Mastering the Nuances of Container Storage

CONTAINER JOURNAL

Container Storage Soars

More than half of organizations (56%) are using containers to run stateful applications.

Source

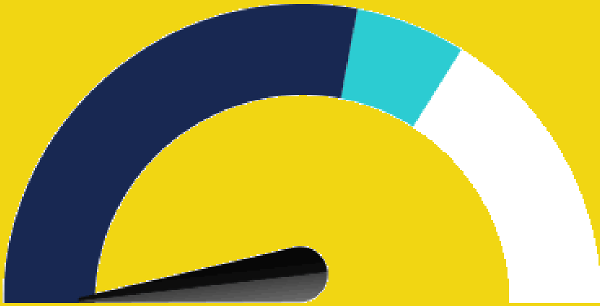

More than 55% of developers had deployed a stateful container application in production, with **another 11% planning to deploy them** in the next 12 months.

Source

80% of the attendees are challenged by the time and effort required to get stateful applications running on Kubernetes clusters.

Source

How Kubernetes Manages Container Storage

CONTAINER JOURNAL

Kubernetes creates permanent storage mechanisms for containers based on Kubernetes persistent volumes (PV). This makes it possible to access data far beyond the lifespan of any given pod. Kubernetes Volumes allows users to mount storage units to expand how much data they can share between nodes. Regular volumes will still be deleted if and when the pod hosting that particular volume is shut down. The permanent volume, however, is hosted on its own pod to ensure data remains accessible. Upon creation, the PV is bound to the pod that requested the PVC.

IT teams manage storage in Kubernetes via three cluster resources:

PersistentVolumeClaim (PVC) to request storage

PersistentVolume (PV) to manage storage lifecycle.

StorageClass defines different classes of storage service

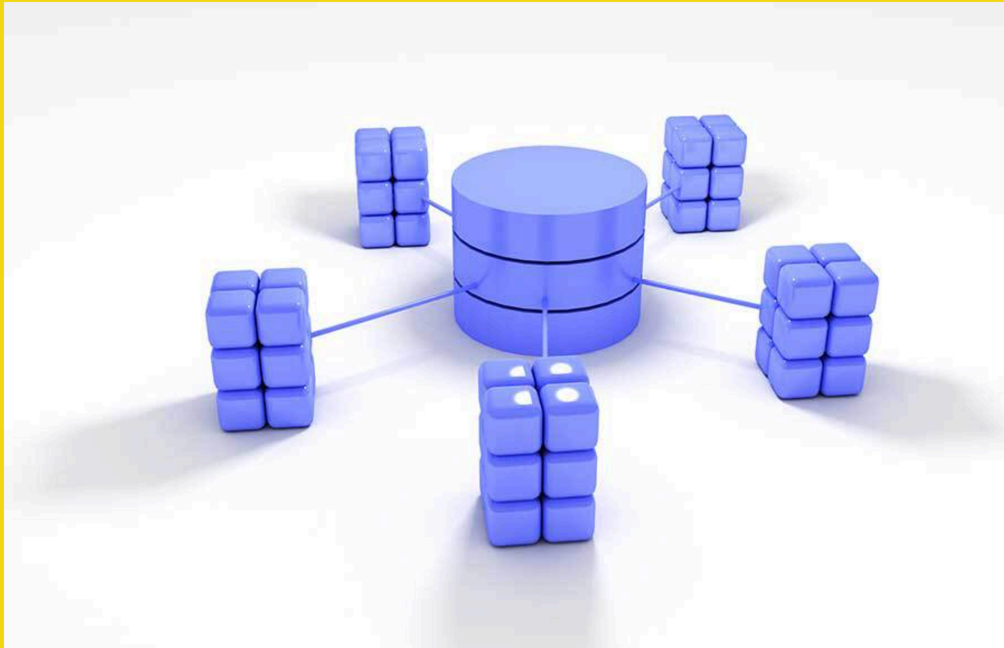

Container Storage Interface (CSI) Provides Access to External Storage

CSI was developed as a standard for exposing arbitrary block and file storage systems to containerized workloads for Kubernetes. It makes the Kubernetes volume layer extensible in a way that allows third-party storage providers to write and deploy plugins exposing new storage systems in Kubernetes.

Persistent Storage Comes in Many Forms

Stateful applications require access to persistent forms of storage that include flash memory, hard disk, tape or optical media. Persistent storage systems store data using file systems, block or object storage. Containers that are employed to build container applications tend to be less ephemeral than containers that are employed to build stateless applications. Of course, even a so-called stateless application needs to eventually store data somewhere.

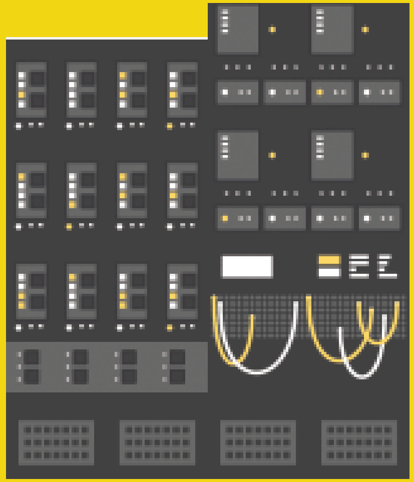

Backup and Recovery for Containers

Most containers are only live for a few minutes at a time, so the need for backup and recovery tools is often overlooked. However, entire clusters, individual nodes and associated persistent storage can suddenly go offline so there is a need to recover from disasters.

Backup and recovery tools are also used to move workloads from a test/dev environment to production, from production to staging before an upgrade and to migrate a Kubernetes cluster.

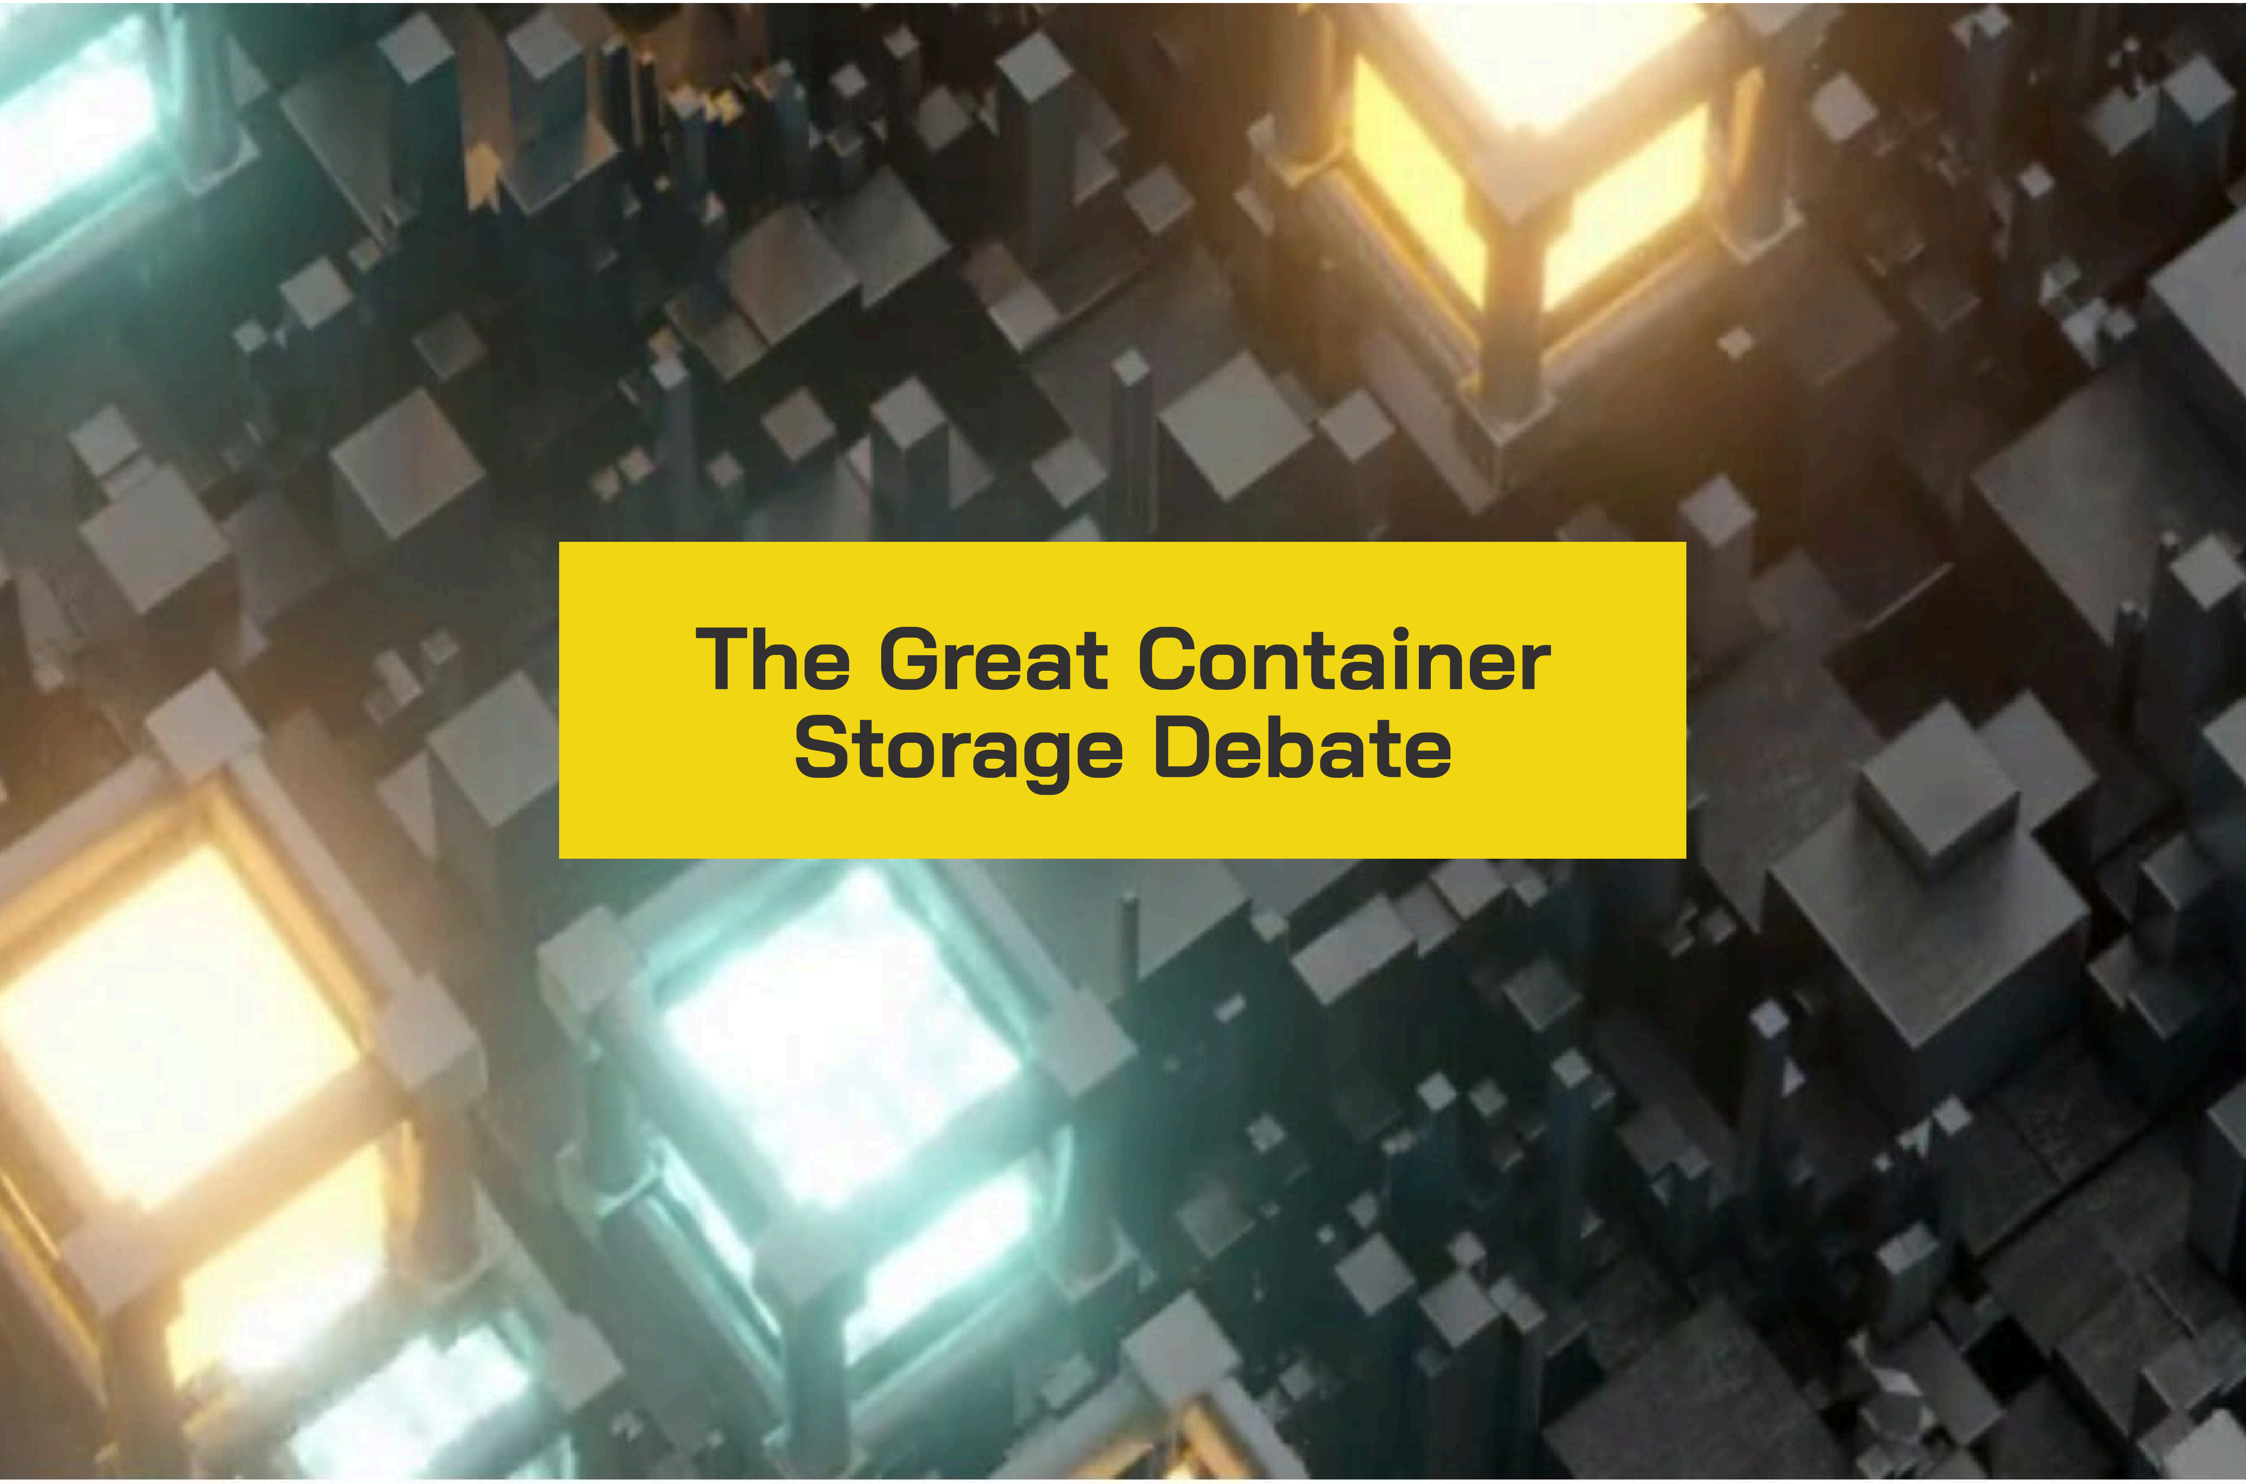

The Great Container Storage Debate

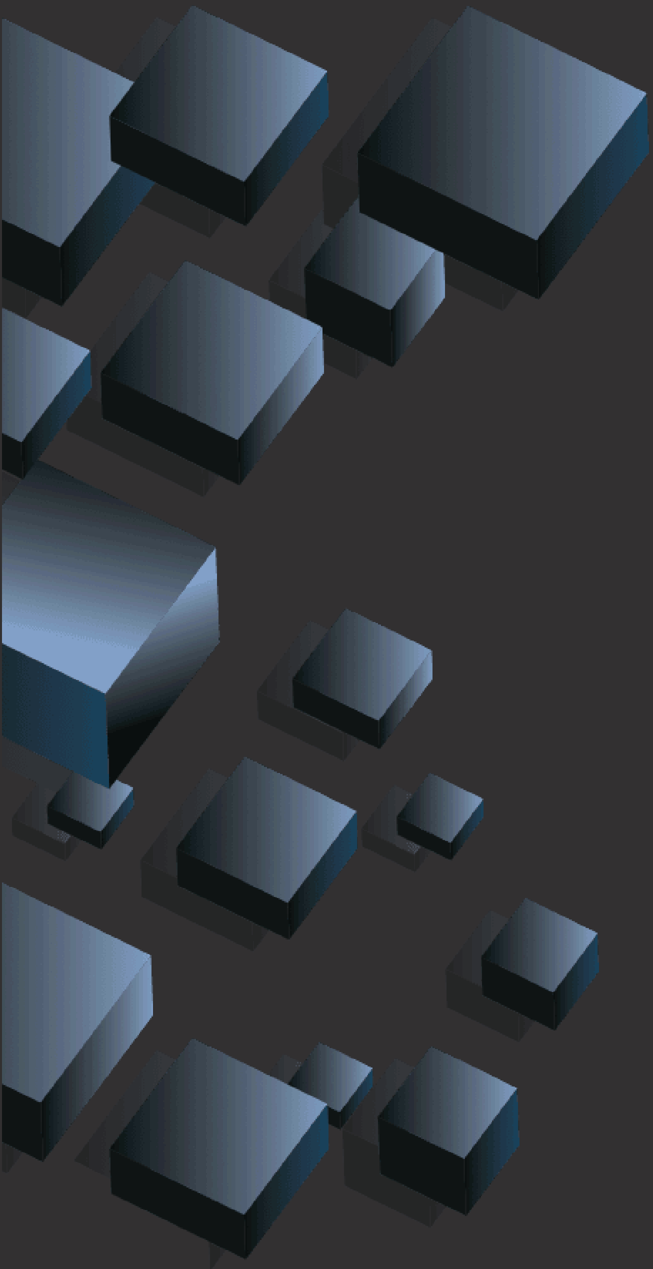

There are many who feel that managing storage systems to enable stateful applications to be deployed on Kubernetes is too complex an endeavor. They argue all container applications should be stateless in the sense that data is stored outside of the cluster. That approach makes it easier to rely on existing storage resources. Conversely, proponents of stateful applications on Kubernetes clusters believe this approach makes it easier to converge the management of computer and storage on the same platform. There is, of course, no one right answer that spans every potential container storage use case.

Thank you for reading

Nuances of Container Storage