



CLOUD STORAGE
TECHNOLOGIES

What's New in Container Storage

Live Webcast
February 26, 2019
10:00 am PT

Today's Presenters



Alex McDonald
NetApp



Keith Hudgins
Docker

- The material contained in this presentation is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
 - ◆ Any slide or slides used must be reproduced in their entirety without modification
 - ◆ The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

SNIA-At-A-Glance



CLOUD STORAGE
TECHNOLOGIES



185

industry leading
organizations



2,000

active contributing
members



50,000

IT end users & storage
pros worldwide

What We Do



Educate vendors and users on cloud storage, data services and orchestration



Support & promote business models and architectures: OpenStack, Software Defined Storage, Kubernetes, Object Storage



Understand Hyperscaler requirements
Incorporate them into standards and programs



Collaborate with other industry associations

What is a Container?

- Starts with a base OS system image
- Container init = 1 process only
- Process runs in its own namespace
- Shared access to system resources (hardware, network, etc)
- Containers Copy-on-Write image is NOT saved when process exits!

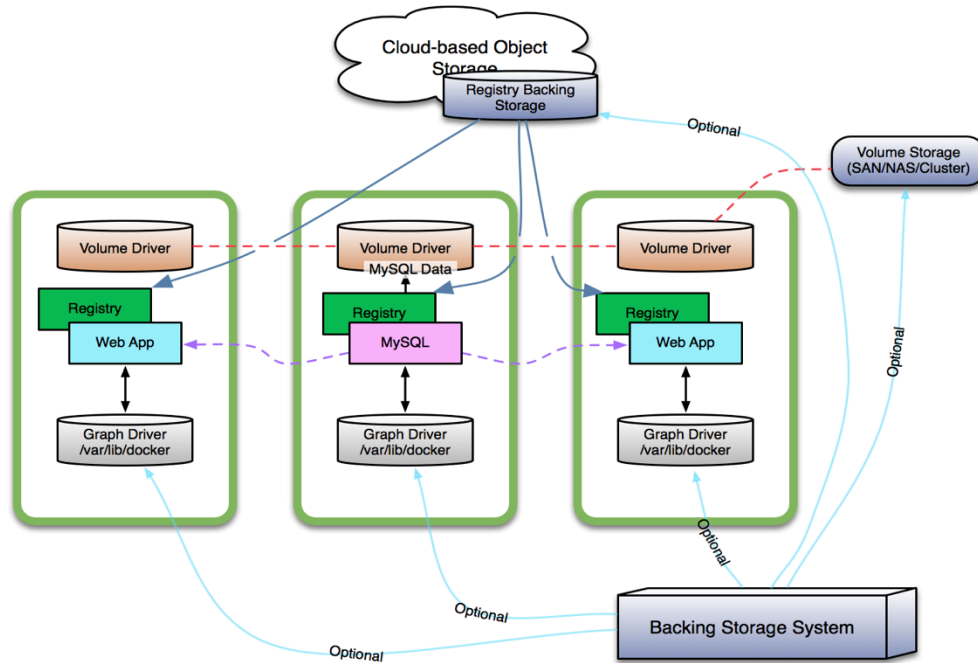
- Cloud Native Computing Foundation (CNCf)
 - ◆ <https://www.cncf.io/>
- Open Container Initiative (OCI)
 - ◆ <https://www.opencontainers.org/>
- Both are part of the Linux Foundation
 - ◆ <https://www.linuxfoundation.org/>

Container Storage Types



CLOUD STORAGE
TECHNOLOGIES

- Runtime Storage - Graph Driver
- Cold Storage for Images - Registry Storage
- Persistent Storage



➤ Integration point determined by container runtime

- > <https://blog.mobyproject.org/where-are-containerds-graph-drivers-145fc9b7255>
- > Basically, two driver styles:
 - Overlay
 - OverlayFS, AUFS
 - Snapshot
 - Btrfs, zfs, devicemapper

➤ Docker (containerd)

- ◆ Default for most container installs, widest user base

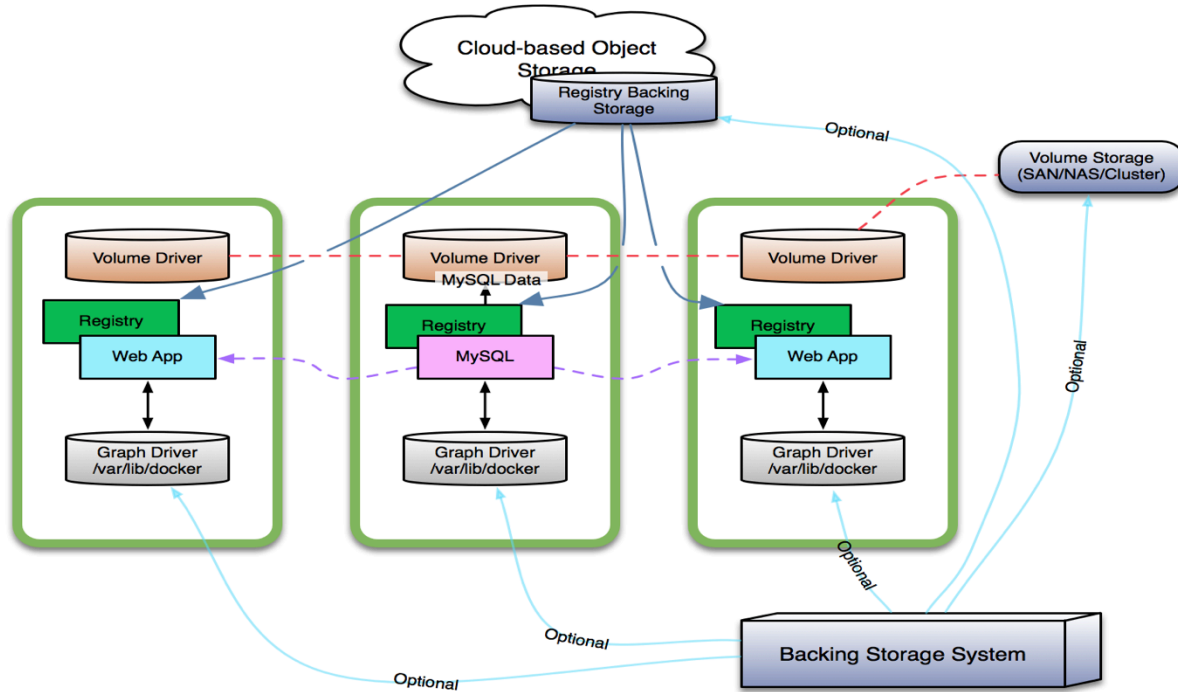
➤ CoreOS (rkt)

- ◆ Aims to be ultra-thin, ultra-simple
- ◆ CoreOS acquired by RedHat May 2018

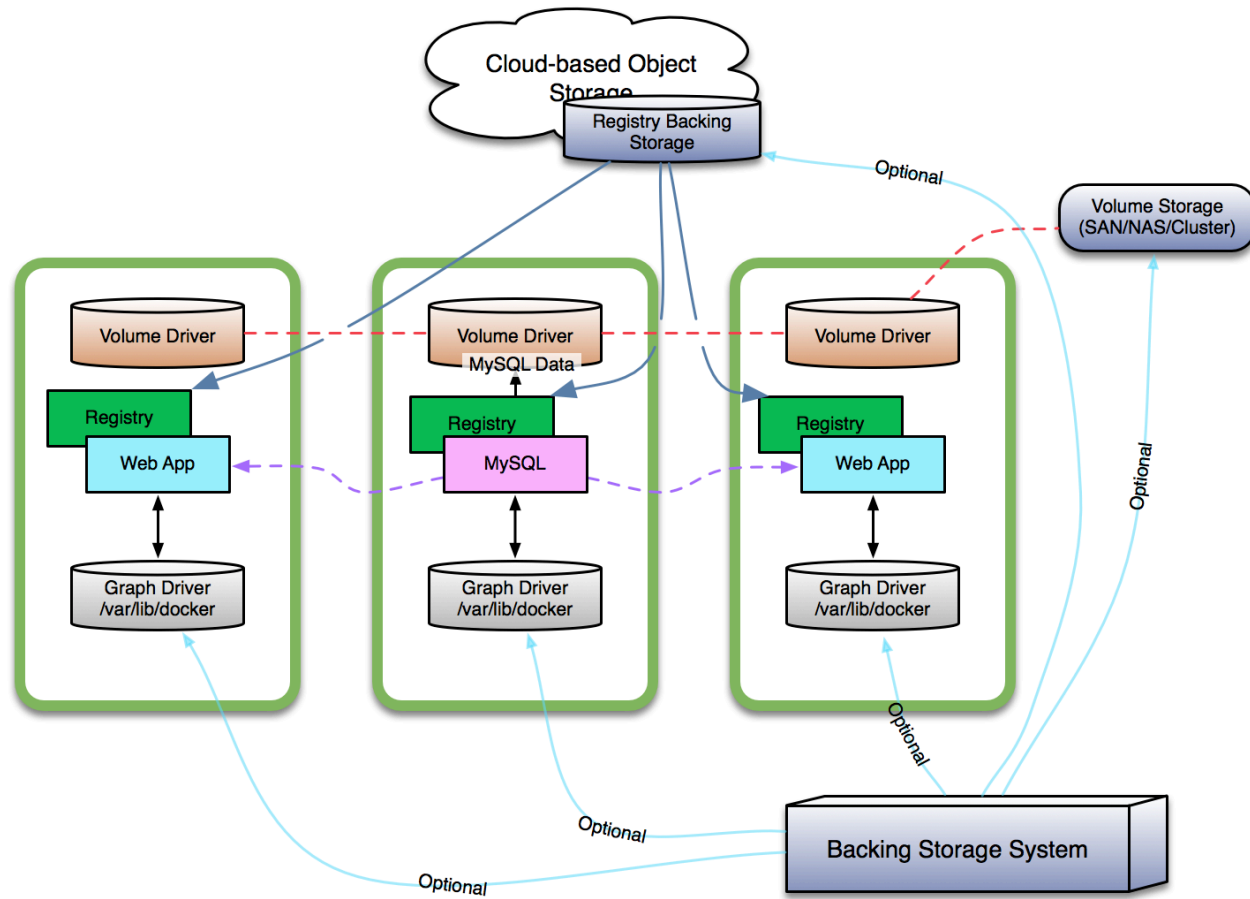
➤ CRI-O

- ◆ Intended to be Kubernetes-native runtime

➤ Upstream of these projects are all CNCF efforts.



- Many private registry options.
- Pretty much all use either
 - ◆ File-based storage
 - ◆ Object storage
- Docker Registry can use either
 - ◆ Supports any mounted filesystem for file-based
 - ◆ S3 or OpenStack Swift APIs for object storage



Persistent Storage

- This category is the single most important integration point for most storage vendors.
- There is little support within the container ecosystem for native persistence, and yet it is a major need for users.

- Every implementation works pretty much the same:
 - ◆ Plugin/integration code receives instructions from the framework to manipulate a file or block device
 - ◆ Plugin reaches out to backing store to mount the new device on the host system and prepare it for use (formatting, etc)
 - ◆ Plugin hands mount point back to framework
 - ◆ Framework bind-mounts it into a container prior to container launch
- APIs are a little different, and some other capabilities (snapshots, RBAC, etc) may or may not be baked into the APIs

Container Frameworks

- Docker Swarm
- Kubernetes
- Mesosphere

- Docker Volume Plugin supports two versions:
- V1: Supports Windows and Linux
 - ◆ Installs out-of-band of container solution
 - ◆ Less secure by default (can be manually tightened down)
- V2:
 - ◆ Install is containerized, so very easy for users
 - ◆ Uses Linux kernel capability controls for security management
 - ◆ Unfortunately, that makes it Linux-only
- https://docs.docker.com/engine/extend/plugins_volume/
- Requires Docker Engine, either Community or Enterprise

Kubernetes Integration - Important Points

Kubernetes is becoming the de facto standard for new container projects. (With growing pains involved)

You should be considering Kube Integration.

Kubernetes is a community project. Like all major community projects, there are many cooks. Which means there's several ways to integrate.

- Kube is operator-centric: admins must prepare volume offerings to users
- Applications are managed in pods
 - ◆ Users don't manage singleton containers in Kube - always app stacks via pods
- Admins prepare resources for users
- Storage Class - Admin prepares
 - ◆ Persistent Volume Claim - Declared in pod file
 - › Persistent Volume - Dynamically created at runtime for pod
- Community links to get started
 - ◆ <https://github.com/kubernetes/community/tree/master/sig-storage>
 - ◆ <https://github.com/kubernetes/community/blob/master/sig-storage/contributing.md>

Kubernetes In-tree Storage Drivers

- These are baked in to the Kubernetes source
- Most distros will ship them
- <https://kubernetes.io/docs/concepts/storage/>
- Example volume driver:
 - ◆ <https://github.com/kubernetes/kubernetes/tree/master/pkg/volume/rbd>
- You WILL need to join the Kubernetes community org and sign their Contributor License Agreement - this may not be a good fit for everyone!
- This is slowly being deprecated for CSI.

FlexVolume Plugins

- 1st released persistence API for Kubernetes
- Fully GA and supported as of Kubernetes 1.8
- Docs are hard to find - Kube docs team are moving things
- <https://github.com/kubernetes/community/blob/master/contributors/devel/flexvolume.md>
- Internal Kubernetes project

➤ CNCF Incubator Project

- ◆ This means more overall community support than FlexVolume
- ◆ However, this is also newer, and only recently GA.
- ◆ Cloud Foundry (not specifically containers) and Mesos have announced support for CSI as well
- ◆ Note that most distros have lag time for supporting Kube releases

➤ <https://github.com/container-storage-interface>

- ◆ Several repos there that are useful.
- ◆ Full specification is in the Spec repository
- ◆ <https://github.com/container-storage-interface/spec/blob/master/spec.md>

Kubernetes in the Cloud

May 2, 2019

11:00 am PT/ 2:00 pm ET

Register at:

<https://www.brighttalk.com/webcast/663/350613>

- Intro to Containers, Container Storage Challenges and Docker
 - ◆ <https://www.brighttalk.com/webcast/663/217971>
- Containers: Best Practices and Data Management Services
 - ◆ <https://www.brighttalk.com/webcast/663/227349>
- Containers and Persistent Memory
 - ◆ <https://www.brighttalk.com/webcast/663/266009>

After This Webcast

- Please rate this webcast and provide us with feedback
- This webcast and a PDF of the slides will be posted to the SNIA Cloud Storage Technologies Initiative website and available on-demand at <https://www.snia.org/forum/csti/knowledge/webcasts>
- A full Q&A from this webcast will be posted to the SNIA Cloud blog: www.sniacloud.com/
- Follow us on Twitter @SNIACloud

Thank You