# DevOps Workflow for Kubernetes

SOLUTIONS BRIEF

## Save Your Sanity, Enable Persistent Data Orchestration For TestDev Workflows

### THE CHALLENGE

One of the top challenges to supporting a DevOps team running TestDev workflows on Kubernetes is that you must support persistent data performance and accessibility requirements while protecting and preserving it. These activities are made more difficult as multiple clusters span hybrid cloud environments which brings complexity for IT while limiting the portability and scalability promised to DevOps by the cloud native ecosystem.

### THE SOLUTION

So, how do we help accelerate TestDev workflows for DevOps without giving up control of the infrastructure? By enabling data orchestration so that DevOps can programmatically self-service their data management needs while maintaining the guard rails set in place for the enterprise data services found on storage infrastructure.

With Hammerspace, this can be done declaratively using service level objectives to describe the desired state, while its global file system non-disruptively moves data live across the infrastructure to meet those objectives. Objectives can describe performance, cost, protection, durability, location, and other aspects; all of which can be simplified even further by IT by creating storage classes in Kubernetes for DevOps to deploy as persistent volumes. Using the global file system approach to enabling data orchestration means that data can instantly be made available between clusters and across sites without first replicating volumes and copying full sets of data.

### AGILITY: DEVOPS CAN SELF-SERVICE DATA MANAGEMENT

Hammerspace abstracts the physical infrastructure, which serves users with instant access to data across clusters, sites, and clouds. This greatly simplifies the orchestration of data, using service-level objectives to make declarative statements about what is needed from data from performance to protection. By removing all the complexity of storage from Kubernetes environments, DevOps can focus on development workflows with the confidence that the management of data and storage will be consistent across whatever environment they use.

## AGILITY, CONTROL & EFFICIENCY

**DevOps can self-service data management**

- Instant access to data across clusters and physical locations
- Orchestrate data using ser vice-level objectives
- The complexity of storage infrastructure is abstracted from Kubernetes

**IT maintains the infrastructure**

- Use any storage
- Implement guard rails to maintain performance and protection
- Reduce IT interventions to support DevOps

**The cost of data is greatly reduced**

- Live data mobility eliminates disruptions
- TestDev generates much fewer copies of data
- Data placement is automated and continuously optimized

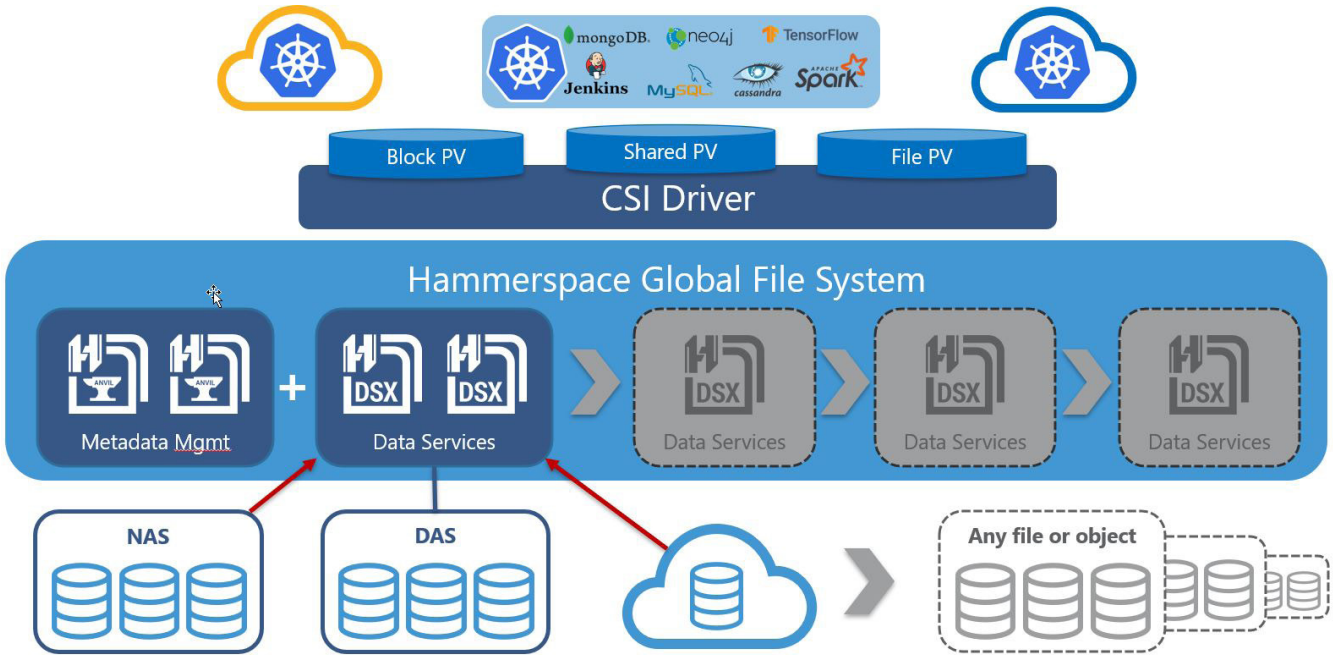


## CONTROL: IT MAINTAINS THE INFRASTRUCTURE

With Hammerspace, IT maintains control and maintenance of the storage infrastructure while giving their DevOps customers the tools they need to remain agile. By putting in place guard rails to maintain performance and protection, IT can use storage classes in Kubernetes as one of these tools. Since data can be moved and managed live, it becomes non-disruptive to make changes to the parameters of these storage classes or changes to the supporting infrastructure, completely transparent to users. This reduces the number of IT interventions required to support DevOps teams as frequent tedious tasks are taken care of by automation.

## EFFICIENCY: THE COST OF DATA IS GREATLY REDUCED

The cost of data can be managed and predicted when you have dynamic control over the infrastructure and the data that lives on it. Live data mobility eliminates disruptive data migrations and rsync jobs. Also, since data is managed at file and container-level granularity, TestDev workflows generate much fewer copies of data overall. 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 mobility 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