

Simplify Hybrid Cloud Data Management for Google Kubernetes Engine (GKE)

Highlights

- Instant data access between clusters
- Auto data tiering with global storage classes
- Change storage classes on-the-fly with dynamic persistent volumes
- Cloud-to-Cloud active-active DR
- Millions of IOPS
- Integrated performance telemetry

Google Kubernetes Engine (GKE) is a critical component of Google Cloud Platform; now with Hammerspace's availability on the GCP Marketplace, there is an easy solution for Hammerspace customers to span data management between these services.

Hammerspace intelligently delivers data on-demand or proactively 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 through metadata-as-a-service. Data mobility is performed live and is automated by machine learning to optimize for cost, performance, and protection.

Availability of Hammerspace Data-as-a-Service on the GCP Marketplace allows customers to easily manage and protect their data across hybrid cloud with global enterprise file data services. With support for any on-premises storage and GCP storage classes, Hammerspace makes it easy for DevOps, application owners, and data operations to consume their data-as-a-service on-premises, cloud, and container environments.

Scale Workloads with Multi-cluster, Hybrid Cloud Persistent Volumes

With Hammerspace, scaling and deploying workloads 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 GKE 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 without application disruption, enabling workloads to scale and grow without storage limits.

Simplify App Data Mobility with Copy-Free Workflows

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. DSX data services scale-out on demand, moving 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, Native Object, 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 enables the full spectrum of cloud-native applications to benefit from its broad capabilities of data management services.

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.


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. The application specific metadata can also be utilized further to provide application -specific protection

Leverage Block 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 multi-cloud.

Hammerspace is a hybrid cloud data control plane where data exists abstracted from storage and is available to any service, container, or cloud. By automating the management of data with metadata-driven machine learning, Hammerspace makes it easy to create more value from data and never get stuck in a silo again. To learn more, visit us at

www.hammerspace.com