Storageless Data Orchestration for Kubernetes

DATA SHEET

Cut The Tethers Of Data Gravity & Stop Worrying About Storage

Hammerspace is storageless data. Achieve the promise of Kubernetes to enable workload portability across any environment by using technology that makes complicated enterprise storage obsolete. Hammerspace enables persistent data orchestration across the hybrid cloud by overcoming the data gravity generated by siloed storage infrastructure.

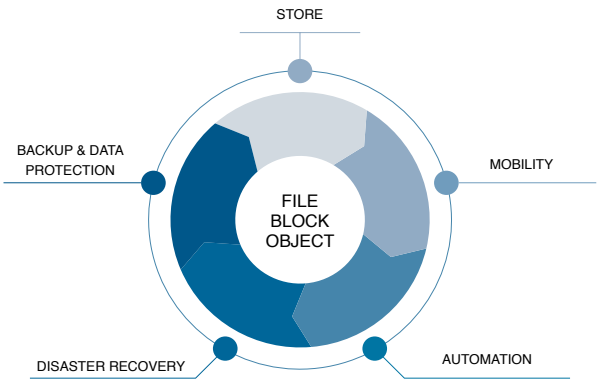

KUBERNETES SOLUTIONS

Disaster Recovery	Backup and Recovery	Test/Dev Workflows for DevOps	Burst to Cloud
Simplify DR to the cloud by making it fully automated for application deployments. With support for AWS, Azure and Google along with data centers, data is made resilient across fault domains.	Replication, snapshots, clones, and global undelete protect app data from accidental data loss, eliminating the need for a full backup recovery process.	Self-service access controls, data management and protection, performance scaling, and portability simplify and accelerate the full CI/CD workflow.	Developers, remote workers, and other data consumers can easily collaborate globally and burst their compute resource needs with simple declarative statements and instant data portability.

Hammerspace makes data storageless by freeing it from the tethers of being defined by the storage infrastructure that your data happens to be on. By using a global file system built on a metadata engine, Hammerspace serves data from block, file, or object storage infrastructure to workloads located anywhere – on-demand. This simplifies data orchestration at scale by making app data portable, protected, and high performance – what Kubernetes does for containers, Hammerspace does for data.

The Advantage of Storageless Data Orchestration

Pay As You Go

An elastic consumption model based on capacity means that you only pay for what you use as you scale up or down.

Ubiquitous Data Services

Enterprise data services are defined by the data, not the storage, allowing users to self-service their data management anywhere, from the cloud to the edge, without refactoring their apps.

Use Any Storage

Data resides on storage systems but is not bounded by the physical infrastructure. So, persistent data management in Kubernetes is greatly simplified by managing through metadata to abstract data from the infrastructure.

Global Data Access

Instant access to data from the cloud to the edge using live data mobility automated by ML-driven data placement optimization. Operating at file and container-level granularity, this allows users self-service their data management while avoiding disruptive data migrations and the expense of unecessary copies. you scale up or down.

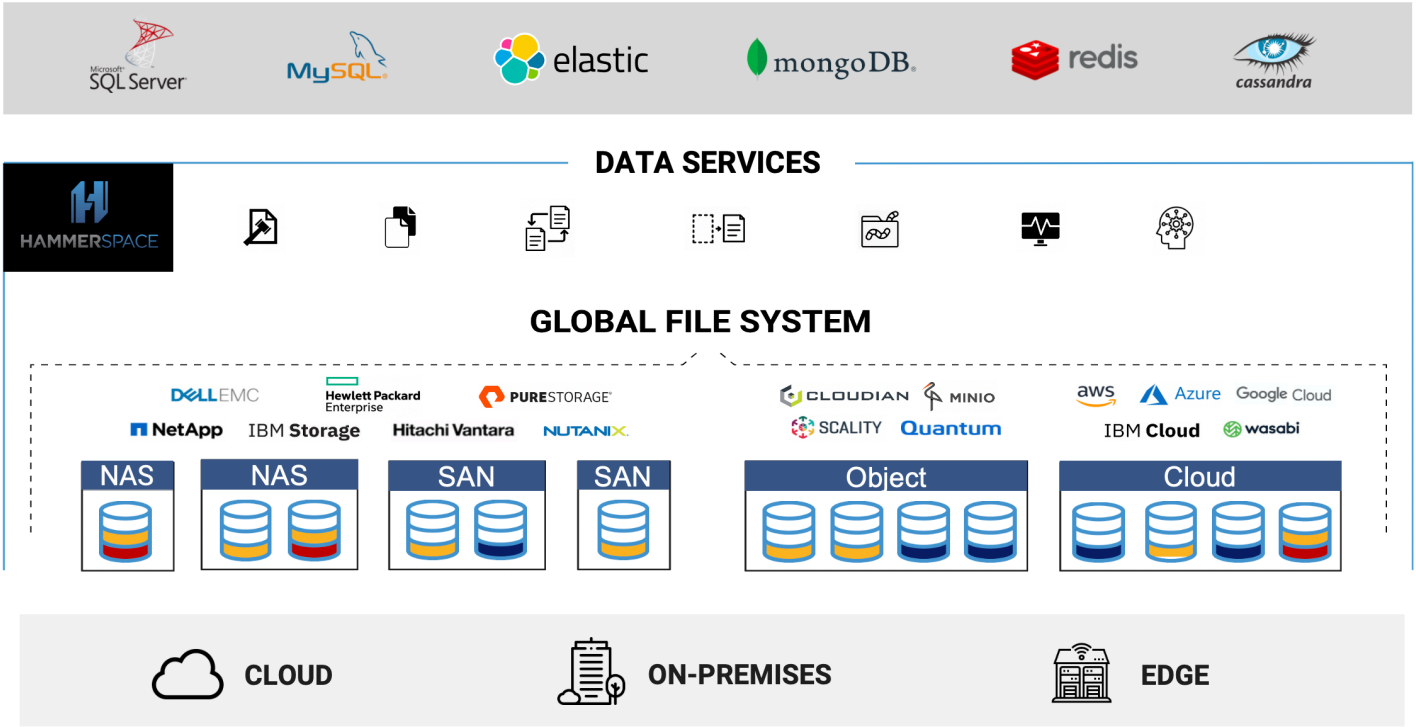

About Hammerspace

Hammerspace is storageless data for hybrid cloud Kubernetes environments. By untethering data from the infrastructure, Hammerspace overcomes data gravity to provide dynamic and efficient hybrid cloud storage as a fully automated, consumption-based resource. Users self-service their persistent data orchestration to enable workload portability from the cloud to the edge.