HAMMERSPACE

Solution Brief

Simplify Database Deployment Across Kubernetes Clusters

Highlights

- Data mobility between clusters
- Global storage classes
- Dynamic persistent volumes
- Cloud-to-Cloud active-active DR
- Millions of IOPS
- Integrated performance telemetry

Enterprises everywhere are going through a digital transformation, leveraging IT technologies across all aspects of their business to improve their competitive advantage. Kubernetes, the container orchestration platform, is one of these technologies used to simplify the deployment and scaling of containerized workloads.

The global accessibility of persistent data across multi-cluster Kubernetes environments is the biggest challenge for scaling popular stateful workloads like MySQL and MongoDB. As soon as data becomes persistent, it must be protected and governed by enterprise data services without sacrificing performance.

Scale Databases with Multi-cluster, Multi-Cloud Persistent Volumes

With Hammerspace, scaling and deploying databases across multiple clusters is as easy as with a single cluster. Data is abstracted from the infrastructure and virtualized, making it instantly available to any cloud or cluster.

Persistent data is available to Kubernetes through dynamic persistent volumes backed by global storage classes which automatically adjust to the available localized resources, keeping the experience consistent. Hammerspace intelligently delivers data on-demand to persistent volumes through standard open protocols, simplifying consumption with native Kubernetes integration.

With Hammerspace, all data is visible and available in a geo-spanning global namespace for intuitive orchestration by metadata management. Data mobility is performed live and is automated by machine learning to optimize for cost, performance, and protection.

Simplify App Data Mobility

App teams can accelerate database deployments by creating templates to define data sets to be rapidly pushed across different locations with pre-initialized clones, saving time and reducing error.

Wherever you deploy and start your containerized app, the data will just be there. Hammerspace makes all data globally accessible through its geo-spanning global namespace. It can move data live to meet workload demands and archive data when it's no longer needed while reducing Kube-sprawl – the unnecessary copying of data.

Hammerspace works on any storage or cloud service that speaks NFS, S3, or block; supporting multi-vendor environments with support for enterprise data services such as snapshots and replication. With unified support for both File and Block interfaces in Kubernetes, Hammerspace lets you mount files as block to support databases while moving database backups to NFS or object so that you don't consume local resource capacity.

High-Availability and Disaster Recovery

Database high-availability and disaster recovery across clusters have never been easier. A simple click configures and enables data protection with the Hammerspace geo-spanning global namespace. Data is automatically orchestrated, active-active, across storage and clouds.

Contact us at sales@hammerspace.com

HAMMERSPACE

Solution Brief

Keep Track of Data Across Clusters with Metadata Management

Using metadata, a user can describe the content of files mounted as block, making it easy to track where data originated from and manage it. Advanced reports can be instantly generated to understand data usage and to track growth over time, specific to the workloads of interest.

Leverage Local Storage for High Performance, While Spanning Clouds

With a Container Storage Interface (CSI) implementation that can deliver both File and Block interfaces from a global namespace, Hammerspace can service workloads that demand block with a file-based solution, without the penalty of NFS networking. Data can promote a block-file into storage local to a Kubernetes worker node to deliver high-performance.

Files served as block can be dynamically and non-disruptively tiered between block, NFS, and object infrastructure to meet performance SLAs while meeting cost requirements and reducing capacity consumption, across clusters and clouds. Every cloud environment will deliver different performance profiles, so Hammerspace will automatically adjust the placement of data behind the scenes to maintain declared intent and SLAs, requiring no intervention from the user.

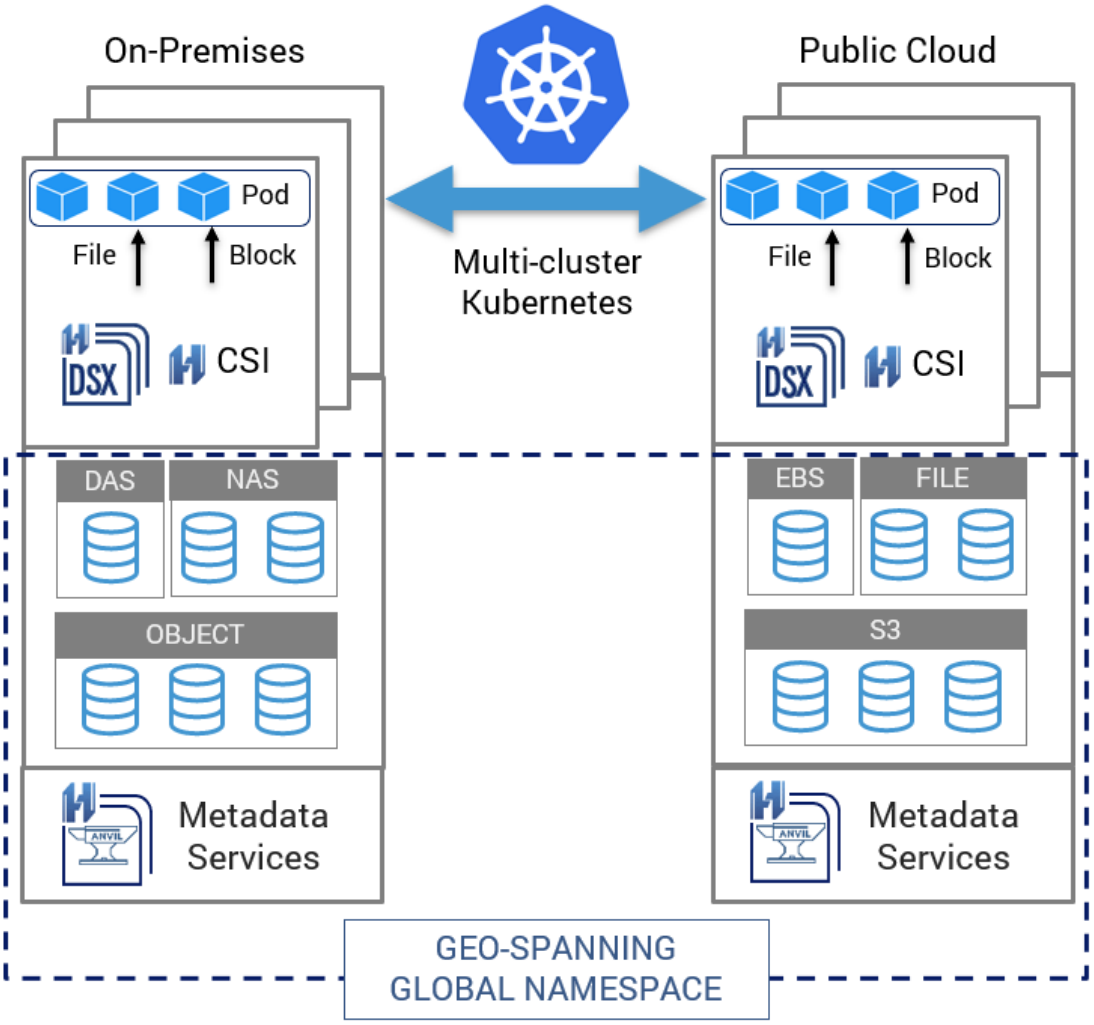

About Hammerspace

Hammerspace is a software company dedicated to enabling fast and easy access to data across the hybrid cloud. Hammerspace is a hybrid cloud data control plane where data exists abstracted from storage and is available to any service, in any cloud or data center. By automating the management of data with metadata-driven machine learning, Hammerspace makes it easy to run more jobs faster and not get stuck in a silo again. To learn more, visit us at www.hammerspace.com

Contact us at sales@hammerspace.com