needs, associating files with specific projects, cost centers, etc.

This metadata tagging includes automatic metadata inheritance of custom tags, to ensure files are appropriately classified with critical information, but without needing to rely on user action to ensure it actually gets done consistently. This powerful ability is unique to Hammerspace, providing value for a number of use cases, including locality of data, data migrations, data protection, disaster recovery, active archiving, burst-to-cloud rendering jobs, simulations, data analytics, and more.

3. Global Automated Cross-Platform Data Services

Hammerspace provides file-granular global data services across all local and remote storage resources, leveraging its metadata-driven policy engine or via user action or on-demand. File-granular services means individual files, or sets of files, can be managed by policies triggered by any metadata attribute including standard POSIX metadata such as the file names, creation dates, modify times, file types, in addition to custom metadata tags.

Hammerspace global data services enable companies to manage their digital business assets in ways that were previously impractical, or even impossible, due to price, performance, and platform incompatibility challenges. Because these data services can be applied globally across all storage resources, the implementation of global control via Hammerspace eliminates the need for IT organizations to manage multiple point solutions to migrate, protect, or \perform other functions, as is typically the case in siloed environments today.

Hammerspace provides:

- **Audit of file system operations -** Hammerspace supports System ACLs across both SMB and NFS, to create an audit log of file system operations such as file/folder deletes, renames and other actions.
- **Anti-virus –** Automatically scan files on-access and in the background; preventing access if a threat is detected. Integrated with industry-leading anti-virus software using the ICAP protocol.
- **Compression & Deduplication –** 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.
- **Replication –** Multi-site replication can be automated using the Hammerspace objectives-based policy engine or simply done on-demand via user/application activity. Protect data globally across multiple locations and into the cloud for policy-based copy management, to implement data redundancy or to reduce latency when data is accessed remotely. File-granular replication ensures that only the files you require at your remote locations are replicated. Once the initial replication is complete, only the file changes are replicated thereafter. And because the file metadata is consistent across all instances via the Parallel Global File System, replication does not proliferate orphaned copies. All instances are maintained in sync, are consistent, and managed globally within the same file system.
- **Snapshots and clones –** Centrally manage global snapshots across otherwise incompatible storage resources, at the share or file level, throughout their life cycle with granular controls. Offload snapshots to the cloud to reduce

Tier-1 storage needs and increase the resiliency of data. Supports snapshots across multiple storage types globally for up to seven years to meet long-term compliance needs without sacrificing granular recovery requirements. Hammerspace supports file-level data recovery, and snapshots may be stored on any storage including object storage and/or the cloud to reduce costs. Snapshots are user-accessible through snapshot directory and Windows VSS (Volume Shadow Copy Service). Snapshots can be recovered in-place or to a new location.

- **Storage Resource Management –** Hammerspace provides a view into the performance, capacity utilization, and availability of the storage systems within the global data environment. The consolidated view and control of all storage systems simplifies IT management and helps organizations more easily visualize their entire data environment to make proactive resource management decisions. Gone are the days of emergency migrations, or reactive capacity management headaches.
- **Tiering (autonomic) –** Hammerspace makes cross-platform tiering a programmable function that ranges from basic policies to sophisticated workflows. You can set tiering policies based on any metadata attribute including file types, file names, creation date, modified data, file size, etc. as well as custom metadata triggers. Without disrupting user or application access, dormant data can be automatically tiered to low-cost storage, ensuring, for example, that only active files are utilizing expensive high-performance storage. Policies may be set to move certain files to remote locations that have additional capacity available, or to restrict the movement of data out of specific geographies for compliance purposes. Tiering can be automated across different storage infrastructures on-premises or in the cloud based on performance, cost, location, or any other metadata attribute. Tier to any storage type, including block, file, object, and cloud. Data is moved without disruption to user or application access, and in real-time to meet business service level objectives, and to avoid potential bottlenecks. And of crucial importance, all data movement is free of proprietary hooks, symbolic links, stubs or other vendor lock-in tactics required by legacy tiering solutions.
- **Workflow automation –** Hammerspace enables the implementation of policy-driven workflows for individual files, groups of files, file types, active files, inactive files, etc., used for collaboration across multiple sites, distributed workforce, local or remote rendering, cloud burst, simulations, analytics and more. This enables just-in-time data placement, to conserve storage resources, and eliminate unnecessary copies.
- **Undelete –** Protect data globally from accidental, intentional, or malicious deletes and to increase the resilience of the data. Policies may be set to store undelete data automatically on any storage type or location. For example, undelete data could be routed to the cloud to reduce capacity pressure on Tier-1 primary storage. In case of accidental or even purposeful deletion, undelete provides data owners with the ability to recover the files exactly as they were just prior to deletion, or prior to corruption by physical or logical causes, ransomware, etc. Files can be recovered easily by users or administrators.
- **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. This provides the ability to roll back to a very fine level of granularity to a previous version for productivity reasons, recovery purposes, and can also help to mitigate ransomware attacks.
- **WORM –** Create a policy to retain data that needs to be immutable. Supports WORM, read-deny, and delete-deny for protection from any data changes.

4. File-Granular Data Orchestration Layer

By elevating the file system above the infrastructure layers, Hammerspace can provide file-granular data orchestration to move data live and non-disruptively between silos in a data center, across multiple data centers, in addition to one or more cloud providers and cloud regions.

Hammerspace's objective-based policy engine enables resource allocation to automate data placement based upon business objectives.

Such business rules can be based upon any combination of POSIXmetadata, plus custom metadata tags. When coupled with included machine learning-driven automation, Hammerspace enables you to tune your data environment to get the best utilization of storage resources globally while providing seamless access to users and applications.

5. Vendor-Neutral Storage Support

Hammerspace is a software-defined solution that was designed to support any storage type, including NVMe, SSD, hard disks, and which can include block, file, object and cloud storage platforms from virtually any vendor. No special integration is required for storage that uses standard protocols.

Hammerspace supports all the major cloud vendors, including AWS, Azure, GCP, Seagate Lyve, Wasabi, Snowflake, and more.

Hammerspace can utilize the available capacity of your existing storage infrastructure, reducing cost and time-to-value, and even extending the life of current storage, to defer or even eliminate the need to expand your Tier-1 storage resources.

You can get up and running with Hammerspace incrementally, without having to acquire any new storage infrastructure. When new storage resources are needed, Hammerspace can access the new storage of any type to seamlessly add it to your Hammerspace global data environment without disruption to users or applications.

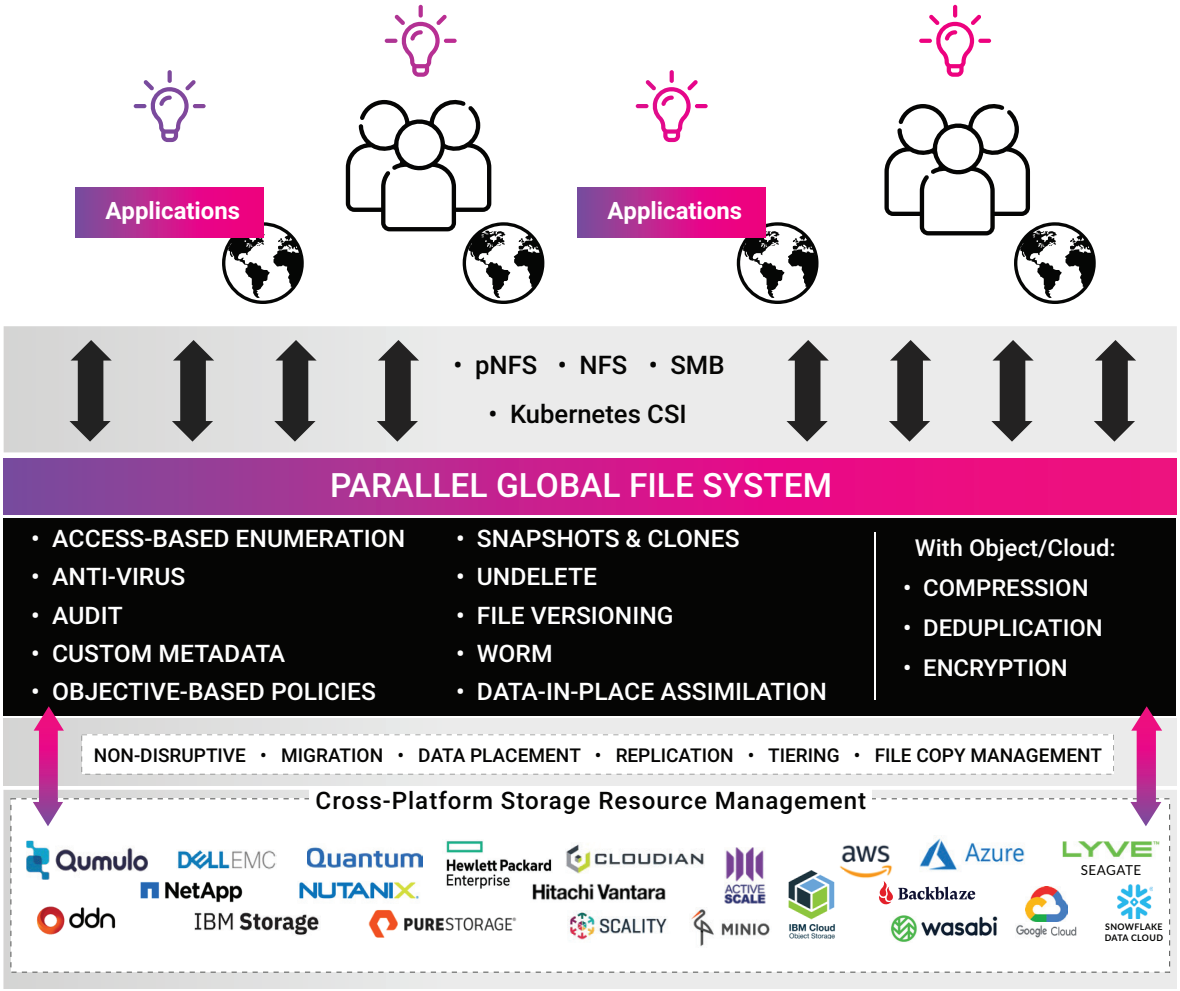

Benefits of Hammerspace

While these above Hammerspace capabilities greatly benefit siloed environments in a single data center, Hammerspace specializes in decentralized environments, where data may need to span two or more sites, and possibly one or more cloud providers and regions.
It creates a global data environment by providing unified data access in a single global file system

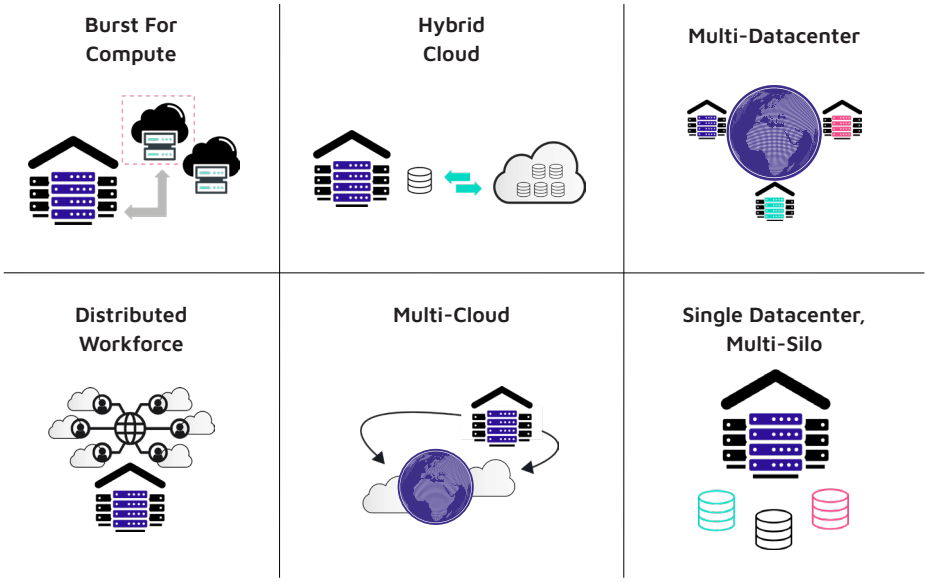

Key Use Cases

Burst for Compute

Video rendering, electronic design automation, genomics pipelines, seismic processing, and other high performance workloads often require thousands of compute cores per project. In many cases, fixed infrastructure becomes cost prohibitive for such jobs, or difficult to stand up or alter for shorter runs.

With Hammerspace, cloud-based compute resources can be rapidly provisioned to do multiple high-performance job runs in parallel, without the limitations of fixed on-premises infrastructure. Since the Hammerspace Parallel Global File System spans both on premises and cloud, this means that jobs can be routed dynamically to whichever cloud region is the most affordable. This is all transparent to users and applications.

As important as rapid provisioning can be, the ability to rapidly decommission cloud compute infrastructure is in some ways more important, as it enables significant savings in third-party application licensing and cloud compute charges. Since the Hammerspace global file system spans both on premises and cloud, this also means that the cloud instances can be rapidly decommissioned without worry of data being orphaned. From a user perspective, all data is where it should be. But from a cost perspective, unnecessary cloud compute and licensing costs are kept at a minimum.

Hybrid Cloud

The vast majority of unstructured data is seldom accessed, and is largely dormant. And yet in most environments, much of this data is still stored on expensive Tier-1 storage systems to ensure it is accessible when needed. Hammerspace enables customers to seamlessly extend their on-premises primary storage environments to the cloud with no interruption to user access, and without the need for gateways, or other proprietary bottlenecks.

Data can be placed transparently to cloud, or other off-premises resources automatically, based upon business rules as set out in easy-to-navigate objective-based policies.

This automates the tiering of data between storage resources in a datacenter or to cloud/object storage easily, intelligently, and quickly, without users even being aware. No changing mount points to new file shares. No fragile proprietary symbolic links left behind, or other old-style HSM techniques.

All users and applications still see all files as they expect in the same file system structure they are used to, regardless of where the files have actually been moved to on the back end. Objective-based policies can be set to tier data based on activity and organizational requirements, whether that needs to happen after one hour, a couple of days, six months, a year, or more.

Distributed Workforce

The way businesses work today has profoundly changed, with many companies no longer requiring their employees to work from office locations. However, providing remote access for employees to have a unified view to all of an organization's network shares is extremely challenging, as data is typically stored in multiple data silos and legacy systems.

Hammerspace makes network shares visible and accessible to anyone anywhere as though they were sitting next to local storage at the data center. This is done using the Parallel Global File System, and leverages automatic file-granular replication to move remote users' files geographically closer to them when needed.

Hammerspace also simplifies IT administration, enabling Admins to globally set up policies so applications and users can access all data in the global data environment. This global control of data policies and orchestration may be monitored and adapted as needed to changing requirements and resources through multiple administrative tools via the GUI, Admin CLI, Hammerspace Toolkit, and its REST API.

Multi-Data Center, Multi-Cloud, Multi-Silo

As described throughout this paper, Hammerspace was purpose built to create a Global Data Environment, eliminating the status quo of a "siloed data environment," where files are locked into storage silos, or a single data center or cloud region. By empowering customers to break free of such silos, Hammerspace provides the ability to scale-across data centers, public and private cloud regions, and multi-cloud to provide the missing link to enable decentralized users, data, compute and storage to truly interoperate effectively.

Ransomware Mitigation

A frequently seen ransomware technique involves creating a new and encrypted temporary file and then deleting the original data. Hammerspace's undelete feature automatically saves a pristine copy of deleted files in the background, making it easy for companies to recover the most recent version of their files as they were prior to the attack.

For ransomware attacks that corrupt, but do not delete data, Hammerspace's file versioning stores the latest versions of files where there is a change, so that customers do not lose any of their data. A ransomware attack that may start to encrypt active data is mitigated by the ability for organizations to turn back the clock to a previous unaltered version of the file prior to the attack.

Hammerspace data services which include snapshots, undelete functions, WORM, and file versioning together create multiple overlapping and complementary protection capabilities to help mitigate and recover from ransomware attacks. In this way, they can be tailored to best suit the individual requirements of our customers globally across all storage resources.

Summary

As noted throughout this White Paper, Hammerspace has been designed from the ground up to solve the problems caused by fragmentation of data across silos in the data center, and increasingly across distributed environments that may span multiple data centers and the cloud.

The transformation into the cloud era has accelerated decentralization of the enterprise in most verticals, bringing into focus the impact of file system fragmentation at the storage layer. And the problem has been additive:

- To be competitive as well as to manage through the changes caused by the pandemic, companies needed to accommodate an increasingly distributed workforce. Storage and data access also needed to be accessible and distributed, while supporting performant workloads at a distance.

- And then to be agile in a dynamic environment where fixed infrastructure may be difficult to acquire due to supply chain issues, the ability for companies to rapidly burst to cloud compute and storage resources has become essential to their survival.

- At the same time, enterprises needed to bridge their existing infrastructure out to these new distributed resources in a way that was cost effective, reduced IT complexity, and that could thus promote greater productivity.

Hammerspace revolutionizes management of data and storage in a world where data cannot be locked into a single vendor's storage ecosystem. It enables organizations to use their existing storage resources plus pick the best of breed from any storage platform, whether local or in the cloud, to create an automated and scalable Global Data Environment.

In this way, Hammerspace provides the immediate benefit of enabling businesses to apply an effective 'data-centric' approach to managing and protecting their digital assets across any and all storage choices. With Hammerspace, these various storage types now become resources that are managed dynamically in the background, according to business priorities for accessibility, performance, cost, protection, and workflow requirements.

No longer do businesses need to be burdened by the complexity, disruption, and costs of a 'storage-centric' approach, where the data gravity associated with a particular storage type or location creates the classic silo problems noted in this paper.

Using the combination of Hammerspace's high-performance Parallel Global File System, with file granular global data services, Hammerspace automates data access and orchestration, and provides global support for any storage infrastructure. In this way, Hammerspace provides local access with global control and data services for all digital assets, across all storage types and locations, while reducing the management complexity for IT staff.

To keep up with the reality of decentralization, a new global paradigm was necessary that effectively bridged the gaps between on-premises silos and the cloud. Such a solution required new technology and a revolutionary approach to lift the file system out of the infrastructure layer, and enable the next wave of decentralization of business in a global economy. It is a revolution as important as lifting the file system out of the operating system was in the 1990s.

This is the Hammerspace innovation.