



# Status and Roadmap of the CernVM-FS Graphdriver Plugin for Docker

CERN, SFT Group Meeting

Nikola Hardi  
nhardi@cern.ch

October 16, 2017

Supervisors:

*Jakob Blomer* : jblomer@cern.ch

*Gerardo Ganis*: gerardo.ganis@cern.ch

# Project Overview

# Recognizing the problem (1/2)



## Joint Blog Post Mesosphere & CERN (03/2016)

*Network traffic gets congested as gigabytes worth of Docker downloads are moving across the pipe [...]. Companies [...] such as Twitter have already experienced this phenomenon.*

## Red Hat, “Containers for Grownups” (02/2016)

*10 things to avoid in docker containers:*

*...*

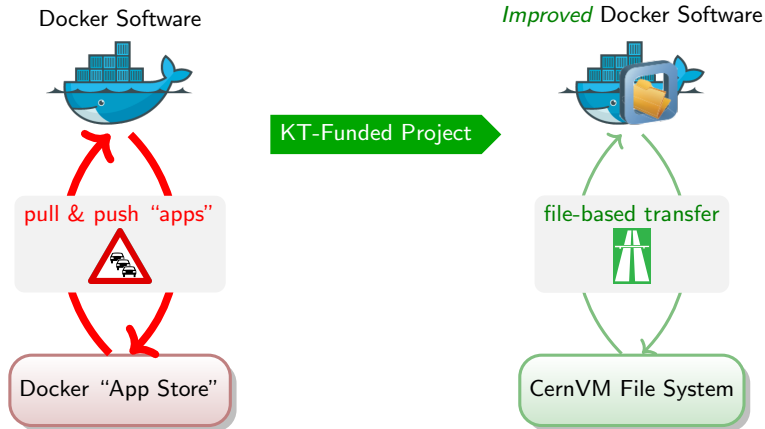
*3) Don't create large images.*

## Medallia (10/2015, CERN KT Screening)

*The problem today with Docker is that distribution of software is a mess, it is a “bottleneck” in our system.*

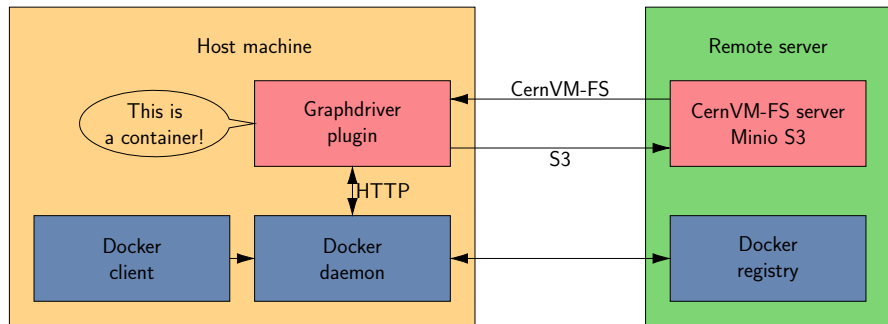
*Jakob Blomer, CernVM-FS as a Distribution Engine for Application Containers, 24/05/2016*

# Recognizing the problem (2/2)



*Jakob Blomer, CernVM-FS as a Distribution Engine for Application Containers, 24/05/2016*

# The CernVM-FS graphdriver plugin for Docker




```
$ docker plugin install cvmfs/overlay2-graphdriver
```

```
Restart Dockerd with --experimental -s cvmfs/overlay2-graphdriver
```

```
$ docker run cvmfs/thin_ubuntu echo 'Good bye EP-SFT!'
```


# How it works

## Regular Docker Image



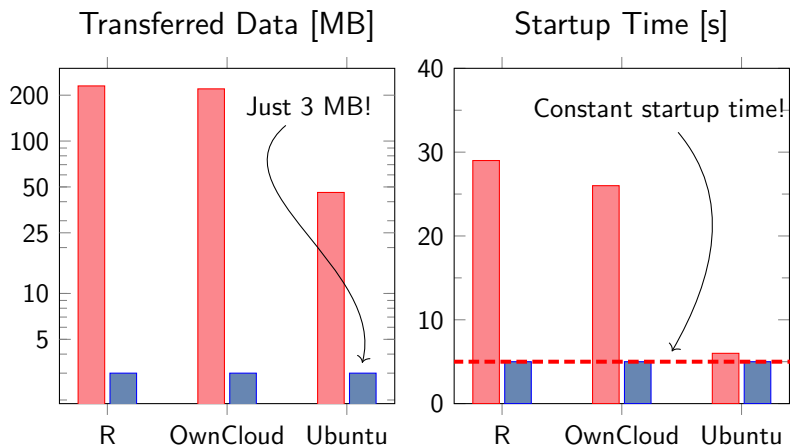
- Scratch layer
- Local read-only layer

## Thin Image



- Thin image descriptor
- CernVM-FS provided read-only layer

## Results



# Status Update



# The software components

- The plugin container, consisting of two derived graphdriver plugins
- The docker2cvmfs utility for creating thin images
- The layer upload and publishing utilities (Minio S3 + web hook) co-located with a CernVM-FS release manager.

# From prototype to pre-production

- ① Support for both relevant union file systems (aufs, overlay2)
- ② Standardized terminology: thin image, image descriptor and layers
- ③ Integration test suite
- ④ Usability improvements
- ⑤ Many bug fixes from real-world tests
- ⑥ Reproducible builds
- ⑦ Packaging and versioning of the different components

# Prospective Collaborations Applications

# Prospective Collaborations Applications

## At CERN:

- CernVM-FS build jobs
- EP-SFT Jenkins
- CERN IT OpenStack team
- (The Batch Services team)

## Outside of the HEP:

- Nexenio
- The Docker team  
(BuildKit and containerd maintainer)

*" Except the hard drives that need to store all those **fat images** [...] and the **network latency** that comes with it. "*


*" **Astonishing results:** Running the application with thin image: **instantaneous.** "*  
***No evident network traffic** when using the CernVM-FS graphdriver plugin.*

# Results of the cloud-benchmark-suite (1/2)

The cluster startup time reduced from 5 min to less than 5 s.




**About 5 minutes until first services are available.**



**Less than 5 seconds until first services available!**

## Results of the cloud-benchmark-suite (2/2)



# Publications and Dissemination

- Published an RFC at Docker Github issues page in November
- Presented the project at the ACAT 2017 in August
- Presentation and demo by CERN IT at the Dockercon Europe (starting today in Copenhagen!)
- Master thesis - starting in November (next week)



# Roadmap and Summary



# From pre-production to release

- ① Merging integration tests with existing CernVM-FS CI routines
- ② Extending S3+webhook ensemble into a "CernVM-FS portals" system.  
E.g. the ability to send "packages" (e.g. RPMs, tarballs with user software, docker layers etc.) via S3 to a CernVM-FS release manager for automatic publishing.
- ③ Easier setup for the server side: cvmfs\_server portal.
- ④ The docker2cvmfs tool push command.
- ⑤ Local only commits on thin images.
- ⑥ Automatically create thin images on demand.
- ⑦ Improve handling of special files in CernVM-FS.
- ⑧ Server side garbage collection

# Summary

- The plugin works and solves a real problem.
- First users have started testing the plugin with encouraging results.
- The plugin is already fully functional but usability can be improved.
- One of the first Docker v2 plugins and especially graphdriver plugins.
- Works on Centos 7.4 and CernVM .
- Requirements:
  - Docker 1.13+ (January 2017) and
  - OverlayFS or AUFS support in kernel.

# Thank you and Goodbye!

Thank you @EP-SFT for an amazing year I had!