
CERN

CSI Driver for CernVM-FS

CernVM-FS Workshop 2022 Amsterdam

Róbert Vašek, CERN

cern.ch

Agenda

Introduction
Mounting CVMFS in Kubernetes
Demo
Conclusion

We are here!

Introduction
Mounting CVMFS in Kubernetes
Demo
Conclusion

Last 60 days

Openstack services statistics

Users	Projects	Kubernetes clusters	Images	Volumes	Volumes size	File Shares	File Shares s...	Object Store ...	Object Store ...
3326	4496	356	3533	7349	3.78 PB	5304	890 TB	452	47.9 TB

Servers				Cores			RAM			Batch		
Physical	Physical in use	Hypervisors	Virtual	Physical	Hypervisors	Virtual	Physical	Hypervisors	Virtual	Servers	Cores	RAM
8658	8226	1996	13299	460 K	57.7 K	87.2 K	1.91 PB	375 TB	206 TB	4885	261651	989 TB

Time series

Physical servers

Last 4 months

Virtual Machines

Virtual Machines created/deleted

Last 7 days

Shared cells availability

Last 30 days

Kubernetes clusters

Volumes

File Shares

Object Store size

CSI Driver for CernVM-FS

Huh?

CVMFS CSI driver driver provides read-only mounting of CernVM-FS repositories in CSI-enabled container orchestrators, e.g. Kubernetes.

Brief history of External Storage in Kubernetes

- Started as persistent volumes pre-created by admin
- Required support inside Kubernetes code-base
- In-tree storage support is difficult to maintain: Flex drivers
- On-demand provisioning with StorageClasses and PersistentVolumeClaims
- Need of standardization: **Container Storage Interface** (CSI)

PersistentVolumes	Mar 2015
Flex plugins	Mar 2016
StorageClasses	Sept 2016
CSI (Alpha)	Dec 2017
CSI (Beta)	Mar 2018
CSI (GA)	Nov 2018

We are here!

Introduction

Mounting CVMFS in Kubernetes

Demo

Conclusion

CVMFS CSI

`github.com/cernops/cvmfs-csi`

- Can be used in anything that speaks CSI
- Exposes CVMFS repositories as Kubernetes-native PersistentVolume objects
- v1 created back in 2018, v2 coming out now!

+

CVMFS CSI v2 changelog

	v2	v1
Local cache volume	Any	Requires `ReadWriteOnce` volume or `emptyDir`
Client configuration	Any time	Deployment time only
Mounting	`autofs` and more	Private mounts with `mount -t cvmfs`

Getting autofs to work in containers

- Driver runs in host's PID namespace, to make `automount` daemon visible
- Mount propagation needs to be set-up properly

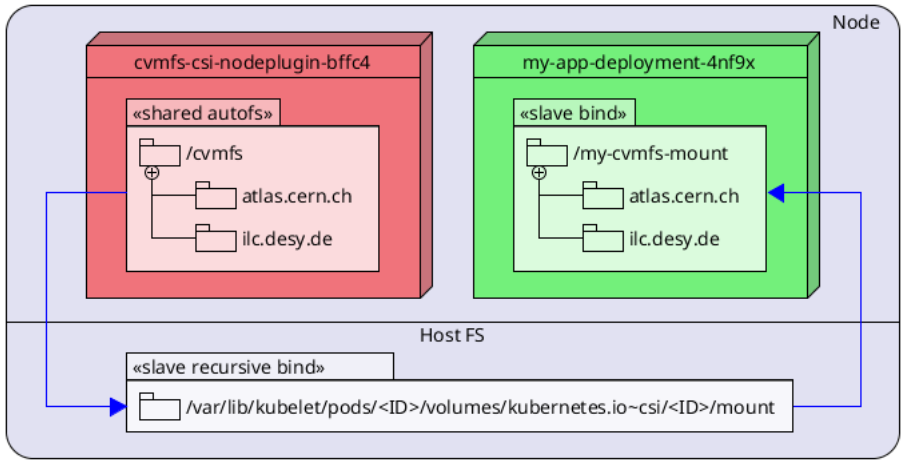

Upgrading from CVMFS CSI v1 to v2

- Almost no downtime™!
- Involves restarting (delete and create) Pods that use CVMFS volumes

1. Set cvmfs-csi Node plugin's update strategy to `OnDelete`
2. Upgrade the driver's Helm chart deployment
3. For every node where Node plugin Pod is running:
   - **3.1** Delete csi-cvmfsplugin v1 Pod and wait until the v2 Pod starts. Note this causes I/O failures in Pods using CVMFS volumes on this node.
   - **3.2** Delete all Pods on this node that are using CVMFS volumes. This refreshes the mounts, Pods should be now up and running again.

We are here!

Introduction
Mounting CVMFS in Kubernetes
Demo
Conclusion

Demo

https://asciinema.org/a/LuDog1Y7VshFJhtwcnDVKHbhT

We are here!

Introduction
Mounting CVMFS in Kubernetes
Demo
Conclusion

Conclusion

- CVMFS CSI v2 patches sent, to be released soon
- Where to go from here:
  - We've dropped a couple of features for now: per-PVC config, preloading.
  - These will be added in the following interations.
  - Need a feature? Open an issue on github.

Questions?

- CernVM-FS: cernvm.cern.ch/fs/
- Kubernetes: kubernetes.io
- Container Storage Interface: github.com/container-storage-interface/spec
- CVMFS CSI: github.com/cernops/cvmfs-csi

- Robert Vasek <rvasek01@gmail.com>