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 solution to provide users and applications with the experience of 'local' access to data that may be stored across widely distributed storage types/locations, while at the same time providing global control for data services transparently across them all at the infrastructure level. This could include multiple otherwise incompatible storage silos in a data center, or perhaps across multiple data centers, and even may include one or more Cloud storage offerings. A GDE directly addresses the problem experienced by most organizations where users and applications need the experience of local, read/write access to all of their file data across all silos and locations based upon their permissions. And IT managers equally need the ability to manage all their storage resources globally, without interrupting access to users and applications, and without being overwhelmed with the complexity of silo-based point solutions for data services.

The key ingredient to creating a GDE is in the ability to leverage the multiple metadata types inherent in all data and storage to create a global metadata-driven control plane. This control plane is abstracted from the physical storage devices, so it can provide a unified view and control of all data across any storage resource of any type. 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.

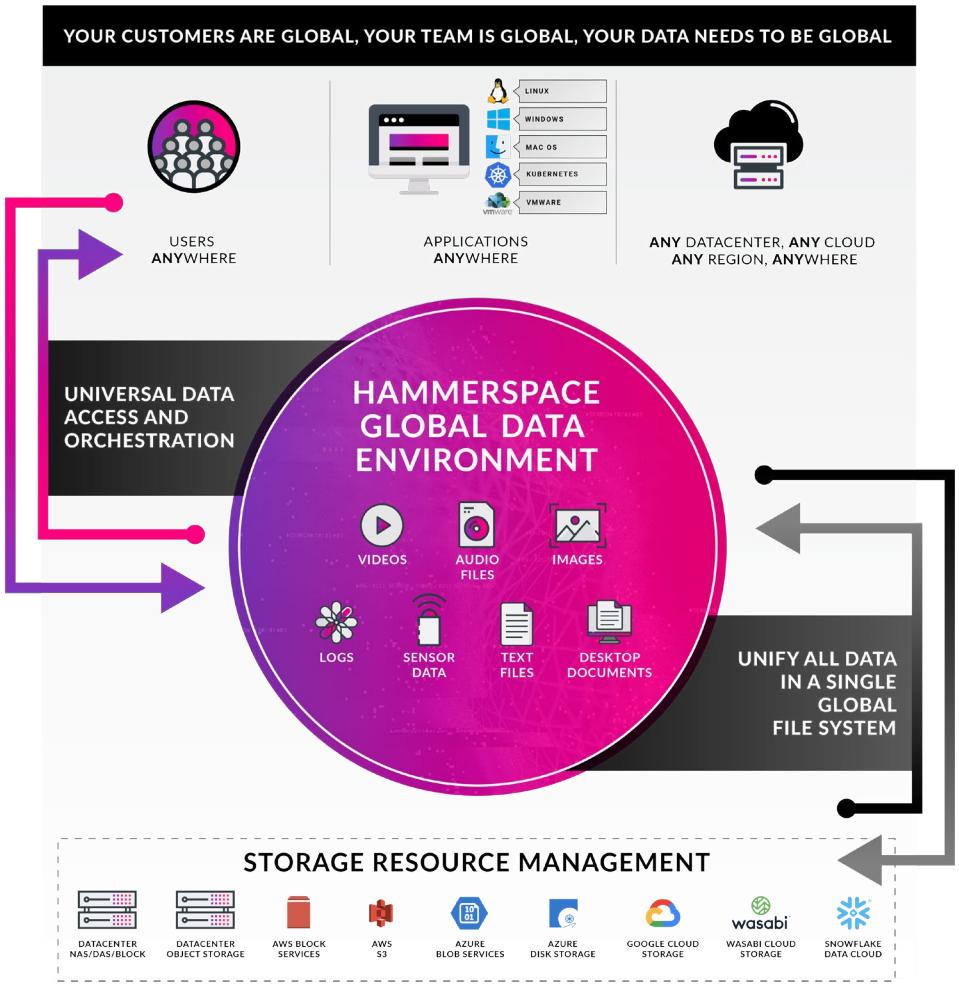

In traditional storage architectures, the metadata that users and applications interact with is typically contained within a local file system, and becomes visible/actionable to them via the standard file or object protocols of each individual silo, such as NFS, SMB, or S3. Storage devices can orchestrate the data internal to their own resources, but cannot, for example, manage data or enable file access outside of that vendor silo without the use of gateways, symbolic links or other point solutions. But in a GDE, by aggregating metadata from all storage resources into a globally accessible metadata control plane, users, applications, and IT data services are now able 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 universal 'local' access to data for users, applications, and compute environments located anywhere in the world. This ability to aggregate metadata into a global control plane 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.

This intelligent metadata aggregation enables the users to have direct access to data stored on any existing storage type – even those that are by definition incompatible – as though the files were local, in a single distributed data environment. Existing DAS, NAS, object storage systems, and cloud block, file, and object services from any vendor can all be leveraged into this architecture, and with access to those data unified through data-in-place metadata aggregation into the GDE.

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 propriety 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 strategy wherever possible to leverage standard network interfaces and kernel-level code built into standard operating systems.

In addition, a GDE must be able to run anywhere (aka software-defined) and has 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 a 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 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 manage manually. Perhaps it is increasing numbers of workers who must access their data from remote locations. Or the fact that most datacenters have more than one classes of storage, with different storage types needed for different performance or cost bands, which applications and users need global access to. Or maybe backups or archives are being pushed to the Cloud or object stores or other cool/cold tiers. The problem is that today all of these elements require significant manual orchestration by IT administrators, and often interruptions or complexity for users or applications to locate and access those files. So enterprises are increasingly facing global data problems with a collection of manual processes and point solutions to manage global data access and protection, but are doing so from the bottom up with tools designed to manage each individual storage platform locally.

Hammerspace implements its Global Data Environment to truly implement a top-down, datacentric approach to solve these problems. It does so with metadata-driven automation to achieve the accessibility and protection requirements of data 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? As such, Hammerspace is a software solution based upon the principles noted above to achieve these datacentric goals, 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.

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 solutions. When a user or application needs to access files that may physically live on remote storage, Hammerspace presents them unified access to the files as though they were local, with the low latency performance needed for read and write operations. This is not just 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.

This is an important distinction: A point solution may be able to replicate file copies from Site A to Site B, and users on each side can see a mirror of the same files. But these are file copies, not the same file. Each of those copies has its own metadata, so they are effectively forked copies. Changes by one user are independent of the other, and must be manually reconciled.

With Hammerspace's Global Data Environment, this complexity is eliminated. Hammerspace aggregates file and storage metadata into a global file system, which then presents multi-protocol access via that universal metadata control plane shared by all users and applications across one or multiple locations. So the SMB user in location A looking for the file would see the same file system view as the NFS user in location B, regardless of which storage the file actually lives on today. The data itself could move by policy from one system to another, or to one or multiple Clouds, but the file metadata is the same, and users/ applications are not even aware it has moved. When users from any location access that file via the global file system, it is live across all via that common global metadata layer. No more synchronizing multiple copies of files that may have proliferated across silos or locations. In the same way that multiple users accessing a local NAS share from a single device see the same metadata and files, multiple users accessing the Hammerspace GDE are seeing the same metadata and files, except those files could be on any storage type in any location.

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 can use business rules to establish objective-based policies to ensure file-granular control over file copy management or multiple data management and file protection services that are typically managed by numerous point solutions in today's siloed environments for workflow provisioning, or multiple data management and file protection services that are typically managed by numerous point solutions in today's siloed environments.

"I can easily connect my team to the data they need, no matter where the individual is located or where the data is stored."

In this way, users or Admins can establish workflow-driven policies based upon these 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, datacenters, and multiple clouds directly. When data needs to be physically moved, only the data that is needed is moved transparently 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 because they're interacting with the global metadata control layer.

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 the GDE software to automate processes for IT, and reduce the number of copies of data while also reducing the number of software applications and point solutions required to ensure all data on all storage types and locations are protected and stored in the most cost-effective manner.

Hammerspace GDE Capabilities Overview

Hammerspace technology delivers a software-defined Global Data Environment so organizations can solve today's data accessibility challenges for distributed data and remote workers. It is a software-defined solution which may be deployed on commodity bare metal servers, 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, to access compute resources, take advantage of the lowest cost infrastructure, and make files locally available to distributed workforces.
- **Applications Anywhere.** Applications can access data stored in remote locations while using automated orchestration tools 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 automates and optimizes the movement of data, and eliminates the need to replicate a full copy of the files at each site, enabling the enterprise to work more efficiently and cost-effectively.

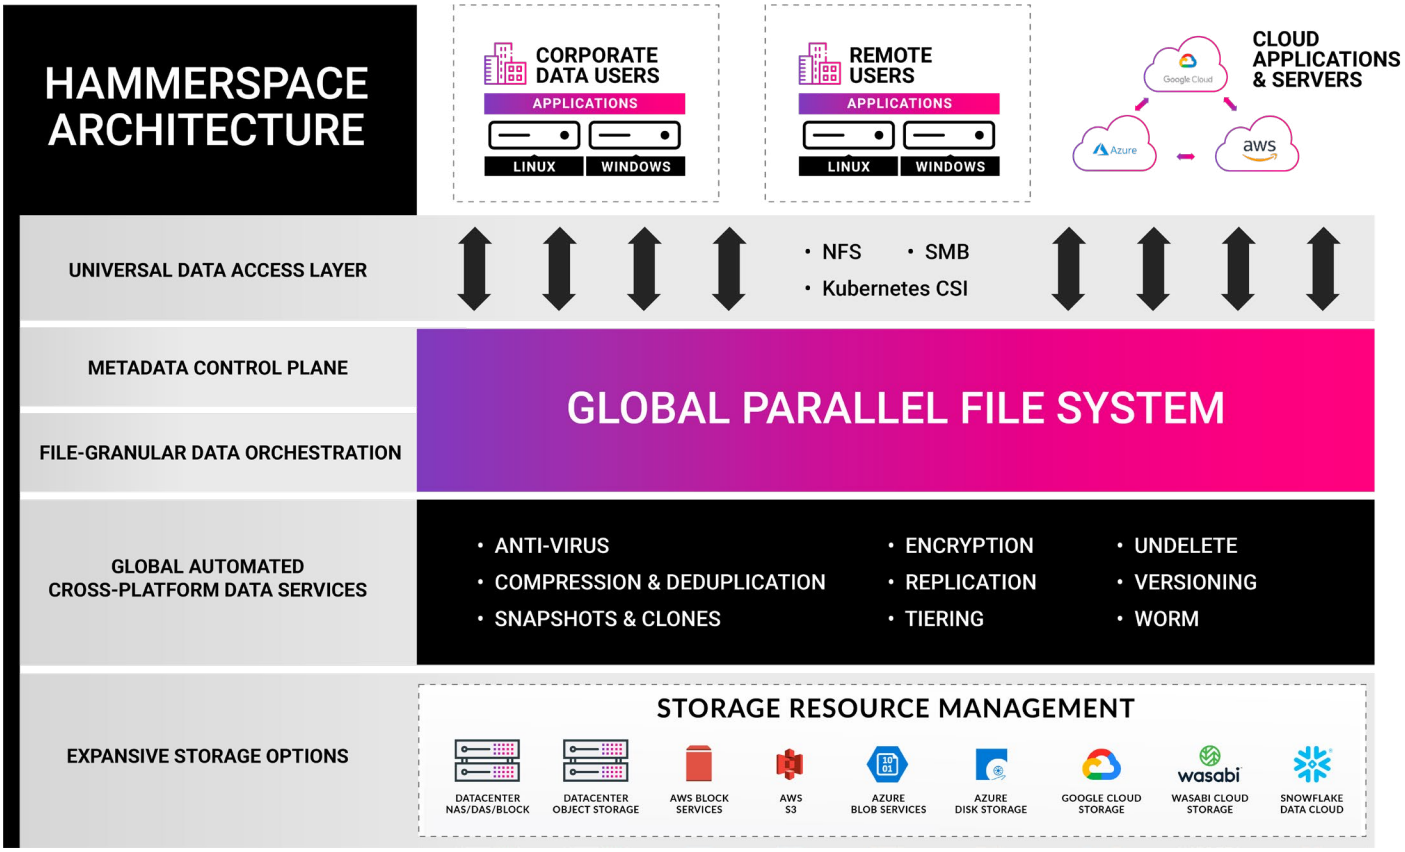

Figure 2: The Hammerspace Software-Defined Global Data Environment

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

- *Universal Data Access Layer*
- *Advanced Metadata*
- *Flexible Data Orchestration*
- *Automated Data Services*
- *Expansive Storage Options*

Universal Data Access Layer

Hammerspace supports multi-protocol frontside access via NFS, pNFS, SMB, and Container/Kubernetes Persistent Volumes for both block and file, individually and in mixed protocol mode. All data may be viewed by any protocol, regardless of the protocols supported by the underlying storage resources.

Additionally, Hammerspace provides a CSI (container storage interface) driver supporting both block and file-based persistent volumes in Kubernetes environments. The Hammerspace CSI driver supports shared-block, file-backed (local file), and shared file (NFS) persistent volumes provisioned from any underlying storage infrastructure. With Hammerspace, data management is efficient and automated. Local volumes do not need to be replicated to additional pods that need to consume data, which would add cost and complexity.

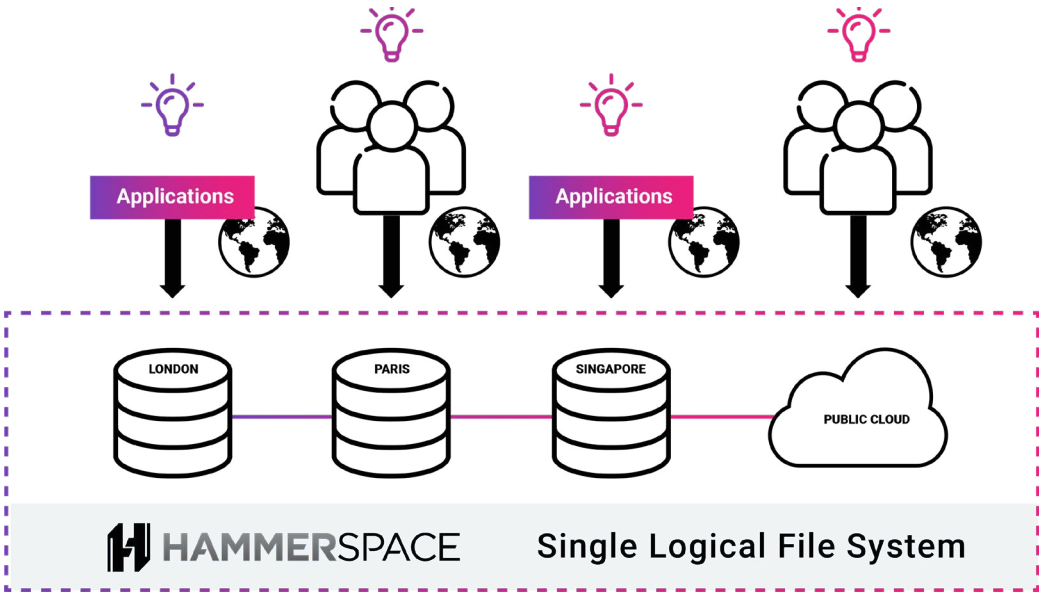

Fig 3. Users anywhere see the same global file system via multiple protocols as though they were local.

Hammerspace enables all to share the same data, and not have to wrangle copies.

Hammerspace data services presented to Kubernetes applications enable storage, backup and recovery, data management, and disaster recovery workflows.

With support for concurrent data access across protocols and full support for RFC 2307bis for Windows and UNIX id-mapping, Hammerspace can present data to any Windows, Linux, macOS, and UNIX operating environment globally.

Metadata Control Plane

Adding existing data into Hammerspace is made easy with our data-in-place metadata assimilation. Data does not need to move anywhere, since Hammerspace can uniquely replace existing file shares without performing a data migration to a new hardware platform. Hammerspace simply scans the share on your 3rd-party NAS platform and aggregates all metadata about the files and storage resources. And this happens almost instantaneously, even for extremely large file shares, since it all happens as a background metadata operation. This means Hammerspace can take over a network share of any size and begin serving read/write requests immediately. Global metadata control and

visibility eliminates the need to do time consuming and disruptive data migrations from one storage system to another, providing global data access to users and applications, with an experience the same as if they were accessing a local primary NAS share. This also means that what was previously multiple otherwise incompatible storage types with individual file shares, can now all be aggregated into the Hammerspace global file system. Users still only see the files they have permissions to see. But file system access is now global, bridging any local, remote or Cloud storage resource from a common file system view.

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.

This metadata tagging includes automatic metadata inheritance of custom tags, to ensure files are appropriately classified with critical information without relying on user action. 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, rendering jobs, simulations, data analytics, and more.

Automatically tagging datasets with department ID, job reference or other business data makes cost allocation for charge back or show-back a simple task, even though the files for those Departments may be stored across multiple storage platforms and locations.

In this way, Hammerspace enables network shares, regardless of size, to be visible and accessible around the world in multiple locations without needing to create and manage multiple copies of data everywhere. Distributed locations can include on-premises storage within one or more data centers and/or across multiple public clouds and cloud regions. This is accomplished by presenting the network shares via the globally distributed file system as a metadata control plane to all remote locations and leveraging just-in-time, file-granular replication and mobility to provide locality of data for better performance and lower latency based upon workflows and policy objectives.

Hammerspace manages data through metadata. Additionally, you can customize your metadata, giving you the power to better organize and utilize data across your enterprise.

Hammerspace network file shares can provide visibility to billions of files across a widely distributed infrastructure, while also managing local instantiations of the specific files needed for user and application workflows. For example, Location A may require 1,000 files, Location B just 200 files, and Location C only one file. Since all three sites are seeing the same file metadata in the global file system, back-end file movement for local access may be policy driven or on demand, and is completely transparent. Unnecessary copies and the resulting wasted local storage capacity are no longer a problem.

The combination of the GDE's ability to make the Hammerspace global file system metadata available to multiple locations with file-granular replication enables companies to work in ways that were previously impractical or even impossible due to price and performance challenges.

File-Granular Data Orchestration Layer

By separating the global metadata control plane from the data orchestration layer, Hammerspace can provide fiel-granular data orchestration to move data live and non-disruptively between silos in a data center, or across multiple data centers, or to public/private Cloud, between different regions within the Cloud, and across heterogeneous Clouds.Hammerspace's objective-based policy engine enables automated resource allocation to accommodate multiple workflows. Data is managed at file-level granularity, efficiently delivering scale across complex heterogeneous and distributed infrastructure without creating unnecessary copies of entire volumes of data.

The extensible metadata that is shared in the global metadata control plane is referenced by these user or Admin-created policy objectives, which are coupled with machine learning-driven automation to move and tier data across storage, sites, and clouds to achieve business requirements, better utilization of storage resources, for multi-instancing of files for data protection, etc. Hammerspace enables the establishment of intelligent policies to orchestrate and manage your data. You can create policies based on any metadata attribute including file name, file type, owner, path, create time, modify time, etc., as well as custom metadata tags at a file-granular level.

These policy objectives are implemented either through the GUI, Admin CLI, Hammerspace Toolkit, or REST API. The GUI provides a variety of checkboxes that instantly determine attributes such as availability, durability, placement, protection, and a host of other options. The same objectives can also be realized through the Hammerspace Toolkit, which is downloadable through GitHub.

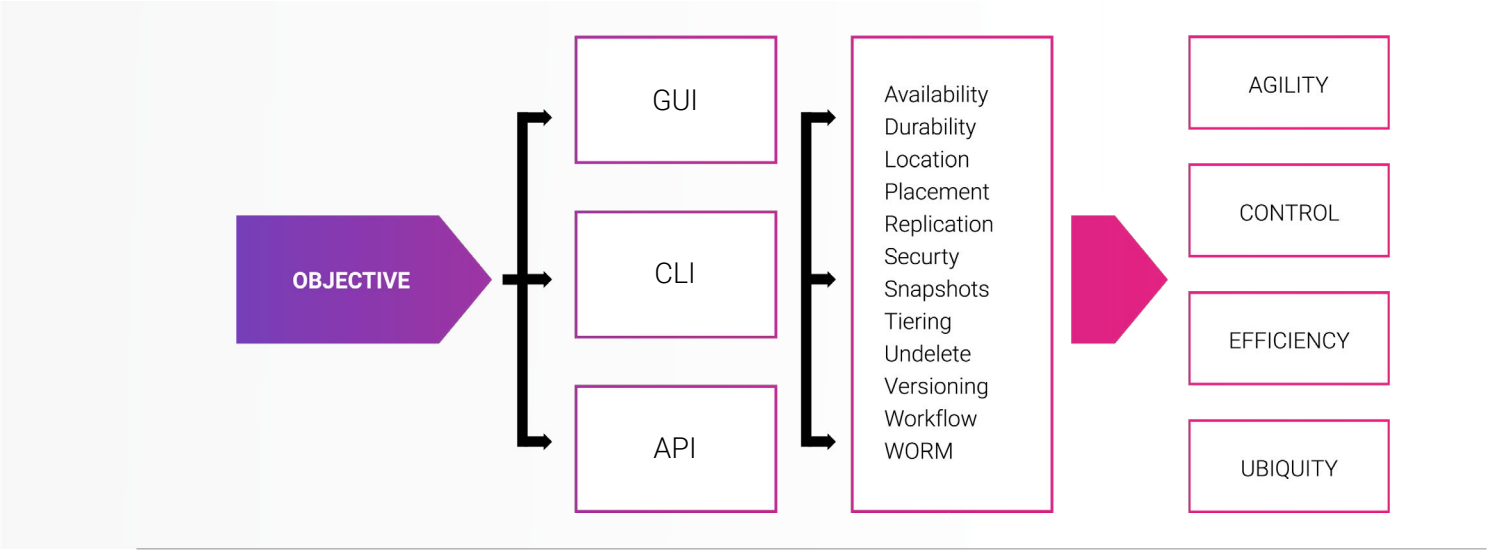

Figure 4: Objective-Based Policies

Global Automated Cross-Platform Data Services

Hammerspace provides file-granular global data services across all local and remote storage resources, leveraging its policy engine or on-demand capabilities. File-granular services give individual files, or sets of files, the ability to be managed by policies for any metadata attribute including 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 globally in ways that were previously impractical, or even impossible, due to price and performance challenges. Because these data services can be applied globally across all storage resources, the implementation of global control via the Hammerspace GDE 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:

- **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 usage. Take advantage of global data reduction to keep capacity minimal across all sites, on-prem, and the Cloud.

- **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 objective-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, for data redundancy, or to reduce latency when data is accessed from multiple sites. File-granular replication ensures that only the files you require at your remote locations are replicated, leveraging our policy-engine and/or on-demand capabilities. Once the initial replication is complete, only the diffs are replicated thereafter. And because the file metadata is consistent across all instances via the global file system, replication does not proliferate orphaned copies. All instances are consistent and managed globally.

- **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 daily snapshot per share across multiple storage types globally for up to seven years to meet long-term compliance needs without sacrificing granular recovery requirements. Hammerspace supports: 4,096 snapshots per share, file-level data recovery, snapshots stored on any storage including object storage and/or the cloud. 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 GDE. The consolidated view of all storage systems simplifies IT management and helps organizations more easily visualize their entire data environment to make pro-active 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 that only active files are utilizing expensive high-performance storage, for example. Policies may be set to move certain files to remote locations that have additional capacity, 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 attributes. Tier to any storage type, including data center block, file, object, and Cloud. Data is moved without disruption to user or application access, and in real-time to meet business SLAs and 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 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 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 corruption by physical or logical corruption, 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 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.

Expansive Storage Options

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 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 GDE. Hammerspace supports any storage type, from any vendor, whether block, file, object, including flash, disk, and cloud-based storage. Hammerspace is in the marketplace for AWS, Azure and GCP, and supports all the major public Cloud vendors.

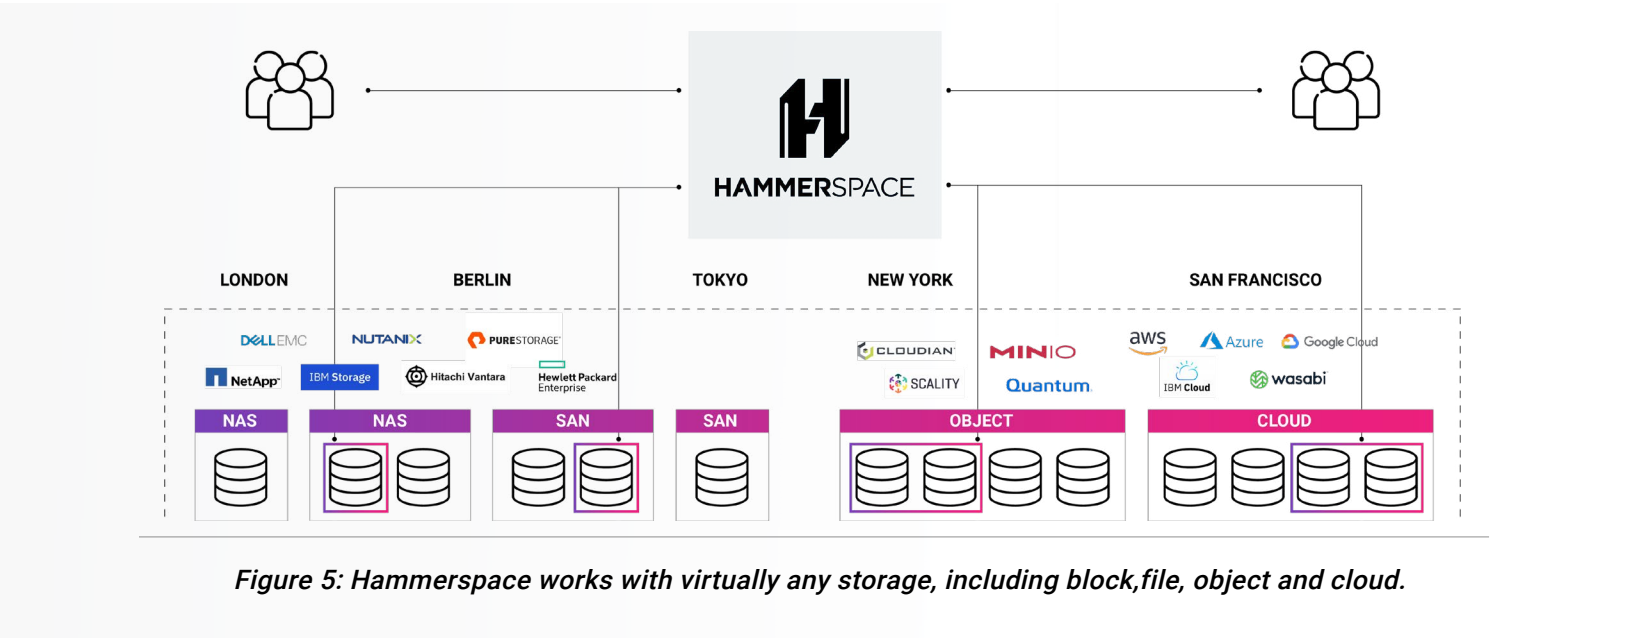

Figure 5: Hammerspace works with virtually any storage, including block,file, object and cloud.

Deploying Hammerspace

Hammerspace sold as a software-defined solutions, and is designed to be simple to install, and to accommodate virtually any IT environment. It installs in minutes as a complete ISO with all software components, and will runs on bare metal servers or within virtual machines, including VMware ESX, Microsoft Hyper V, KVM, and Nutanix AHV.

The software is deployed with two components: the Anvil metadata nodes, and the DSX data nodes.

- Anvil metadata nodes are configured in highly available (HA) pairs and are fundamentally the control plane of the system.

- DSX nodes are the data plane, and can scale out on-demand, supporting from one to 64 nodes in a single cluster, to scale out to accommodate even extreme IO requirements. The DSX nodes provide data storage, replication, and mobility services.

Hammerspace can store data on any storage type, including on-premises block storage, DAS, NAS, object storage, and in the public Cloud on block or object services. Hammerspace can use all of these storage types at the same time, intelligently, and automatically tier between them based on user-defined policies providing universal access via the global file system.

Benefits of the Hammerspace Global Data Environment

Use Cases

While Hammerspace capabilities greatly benefit siloed environments in a single data center, Hammerspace specializes in use cases where data may need to span two or more sites, including one or more Cloud providers and regions. It creates a Global Data Environment by providing unified data in a single global file system across sites.

Active Archive and Cloud Tiering

The vast majority of unstructured data is seldom accessed, and largely dormant. And yet in most environments, much of of this data is still stored on expensive Tier-1 storage systems. Hammerspace 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 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 be 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.

Hammerspace can tier data across multiple systems and locations, and will deduplicate and compress data in-flight and at rest. This enables customers to store offsite copies of data for redundancy, leverage the cloud to reduce on-premises infrastructure footprint, and utilize storage assets that might have free capacity at other locations for intelligent, proactive capacity management. This is done transparently, without interrupting direct user or application access to the data.

Cloud Burst for Compute

Video rendering, data analysis, simulation, and other workloads often require thousands of compute cores per project. As such, most companies do not own the required compute power, and need to leverage burst capacity in the public Cloud to take advantage of its ability to provide a massive number of cores.

Utilizing the public Cloud to process data can be extremely expensive. Organizations benefit from having complete control of which Cloud region they want to use for a given workload, as some Cloud regions are less expensive, or not as busy, as others.

Hammerspace can easily and transparently move data to the lowest cost regions in the Cloud to process organizations' data and greatly reduce expenditure. This can be orchestrated by Hammerspace, but most typically customers use their own scheduling applications to control this directly, enabling them to dynamically route jobs on demand to the lowest cost Cloud region as needed using tools they already are used to.

Data Management

With Hammerspace's rich metadata layer, all users and applications have read and write access to the same network shares and data regardless of their location, and regardless of where the actual files are located. This makes it very simple to collaborate on projects. When data needs to be actually accessed and worked on instead of just viewed, file-granular replication ensures that only the data that is needed in each location is replicated to the

appropriate sites. Automated data orchestration manages the live and transparent mobility of data across all storage resources in the global namespace. This is not copying data to another site, like traditional point solutions. This is exposing direct access to the same data across multiple sites via the global file system.

Data Protection

Modern data protection requires mitigating multiple threats to data. Hammerspace provides multiple mechanisms to protect data from disasters and attacks through immutable snapshots, undelete, WORM (Write Once Read Many), and file versioning, which provide comprehensive, layered protection to ensure data availability. In addition, file copy management may be automated, to ensure global DR policies are applied across files living on any storage type and location.

In this way, Hammerspace provides a rich palate of data protection capabilities that may be applied across all data on any storage type globally by IT administrators. No longer do IT staff need to wrangle multiple point solutions that are specific to a particular storage type.

Hammerspace can also complement existing backup solutions, enabling more back-end storage choices for backup images, snaps and other needs that might otherwise consume more expensive local resources. And when combined with policy-based copy management and tiering, it means Hammerspace provides IT administrators unprecedentd control globally, across otherwise incompatible storage resources. In this way, data that may be distributed across multiple silos can be managed with unique policy-based objectives mapped to the unique use case for those files, to tailor file-level protection and compliance policies to business needs.

Disaster Recovery

Hammerspace reinvents the data management component of planning for disaster recovery. If a site or system goes down, no failover or failback procedure is needed as the same global file system metadata exists in multiple locations providing continuous online access. Users and applications can be automatically and transparently redirected to the alternative location and quickly continue operations without needing to re-point applications, or mount different shares. In the event of an outage at one site, the global file system will still be available to users and applications to view, and data is still accessible as long as there is an instantiation of the data in an alternate location or DR site. To ensure the desired redundant instances of data are created, automated policies can be configured to replicate active data to the desired alternative location(s). Again, this is not like solutions that require IT Admins to re-point applications or users to different IP addresses or shares where copies might be stored. This is the same global file system view to the same data, with the same mount points via the global Hammerspace Metadata Control Plane. The physical location of the data could be anywhere, or in multiple instantiations for redundancy.

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 of all an organization's network shares is extremely challenging, as data is typically stored in multiple data silos in legacy data center 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 its metadata-based global file system and invokes 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, or REST API.

Multi-Cloud

New data services for compute resources, application processing, machine learning, and data storage can be created on a continual basis as needed. Hammerspace provides live and transparent data mobility across Cloud providers, and Cloud regions, to make it easy to leverage the Cloud service or application best suited for the job.

The Hammerspace multi-Cloud functionality also helps future proof your Cloud vendor decisions. If a change in organization policy or preferred Cloud partner changes in the future, migration of data is simple and transparent.

Ransomware Protection

Ensuring data is protected from ransomware is increasingly a priority for companies. Most organizations are seeking more sophisticated ways to combat or recover from attacks.

A popular ransomware technique involves creating a new and encrypted temporary file and then deleting the original data. Hammerspace's undelete feature automatically saves a copy of the deleted files, 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 do not delete data, Hammerspace's file versioning stores the latest versions of files 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 provides a combination of snapshots, undelete functions, WORM, and file versioning, to 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.

Storage Standardization

Hammerspace is software-defined and can be deployed on a wide range of infrastructure. With Hammerspace, you pick the deployment method that best suits your needs. Installation and configuration are simple and fast on bare metal, hypervisors, or the cloud in any combination required.

- Bare metal installation is supported on enterprise hardware platforms.
- Supports hypervisor installations on VMware ESX, Microsoft Hyper-V, KVM, and Nutanix AHV.
- Available on AWS, Azure, and GCP marketplaces.

With Hammerspace you can standardize global management of otherwise incompatible storage resources. This is important because it has a direct impact of simplifying infrastructure management tasks for IT, and in the process reducing OPEX. This may include tasks such as the end-of-life process of aging hardware, accelerating and automating cloud migrations, unifying the data environment across block, file and object storage systems, and much more. When the time comes to retire a storage system, for example, an Objective can be created to transparently migrate data by policy to any existing or new storage resource without interrupting user or application access. With data migration as a policy-based background operation, this puts an end to the headaches, outages and pressure on IT staff to minimize interruptions to the business. Migrations no longer need to be reactive and distruptive, and be proactive and transparent.

Hammerspace can work with a wide range of storage infrastructure implementations, including block storage, NAS, object stores, and Cloud from different vendors. The outcome is a unified data access layer to users, applications, and compute environments for all your data.

Summary

Hammerspace revolutionizes data management in a global data world. It enables organizations to use their existing storage resources to create an automated and scalable Global Data Environment to provide users with local access to data that may live on any storage tier in multiple locations. 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 platforms. The various storage types therefore are now resources that are managed dynamically in the background as needed according to business priorities for accessibility, performance, protection, and so on.

No longer do business 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 Metadata Control Plane, with file granular global data services applied via data orchestration layer, Hammerspace automates data access and orchestration, and provides global support for any storage infrastructure. In this way, a Hammerspace Global Data Environment provides local access, and global control and data services for all digital assets, across all storage types and locations, while reducing the management complexity for IT staff.