

Agile Server Infrastructure for Kubernetes

The DriveScale Composable Platform

Kubernetes container orchestration is transforming the way software is developed and deployed. Kubernetes provides the tools needed to build and deploy reliable, scalable, distributed applications. And with the addition of storage volumes, Kubernetes solves the need for storage availability by requesting and attaching volumes to containers making the deployment of stateful applications, such as databases, in containers a popular solution.



With Kubernetes, IT can deliver data-intensive applications such as Big Data, Machine Learning and NoSQL databases using lightweight, flexible containers with the ability to deploy, scale and manage them seamlessly. To take advantage of this newfound flexibility, legacy server infrastructure needs to evolve from the static, direct-attached storage servers used today to programmable and adaptable resources.

DriveScale Delivers Agility with Persistent Storage and Performance for Kubernetes

Traditional deployment models constrain and limit the flexibility of Kubernetes. When deploying Kubernetes using local storage in direct-attached storage (DAS) servers, applications and containers are captive inside the DAS server. In this configuration, local data is unavailable when a server fails so data must be replicated to multiple locations within the data center to ensure the container can be re-instantiated in a location where it can access the replicated copy of its data. In addition, when compute is tied to its local disks and data, compute resources can be underutilized or remain idle when a containerized workload completes.

To overcome the constraints of local storage, a distributed file system such as GlusterFS, Cinder or CEPH can be used in a container with Kubernetes which is often called container-native storage. Like HDFS (Hadoop file system), these file systems provide the data management and replication for the application. When high I/O performance is a critical requirement, raw volumes can also be provided to Kubernetes.

DriveScale offers a unique approach to server infrastructure for Kubernetes and containers. With DriveScale, users create, deploy and adapt servers on the fly to meet the demands and distributed nature of containerized applications and Kubernetes application pods. DriveScale's flexibility and adaptability is the perfect match for Kubernetes. And DriveScale automatically



Agile Server Infrastructure for Kubernetes Environments

provides persistent, raw storage volumes to Kubernetes or a distributed file system to provide both the performance and flexibility required. By providing persistent storage, DriveScale eliminates the confines of local storage enabling a container or pod to move anywhere in a server cluster.

DriveScale's integration with Kubernetes enables the DriveScale Composable Platform to automatically create, delete and attach raw storage volumes on demand in concert with a distributed file system or natively to a containerized application and ensure persistence. In this way, Kubernetes can spin up, spin down or move containers around server clusters and the associated storage remains connected.

When Kubernetes requests a volume, DriveScale creates the volume and attaches it to the compute node. When a compute node fails, the container can be re-instantiated where it is best placed in the server cluster, and its storage reconnects automatically which allows continuous operation and real-time allocation of compute resources to the jobs as needed. DriveScale also provides added intelligence and visibility into Kubernetes pods enabling a detailed view of which Kubernetes pods are running on each server and which storage volumes are connected to each pod.

DriveScale is the only server infrastructure that scales and adapts compute and storage resources to meet the needs of each application or workload. Heterogeneous, low-cost compute nodes and GPU nodes can be composed as part of the server infrastructure along with 100G connected NVMe flash drives or 10G connected hard disk drives in dense eBODs (Ethernet Box of Drives). With DriveScale, users can deploy server and storage infrastructure in minutes not months, maximize resource utilization and eliminate wasted spend with independent compute and storage upgrades.

Why DriveScale

IT teams use the DriveScale platform to quickly and easily compose heterogeneous compute, storage and networking resources into high-scale, adaptable servers for Kubernetes environments.

The DriveScale platform for Kubernetes provides:

- Automatically provisioned persistent storage
- Direct-attached like performance
- Optimized compute and storage utilization, eliminating overprovisioning
- Highly scalable, distributed server infrastructure

By optimizing resource utilization using DriveScale, companies can deploy server infrastructure at a lower cost than alternatives while ensuring they have the flexibility to quickly scale up or down compute and storage resources as needed.

Composable Infrastructure

Composable Infrastructure is next-generation server infrastructure that provides the ability to flexibly create, adapt, deploy and later redeploy servers using pools of disaggregated, heterogeneous compute, storage and network fabric. According to IDC, the composable infrastructure market is estimated to grow from \$752 million in 2018 to \$4.7 billion in 2023.