Native Persistent Data Orchestration

SOLUTIONS BRIEF

Kubernetes Data Management

THE CHALLENGE

Enterprises everywhere are going through a digital transformation, adopting new IT technologies like Kubernetes to improve their ability to rapidly react to business demands. Kubernetes is the most popular container orchestration platform, simplifying the development, deployment and scaling of containerized workloads. As more stateful workloads are managed through Kubernetes, the associated persistent data must be properly managed (performance, protection, compliance) and with native Kubernetes integration for developers who need to orchestrate their data but know nothing about enterprise storage. The concerns multiply when you add multi-cluster and hybrid multi-cloud to the equation. Kubernetes was built to be infrastructure agnostic, but as soon as data becomes persistent the challenges are often mischaracterized as storage infrastructure problems. Specialized storage solutions for Kubernetes are not the answer; they just create new data silos and compromise data agility by forcing developers to navigate storage infrastructure.

THE SOLUTION

Hammerspace takes a data-centric approach to file data in the cloud, serving and managing it independently from the infrastructure. Built for the hybrid multi-cloud, Hammerspace serves data at high-performance to any cluster across the hybrid multi-cloud. Using standard protocols, Hammerspace abstracts data at file-level granularity from the infrastructure so that any storage system or service can natively serve data to Kubernetes persistent volumes.

To span data management across the hybrid multi-cloud, Hammerspace separates the control plane (metadata) from the data plane (data) reading, writing, and moving data across sites through a Universal Global Namespace, at file level granularity. Hammerspace metadata servers (Anvil) are present at each site, replicating metadata so that every site has a complete view of all data, with the assistance of machine learning-driven automation to direct resource optimization. When non-local data is accessed, Hammerspace data services (DSX) moves data live to where it needs to be, even while actively being read/written. DSX data services are architected to scale-out on-demand so that performance

AGILITY, CONTROL & EFFICIENCY

Accelerate DevTest to Production

- Access data from any cluster through single data service
- Copy-free multi-cluster data workflows
- Non-disruptive app data mobility between clouds

Data Protection & Disaster Recovery

- Use snapshots for periodic checkpoints
- Pod granular protection policies
- Deliver high availability across clouds

Serve Data From Any Storage

- Leverage block or file for high-performance
- Optimize resources with auto data tiering
- Change storage classes on-the-fly

HAMMERSPACE

HAMMERSPACE.COM/CONTACT

is parallelized to meet application SLAs. Hammerspace key management server (KMS) integration encrypts all data stored and moved across the cloud; and data is protected by services like snapshots, undelete, and replication defends against the loss of infrastructure.

AGILITY: ACCELERATE DEVTEST TO PRODUCTION WORKFLOWS

Through the Hammerspace Universal Global Namespace Kubernetes clusters can share and access data from anywhere across the hybrid cloud. Data automatically tiers and moves as necessary to the most appropriate storage resource, completely transparently to any PV mounted in a pod. This makes it fast and easy for DevTest workflows to gain access to data without having to copy it first.

CONTROL: AUTOMATED DATA PROTECTION AND ACTIVE-DISASTER RECOVERY

When Kubernetes persistent volumes access data through the Universal Global Namespace, metadata enables security and data protection services to keep data safe anywhere across the infrastructure. DevOps teams can orchestrate data protection policies through services such as undelete, snapshots, and disaster recovery are all available globally across the infrastructure. Hammerspace enterprise data services are infrastructure agnostic and file, pod, or site granular.

EFFICIENCY: SERVE DATA FROM ANY STORAGE INFRASTRUCTURE

Hammerspace supports mixed storage infrastructure on-premises, making any storage container-native. When the infrastructure is abstracted away with the Universal Global Namespace, it becomes easy to consolidate storage resources, non-disruptively tier data to cloud storage, and change storage classes on-the-fly without stopping/starting a running pod. With unified block or file service into a PV, Hammerspace serves data from anywhere in the cloud using local storage for high-performance.

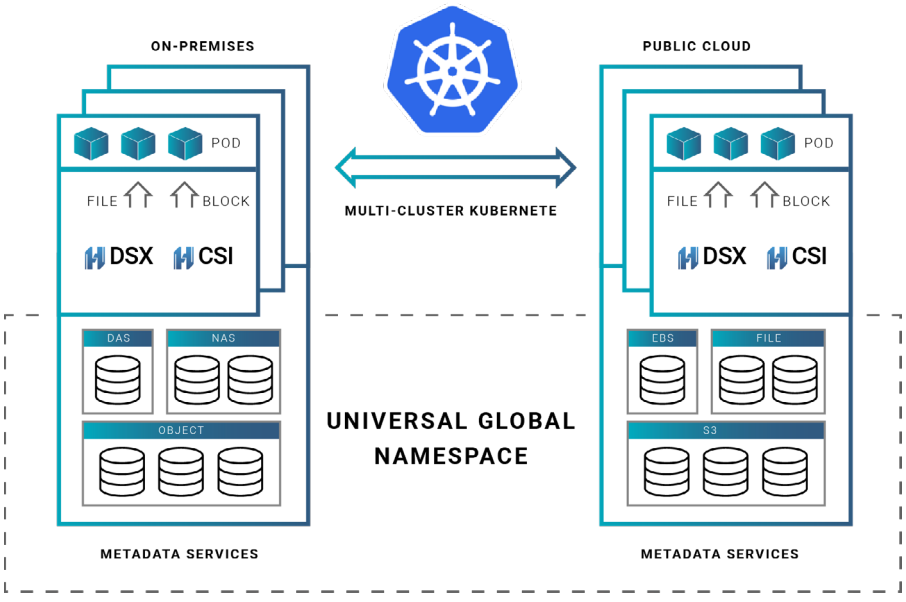

ABOUT HAMMERSPACE

Hammerspace is a hybrid multi-cloud file service that smashes the complexity of managing and protecting data on the hybrid-multi-cloud, eliminating the challenges of making unstructured data cloud-native and independent of the infrastructure. With non-disruptive, ML-driven data management, Hammerspace reduces the complexity of adopting hybrid, multi-site or Kubernetes workflows.