



# Container Storage Improved: What's new in CVMFS and EOS Integrations

CERN IT Container Service Webinar

Robert Vasek, CERN



[cern.ch](https://cern.ch)

# Agenda

**Introduction (to CSI)**  
**CSI driver for CVMFS**  
**CSI driver for EOSxd**

Bonus: eosxd-csi + oauth2 = <3  
**CSI driver for NFS**  
**Conclusion**

# We are here!

## Introduction (to CSI)

CSI driver for CVMFS

CSI driver for EOSxd

Bonus: eosxd-csi + oauth2 = <3

CSI driver for NFS

Conclusion


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



# 80% of the slides are related to CSI... What is it?

Introduction (to CSI)  
CSI driver for CVMFS  
CSI driver for EOSxd

Bonus: eosxd-csi + oauth2 = <3  
CSI driver for NFS  
Conclusion



Not this CSI.

# 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)**

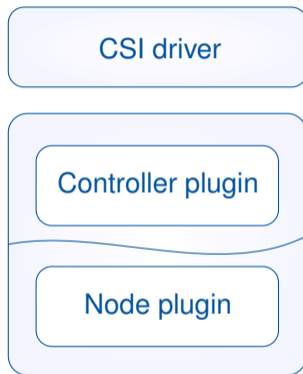
---

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

# Container Storage Interface

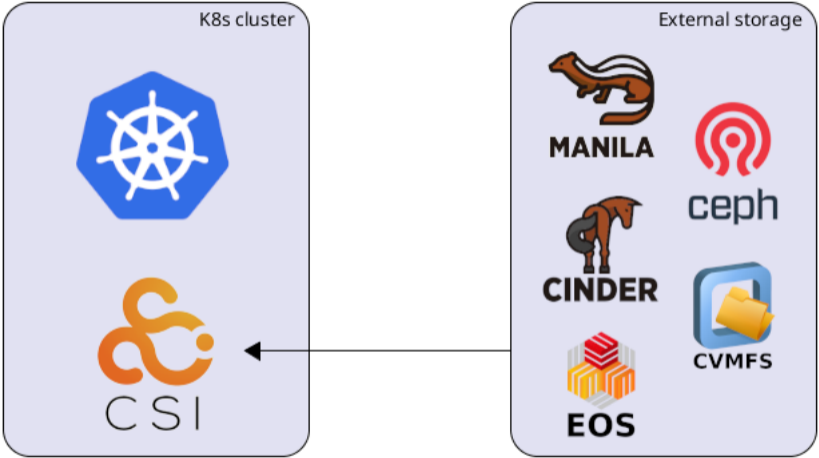
[github.com/container-storage-interface/spec](https://github.com/container-storage-interface/spec)

- ▶ Industry standard for cluster-wide storage plugins
- ▶ Collaboration of communities incl. Kubernetes, Mesos, Docker and Cloud Foundry
- ▶ Defines the protocol between a CO and a plugin
- ▶ Plugins are CO-agnostic
- ▶ Write once, use everywhere





# Storage support in CERN's Kubernetes service



# We are here!

Introduction (to CSI)  
**CSI driver for CVMFS**  
CSI driver for EOSxd

Bonus: eosxd-csi + oauth2 = <3  
CSI driver for NFS  
Conclusion

# CVMFS in Kubernetes

- ▶ Need of accessing CVMFS repos from within Kubernetes workloads
- ▶ `docker-volume-cvmfs`
  - ▶ Deployed as a *DaemonSet*
  - ▶ Relies on FlexVolume
- ▶ CVMFS CSI
  - ▶ Cleaner integration into Kubernetes ecosystem

---

<code>docker-volume-cvmfs</code>	•	Feb 2016
<code>cvmfs-csi v1.0.0</code>	•	Jun 2018
	⋮	
<code>cvmfs-csi v2.2.0</code>	•	May 2023

# What's new in CVMFS CSI v2.2

[github.com/cvmfs-contrib/cvmfs-csi](https://github.com/cvmfs-contrib/cvmfs-csi)

- ▶ Automounts
- ▶ Dynamic config loading
- ▶ Per-volume configs
- ▶ Arbitrary volume as client's local cache
- ▶ v2.0 available in v1.22 clusters, v2.1 in v1.24, v2.2 next

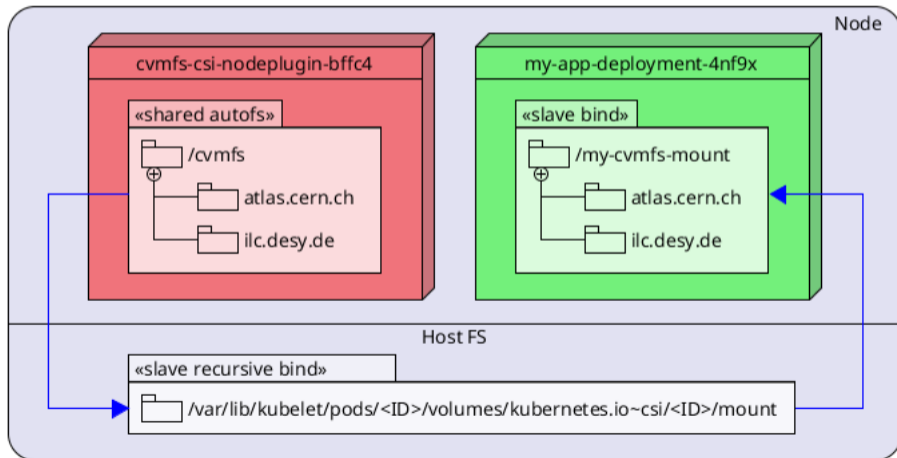


# Upgrading from CVMFS CSI v1 to v2

- ▶ 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-cvmfplugin 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.

# 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



# We are here!

Introduction (to CSI)  
CSI driver for CVMFS  
**CSI driver for EOSxd**

Bonus: eosxd-csi + oauth2 = <3  
CSI driver for NFS  
Conclusion

# EOSxd in Kubernetes

## EOSxd DaemonSet

- ▶ EOSxd mounts provided by eosxd DaemonSet (EOS root exposed as a hostPath in /var/eos)

[gitlab.cern.ch/kubernetes/storage/eosxd-csi](https://gitlab.cern.ch/kubernetes/storage/eosxd-csi)

- ▶ Replacement for eosxd DaemonSet
- ▶ Automounts
- ▶ Support for oauth-based authentication
- ▶ Room for improvements (provisioning like we do with cephfs?)





# We are here!

**Introduction (to CSI)**  
**CSI driver for CVMFS**  
**CSI driver for EOSxd**

Bonus: eosxd-csi + oauth2 = <3  
**CSI driver for NFS**  
**Conclusion**

# oauth2-refresh-controller

[gitlab.cern.ch/kubernetes/security/oauth2-refresh-controller](https://gitlab.cern.ch/kubernetes/security/oauth2-refresh-controller)

- ▶ Expose OAuth2 tokens as Secrets
- ▶ Refresh access tokens before they expire
- ▶ Inject access tokens into your Pods
- ▶ Annotation-based config
- ▶ ... works with eosxd-csi!



# We are here!

Introduction (to CSI)  
CSI driver for CVMFS  
CSI driver for EOSxd

Bonus: eosxd-csi + oauth2 = <3  
**CSI driver for NFS**  
Conclusion

# NFS in Kubernetes

[github.com/kubernetes-csi/csi-driver-nfs](https://github.com/kubernetes-csi/csi-driver-nfs)

- ▶ NFS integration is available in-tree but unmaintained
- ▶ nfs-csi as a replacement
- ▶ Important use-case for ATS, enabled by default in v1.26+ TN clusters

# We are here!

Introduction (to CSI)  
CSI driver for CVMFS  
CSI driver for EOSxd

Bonus: eosxd-csi + oauth2 = <3  
CSI driver for NFS

**Conclusion**

# Conclusion

- ▶ Improvements in many areas
- ▶ Unifying the last remaining storage bits to CSI

Questions?

- ▶ Robert Vasek <rvasek01@gmail.com>, <robert.vasek@cern.ch>