Backup & Recovery for Kubernetes
(Files, Containers, or Clusters)

SOLUTIONS BRIEF

Reduce The Need For Formal Backup With Hybrid Cloud Data Services

THE CHALLENGE

Backup and recovery are essential IT activities to protect against data loss, but the technology has not adapted itself well to the needs of DevOps running Kubernetes environments. Backup often has rigid workflows and tends to capture a lot of redundant and unnecessary data, consuming capacity, and licensing costs. Restore is slow and usually requires the intervention of IT. Both issues are not very compatible with the ethos of Kubernetes, which is to make workloads cloud-native for portability, scalability, and efficiency.

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 site across the hybrid cloud. Hammerspace abstracts data from the infrastructure, making it easy for Kubernetes users to self-service backup and recovery data services that are consistent across the hybrid cloud, regardless of what storage the data happens to live on.

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 Global File System, at file level granularity. Hammerspace metadata servers 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 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 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 to defend against the loss of infrastructure.

AGILITY: PUT BACKUP AND RECOVERY SERVICES IN THE HANDS OF USERS

Allowing Kubernetes users to self-service data services that follow the data and

AGILITY, CONTROL & EFFICIENCY

Put backup and recovery services in the hands of users

- Share level snapshots for rapid and frequent backup
- Global undelete for quick restore at the user-level across hybrid cloud
- Data vault is a safe place to protect data, enabled by WORM data-lock

Manage the data that gets backed up using platform intelligence

- Backup as an integrated data management policy
- Build awareness of data types found in workflows

Reduce the overall cost of formal backup

- Reduce unnecessary copies
- Optimize data placement for global dedupe and compression
- Centralize copies of backup data in the cloud

are not tied to the infrastructure gives them the agility they need to keep up the pace of their work while also reducing the support they need from IT. Services include features such as snapshots, undelete, replication, and data-vault with WORM. Frequent snapshots can store years of data demanding enormous capacity requirements, while undelete is a fast and easy tool to undo accidental data loss. The more services that DevOps can use themselves without impacting the underlying infrastructure, the easier it becomes for IT to support their activities without needed to make huge changes to how they already work.

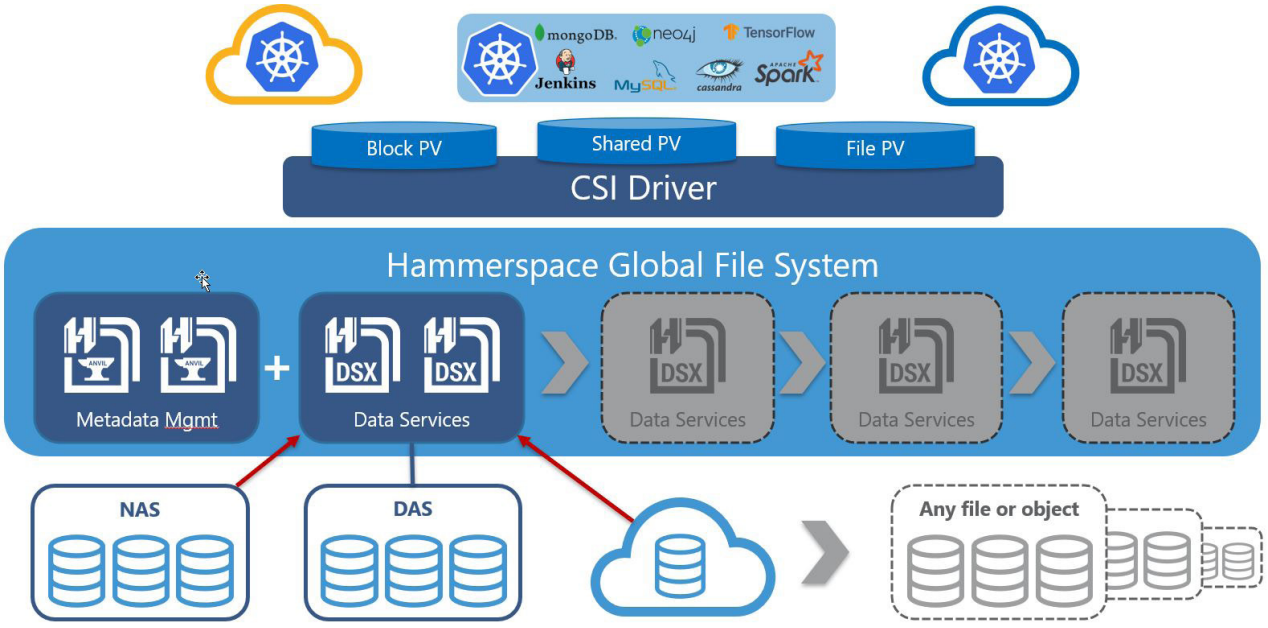

CONTROL: MANAGE THE DATA THAT GETS BACKED UP USING PLATFORM INTELLIGENCE

When backup and recovery are made available as data services, then integrated data management policies can bring control to what data finds its way into formal backup workflows. There is even more control over what data gets backed up when the system builds awareness over the types of data that the system backs up. This is all managed through the metadata as service-level objectives, giving global control and overcoming the siloed nature of storage infrastructure.

EFFICIENCY: REDUCE THE OVERALL COST OF FORMAL BACKUP

It can help keep the data copy sprawl that can build up from TestDev and CI/CD workflows by using a data orchestration platform that makes data available to container workloads and clusters without copying full volumes first. Data is managed at file and container level granularity. So, there is much more control over how many copies of data find their way into backup workflows, reducing the capacity costs of backup. The platform optimizes data placement across the infrastructure for cost, performance, and protection but this also has the effect of optimizing the cost of backup by moving things like snapshots to cheaper object storage. Finally, user data and backup data can be centralized in the cloud to reduce the capacity and licensing costs of doing backup at the edge.

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. To learn more, visit us at www.hammerspace.com or on Twitter @Hammerspace_Inc