

MAKING CONTAINERS LAZY WITH DOCKER AND CERNVM-FS

Nikola Hardi, Jakob Blomer, Gerardo Ganis, Radu Popescu

CERN

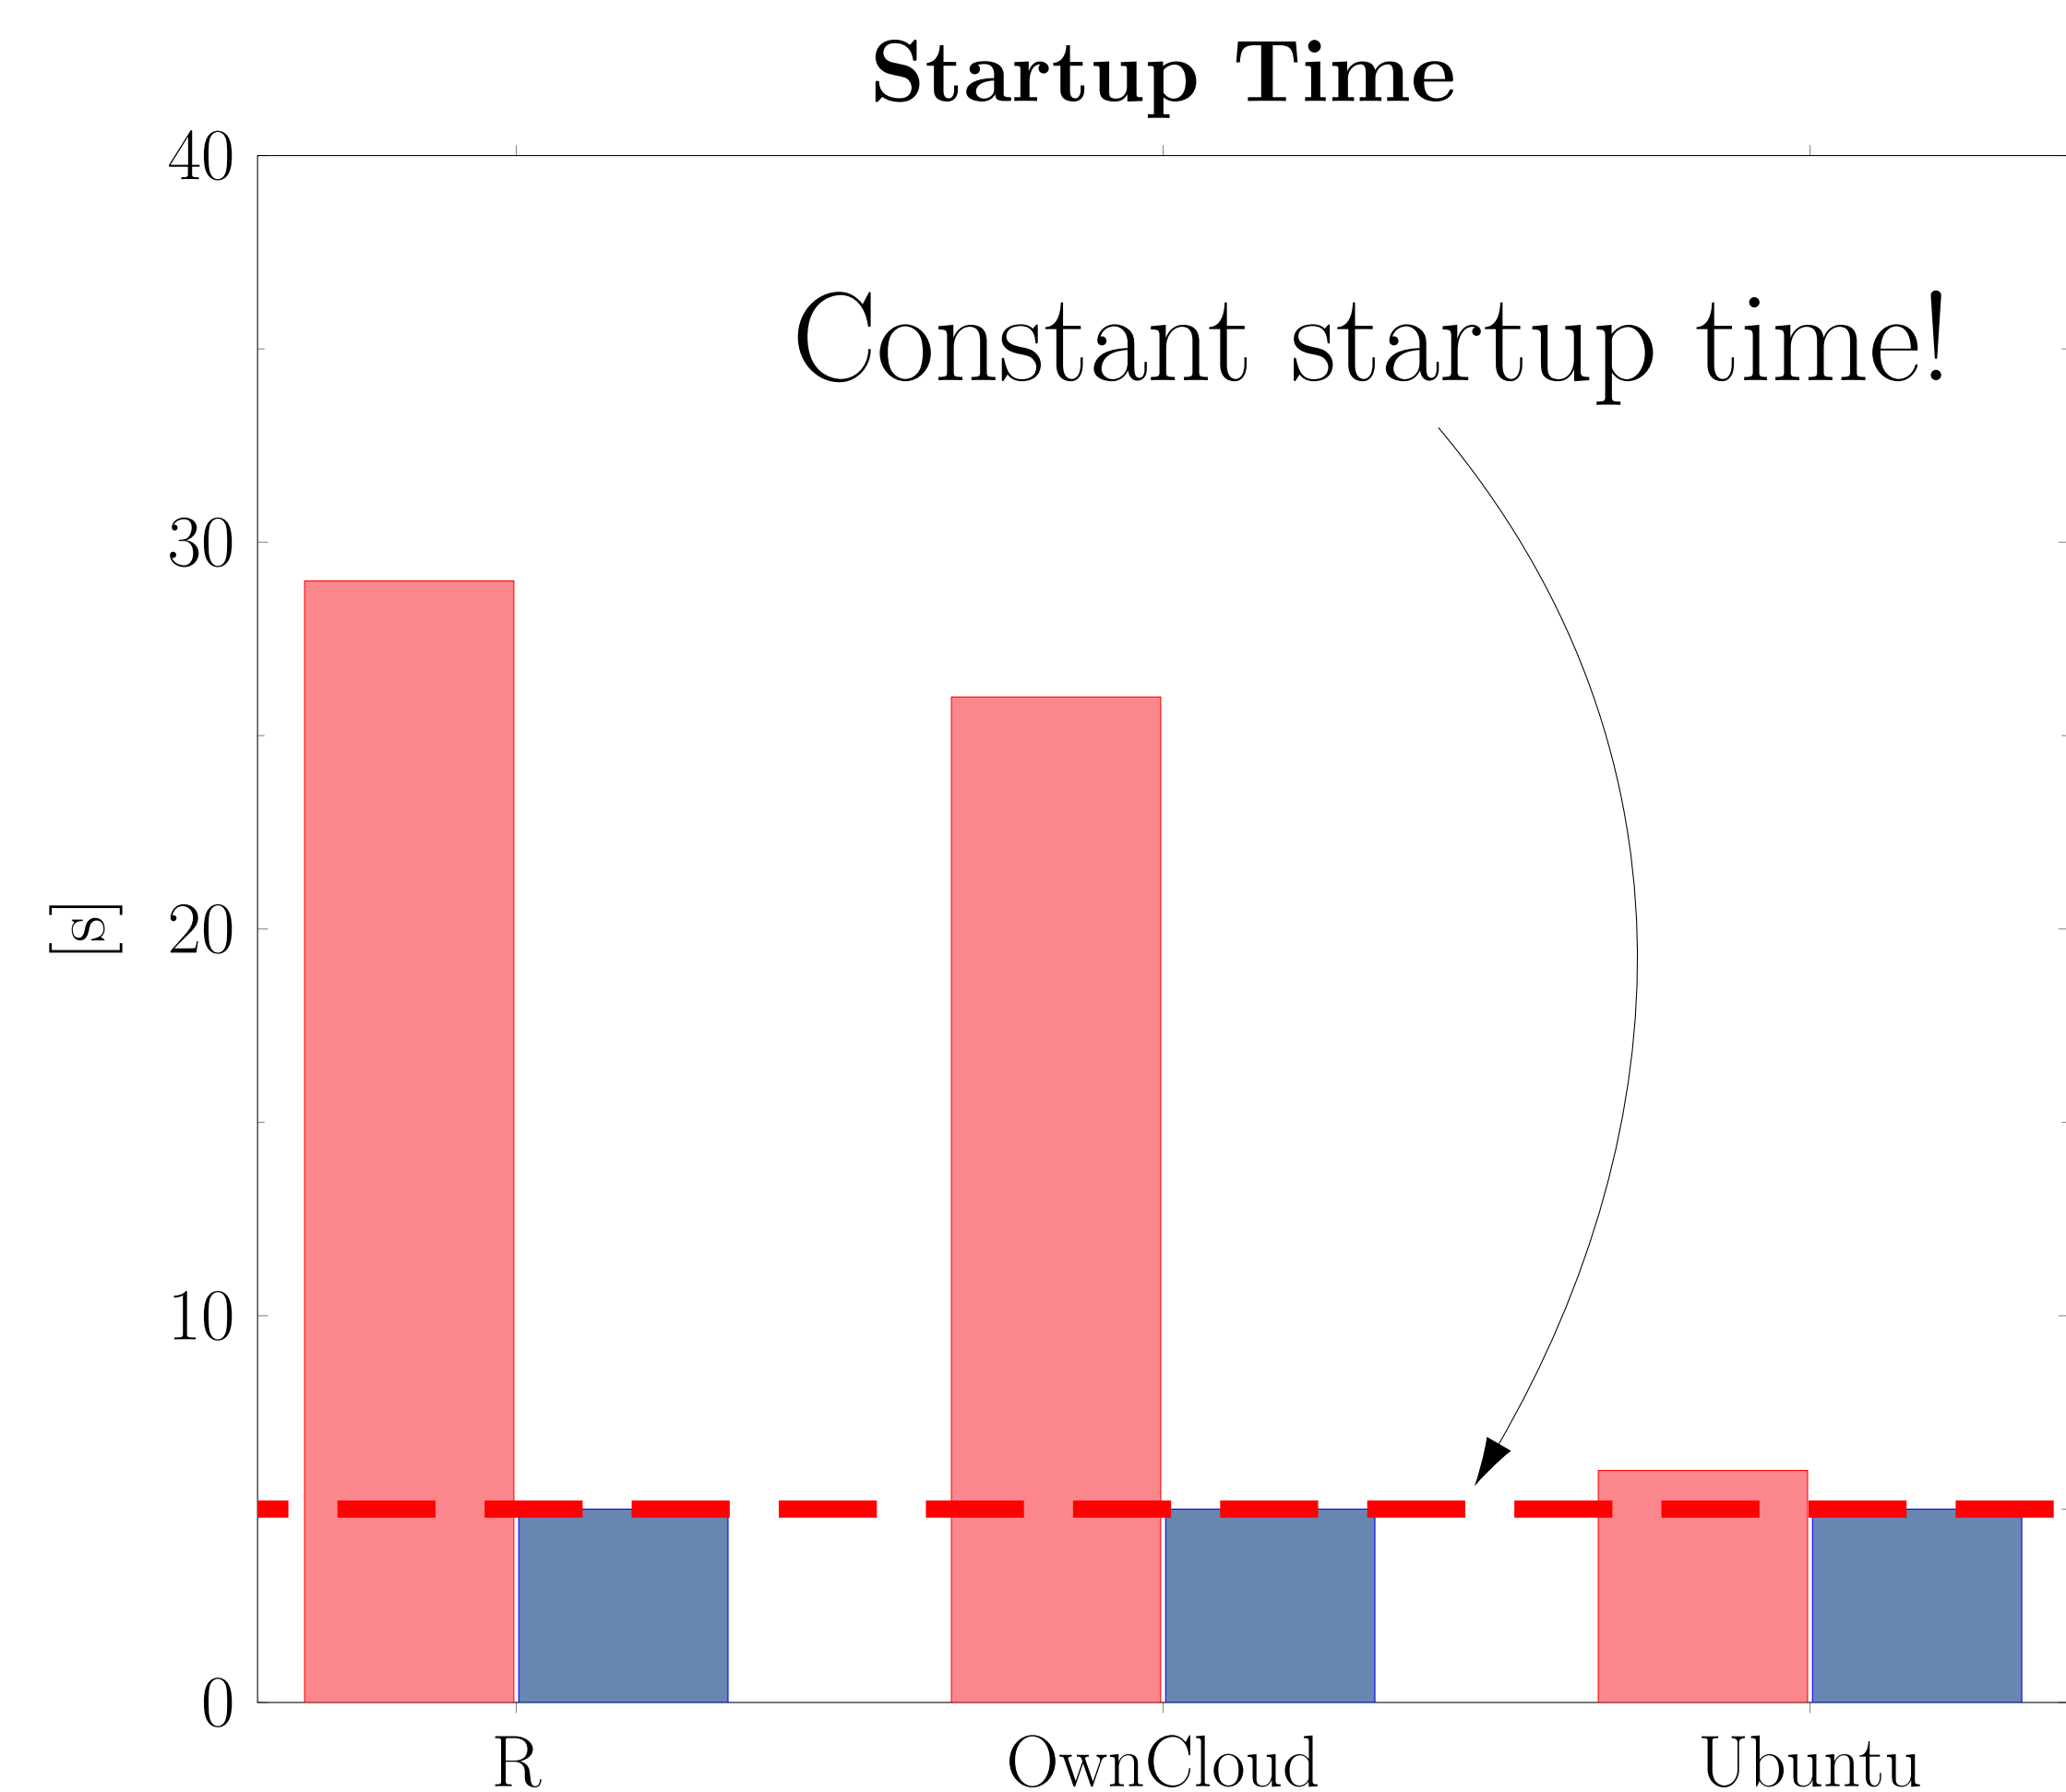

Docker Images Clutter Data Center Networks!

Docker is a popular containerization technology both in industry and in science because it makes containers easy to create, manage and share. Docker images are composed of shared layers to save space and bandwidth, but **the price for obtaining image is very high**. For instance, to start a cluster of 1000 nodes using an image of 1 GB your network has to transfer 1 TB at once and as quickly as possible.

CernVM-FS Offers a Solution!

On average only 6% of the image is required [1]. We developed a Docker plugin [2] to **refine data reuse granularity from layers to files**. Additionally, the **download is delayed** until the file is accessed for the first time. Using this approach and CernVM-FS [3], only a fraction of the image is initially transferred and containers start instantly.

Instant Startup by Downloading Only What is Needed



Our measurements show that by using the CernVM-FS plugin **containers start instantly**. Less data is transferred and the transfer is overlapping with the lifetime of the container.

■ Standard Docker
■ Docker with CernVM-FS plugin


Thin Images

A thin image is a regular Docker image except the data is organized differently. The thin image descriptor (a JSON snippet with metadata) replaces local read-only layers, and the contents of the layers are stored in a CernVM-FS repository. The thin images are inspired by Slacker [1].

Compatibility


- Thin images seamlessly integrate with other Docker components.
- Any Docker image can be converted to the thin image.
- Our graphdriver also works with all regular Docker images.
- CernVM-FS graphdriver is required to run thin images.

Regular Docker Image



■ Scratch layer
■ Local read-only layer

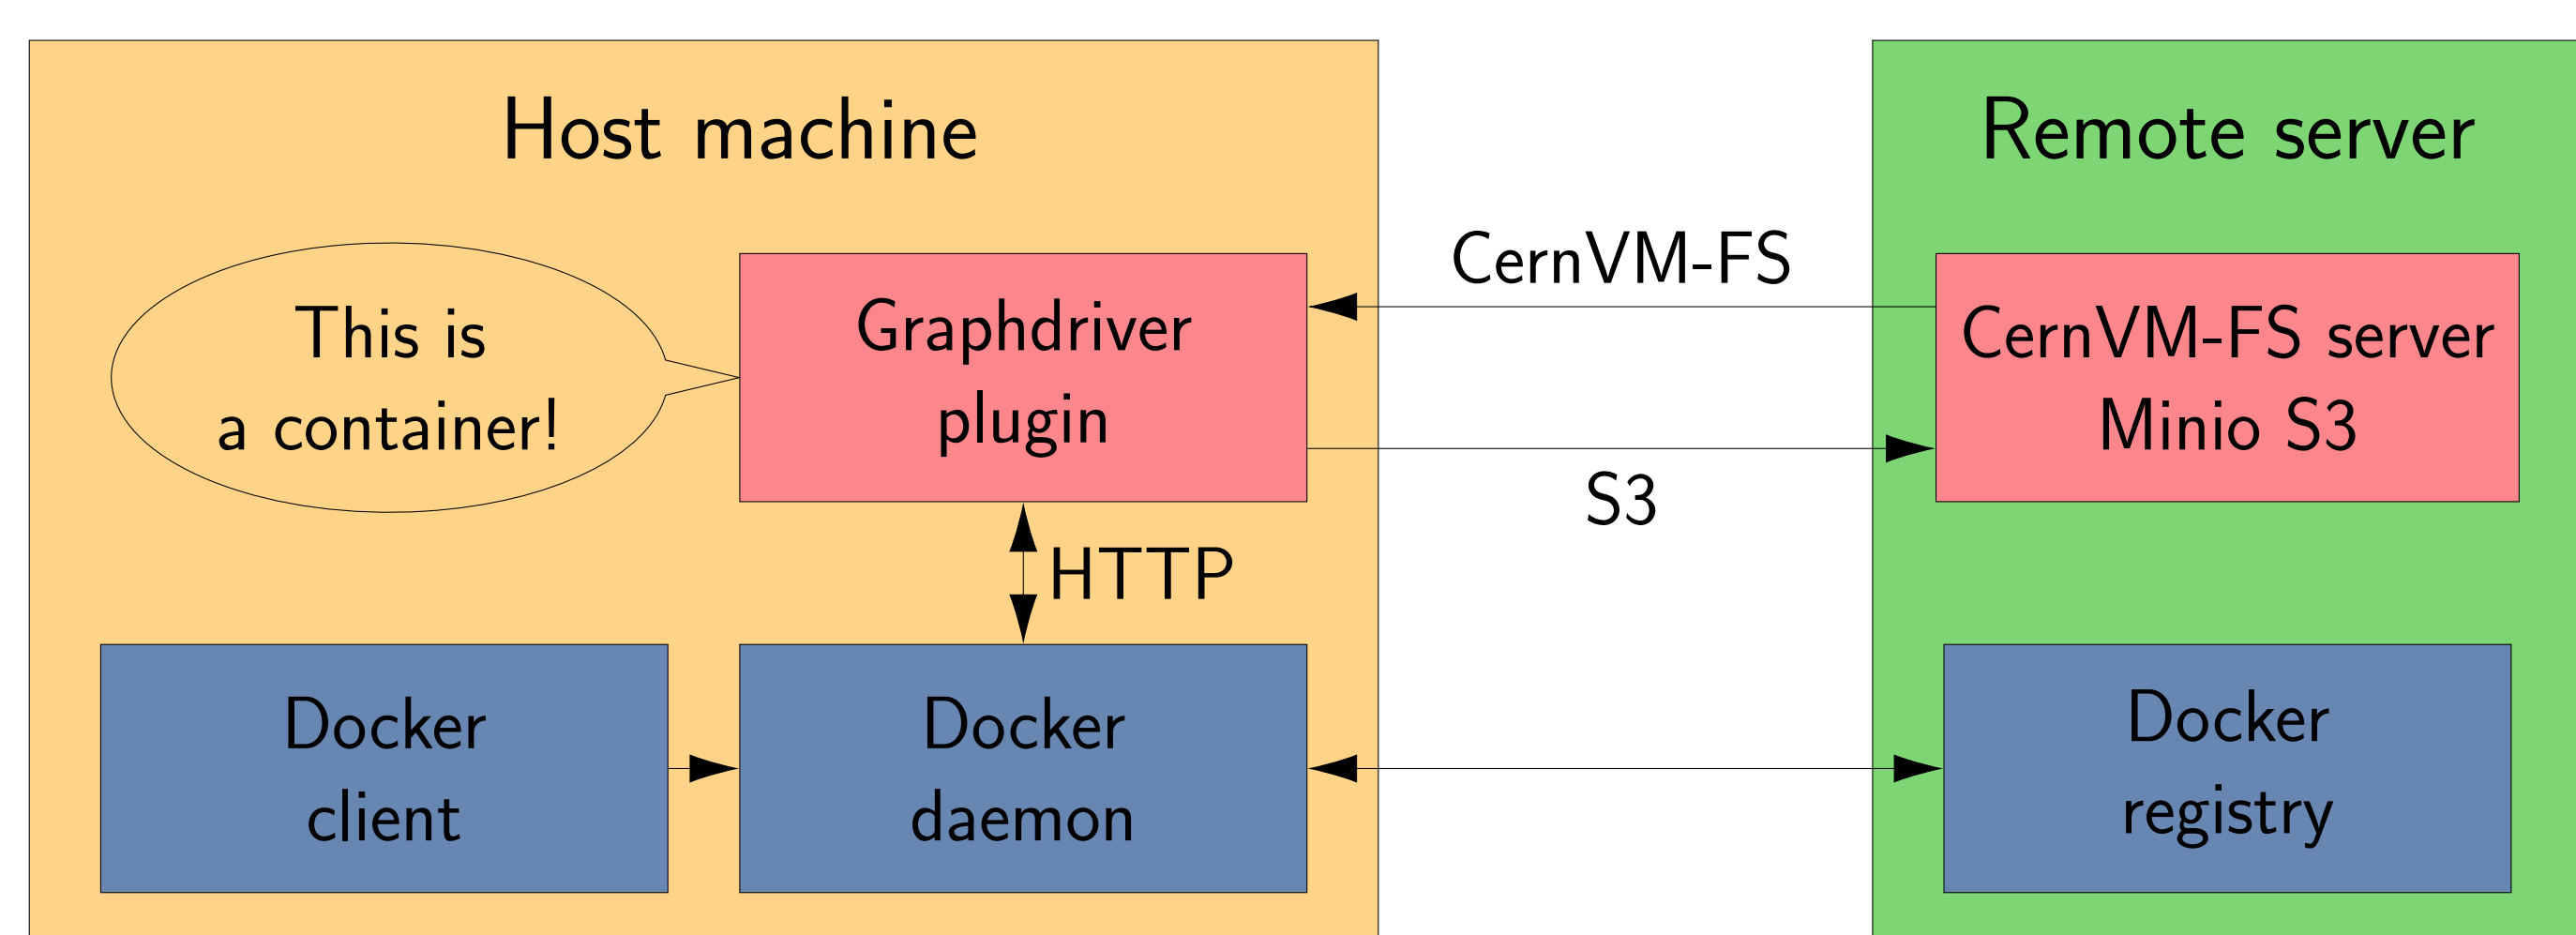
Thin Image



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

System Overview

Docker plugins are running in a containerized environment and communicate with the Docker daemon through an HTTP API. The CernVM-FS graphdriver plugin comes with built-in CernVM-FS and Minio S3 clients.



Try it!

The CernVM-FS plugin is easy to obtain, easy to use and does not interfere with your existing Docker installation. You can test a read only configuration or set up your own CernVM-FS powered Docker registry and publish a new thin image.

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

Restart Dockerd with --experimental -s cvmfs/overlay2-graphdriver
$ docker run cvmfs/thin_ubuntu echo "Hello ACAT2017"
```

Related Work

CernVM-FS Docker volume driver [4] lets users mount any CernVM-FS repository in a Docker container. Volume drivers manage user data while graphdrivers manage container layers.

Singularity [5] can start containers from images extracted into directories on CernVM-FS. This approach also results in lazy containers.

This project was funded by the **CERN Knowledge Transfer Fund**.

References

- [1] Slacker, USENIX FAST '16
- [2] <https://github.com/cvmfs/docker-graphdriver>
- [3] CernVM-FS, Computing in Science and Engineering, '17
- [4] <https://gitlab.cern.ch/cloud-infrastructure/docker-volume-cvmfs/>
- [5] <http://singularity.lbl.gov/>

