Burst To Cloud & Global Collaboration For Kubernetes

SOLUTIONS BRIEF

Make Data Instantly Available Anywhere For Extra Resources Or Sharing

THE CHALLENGE

Burst-to-cloud and collaboration workflows to make data accessible across hybrid cloud environments usually requires full copies of volumes to be made between sites and knowledge of the particular infrastructure available. This requires Kubernetes users to coordinate their needs with IT who then perform data migration or synching tasks to copy data to the remote site (Cloud, edge, or user). This entire process can be disruptive, slow, labor intensive, prone to error and failure, and costly.

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 their needs for extra compute resources in the cloud or sharing data to remote users.

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. 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: SIMPLIFY HYBRID CLOUD WORKFLOWS WITH INSTANT DATA ACCESS ACROSS SITES, CLUSTERS, AND CLOUDS

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

AGILITY, CONTROL & EFFICIENCY

Simplify hybrid cloud workflows with instant data access

- A global file system makes data available across sites, clusters, and clouds
- Infrastructure is abstracted, eliminating the worry about reconfiguration
- Data portability enables remote users and elastic compute resources

IT stays in control of the infrastructure

- DevOps teams have the tools they need for data orchestration
- IT manages the underlying infrastructure to service their customers
- IT sets guard rails for how data is protected and accessed while supporting agile dev teams

Reduce the overall cost of data

- Live data mobility eliminates disruptive data migrations
- Reduce capacity costs with fewer full copies of data
- Data placement is automated and continuously optimized

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. The Hammerspace global file system makes it fast and easy to make data available in the cloud or between sites regardless of the underlying storage infrastructure. A global file system has the added benefit of true cloud-native portability, meaning that containerized workloads will require no refactoring or reconfiguration as they are moved to new environments. This essentially makes app data portable so that remote users or elastic cloud resources can quickly have access to the data they need.

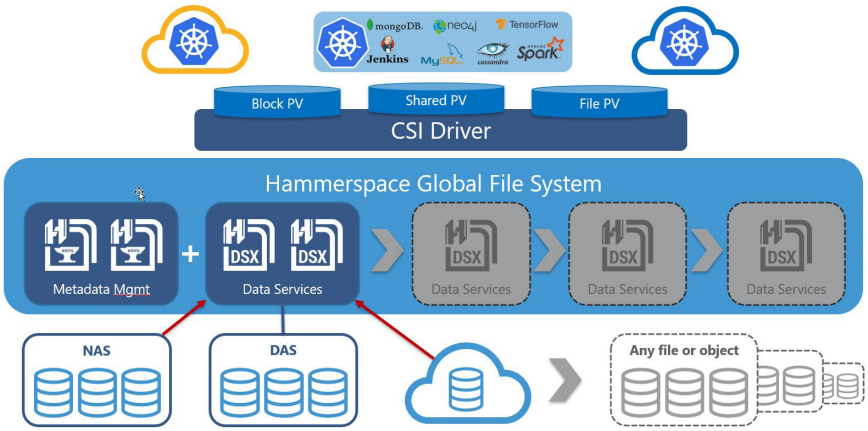

CONTROL: IT STAYS IN CONTROL OF THE INFRASTRUCTURE

DevOps teams and remote users are looking for the fastest way to get access to the data they need either to process it in the cloud or somewhere at the edge. With Hammerspace, a global file system and fully automated data management gives IT control across the infrastructure to meet SLA requirements while enabling true data orchestration for data users. Data is protected, performance is maintained, and cost is managed by data policies set by IT while DevOps teams have the tools they need to non-disruptivly move and access data anywhere. This reduces the load on IT to support these teams without giving up control over the things they are most measured on.

The cost of data is greatly reduced

- Live data mobility eliminates disruptive data migrations
- Reduce capacity costs with fewer full copies of data
- Data placement is automated and continuously optimized

EFFICIENCY: REDUCE THE OVERALL COST OF DATA

The cost of data can be managed and predicted when you have dynamic control over the infrastructure and the data that lives on it. With Hammerspace, data is managed at file and container-level granularity and data is moved live and non-disruptively effectively eliminating data migrations. This means that burst-to-cloud workloads can generate fewer copies of data overall as only the data needed is copied between sites, on-demand. Finally, data placement is automated and continuously optimized to tier data across the infrastructure from high-performance flash, to object, to cloud storage services. This intelligent optimization of data keeps costs low, while unlocking the data agility required for a cloud-native hybrid IT environment.

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