CernVM-FS Container Integration

Jakob Blomer

December 13, 2017

HSF Technical Forum – Packaging Discussion

Containers Grow Big

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.

What we want

1. Docker (container) for the isolation and the tooling
2. CernVM-FS for the image contents distribution
   and node's automatic cache management

   Works well because only a few percent of typical images are needed at runtime

CernVM-FS Graph Driver Plugin

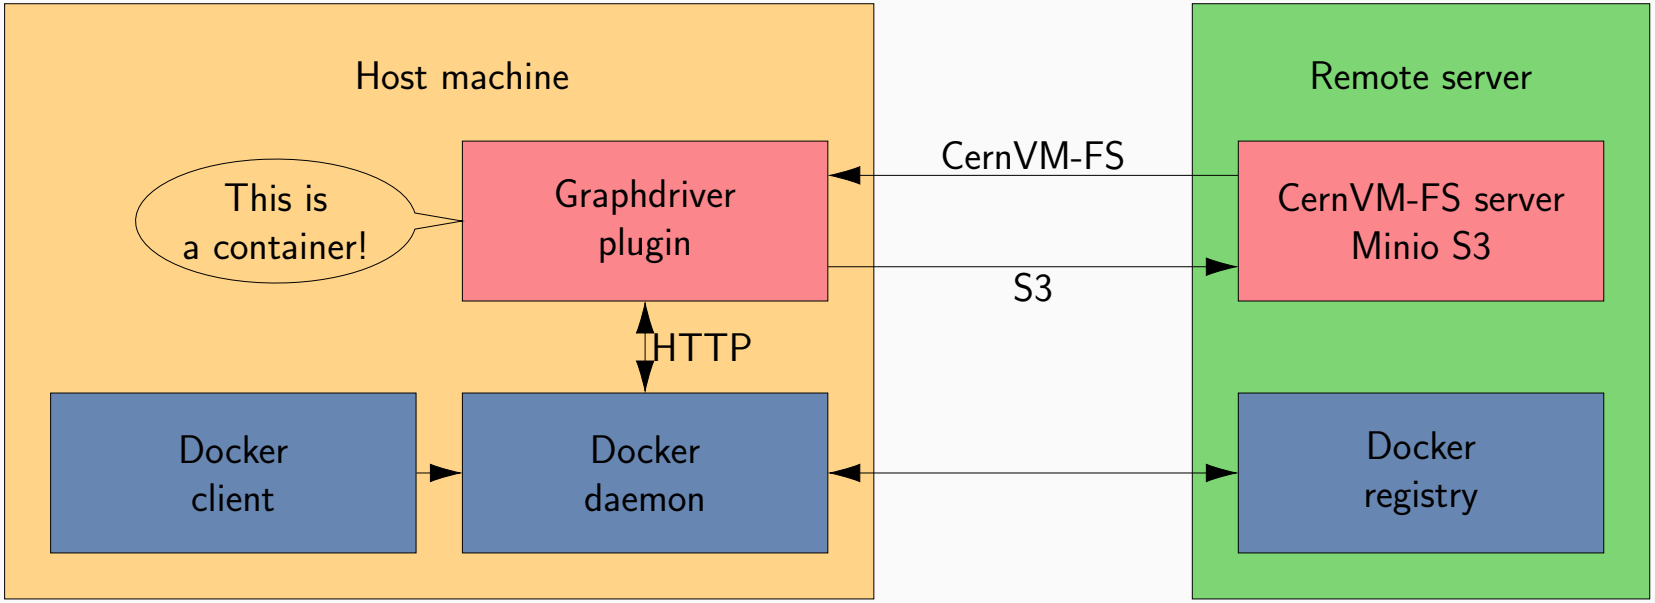

Note: Singularity and CernVM-FS work together out of the box!

Thin Images

Regular Docker Image

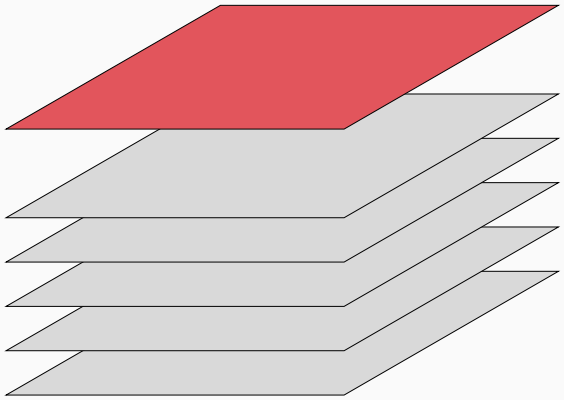

- Scratch layer
- Local read-only layer

Thin Image

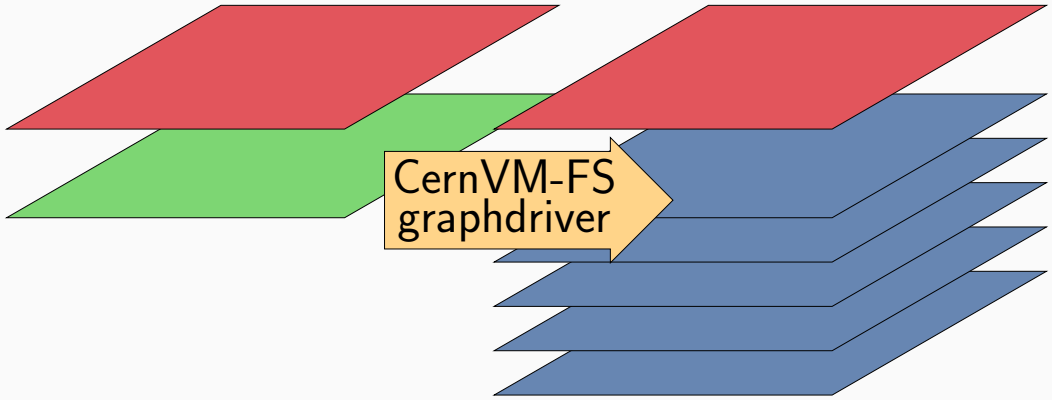

- Thin image descriptor
- CernVM-FS provided read-only layer

Demo

Benchmarks (Carried out by CERN IT)

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

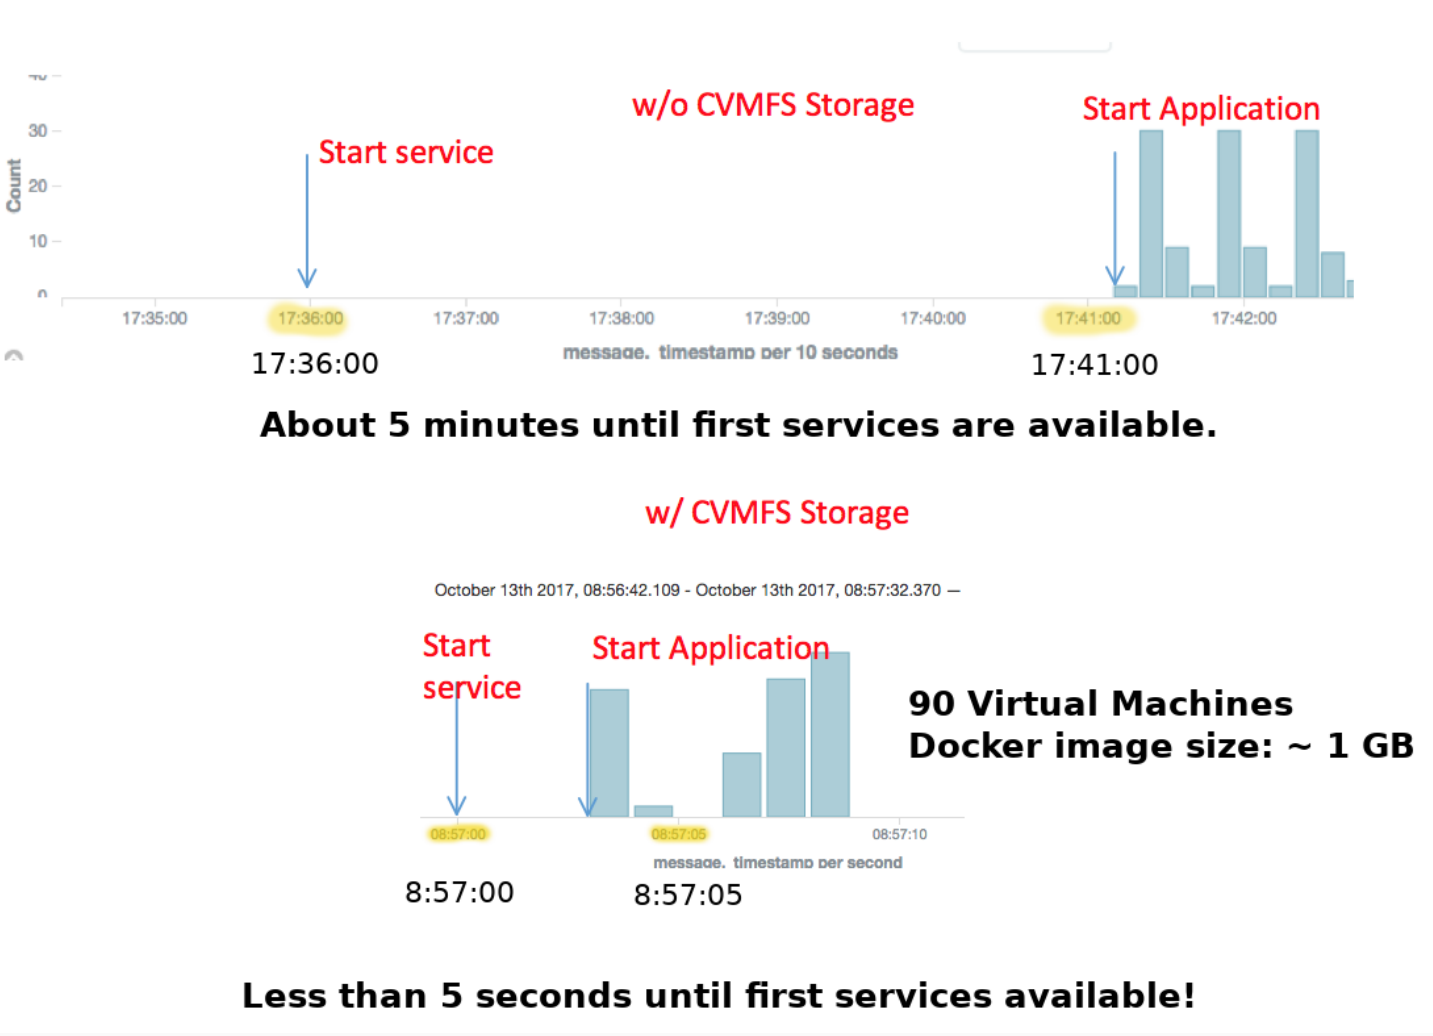

Status and Future Work

Plugin container	pre-production
Server-side portals (S3 endpoints for publishing)	prototype
docker2cvmfs	prototype
`docker push` support	prototype
containerd integration	work with upstream[1]

[1] containerd is the engine of Docker. It can be used independenly of Docker. Kubernetes might move to a containerd-only deployment. Containerd maintainers are working on file-granularity image standards, which would benefit us greatly.

Using the existing S3 Support in CernVM-FS, we could provide an **HSF container service** in AWS.

Say: **hsfhub.cvmfs.io**