# DUCC Daemon to Unpack Containers in CVMFS

**Authors: Simone Mosciatti (simone.mosciatti@cern.ch)**, Jakob Blomer, Gerardo Ganis & Radu Popescu

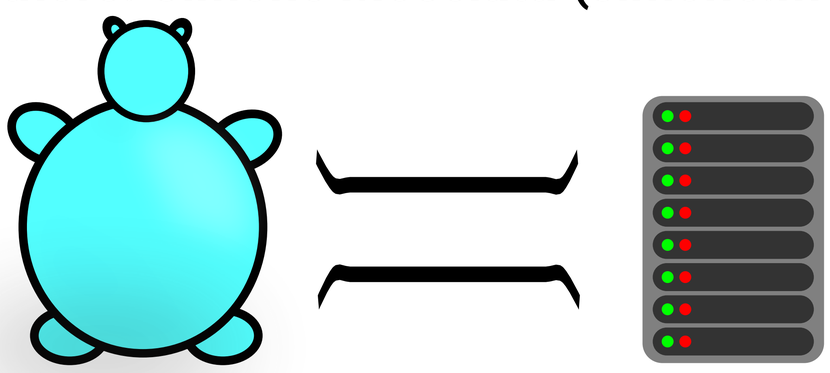

## 1. The Problem

**Scientific container images can be very large (>10 GB) and are hard to deploy at scale where bandwith is limited.**

## 2. Our Approach

**CVMFS can efficiently distribute container images in their unpacked form [1, 2]. DUCC manages the lifecycle of such images. It publishes a Docker image as a tiny ~KB "thin" image, a set of references to the original content of the image. The thin image is then sent to the machines executing the workload.**

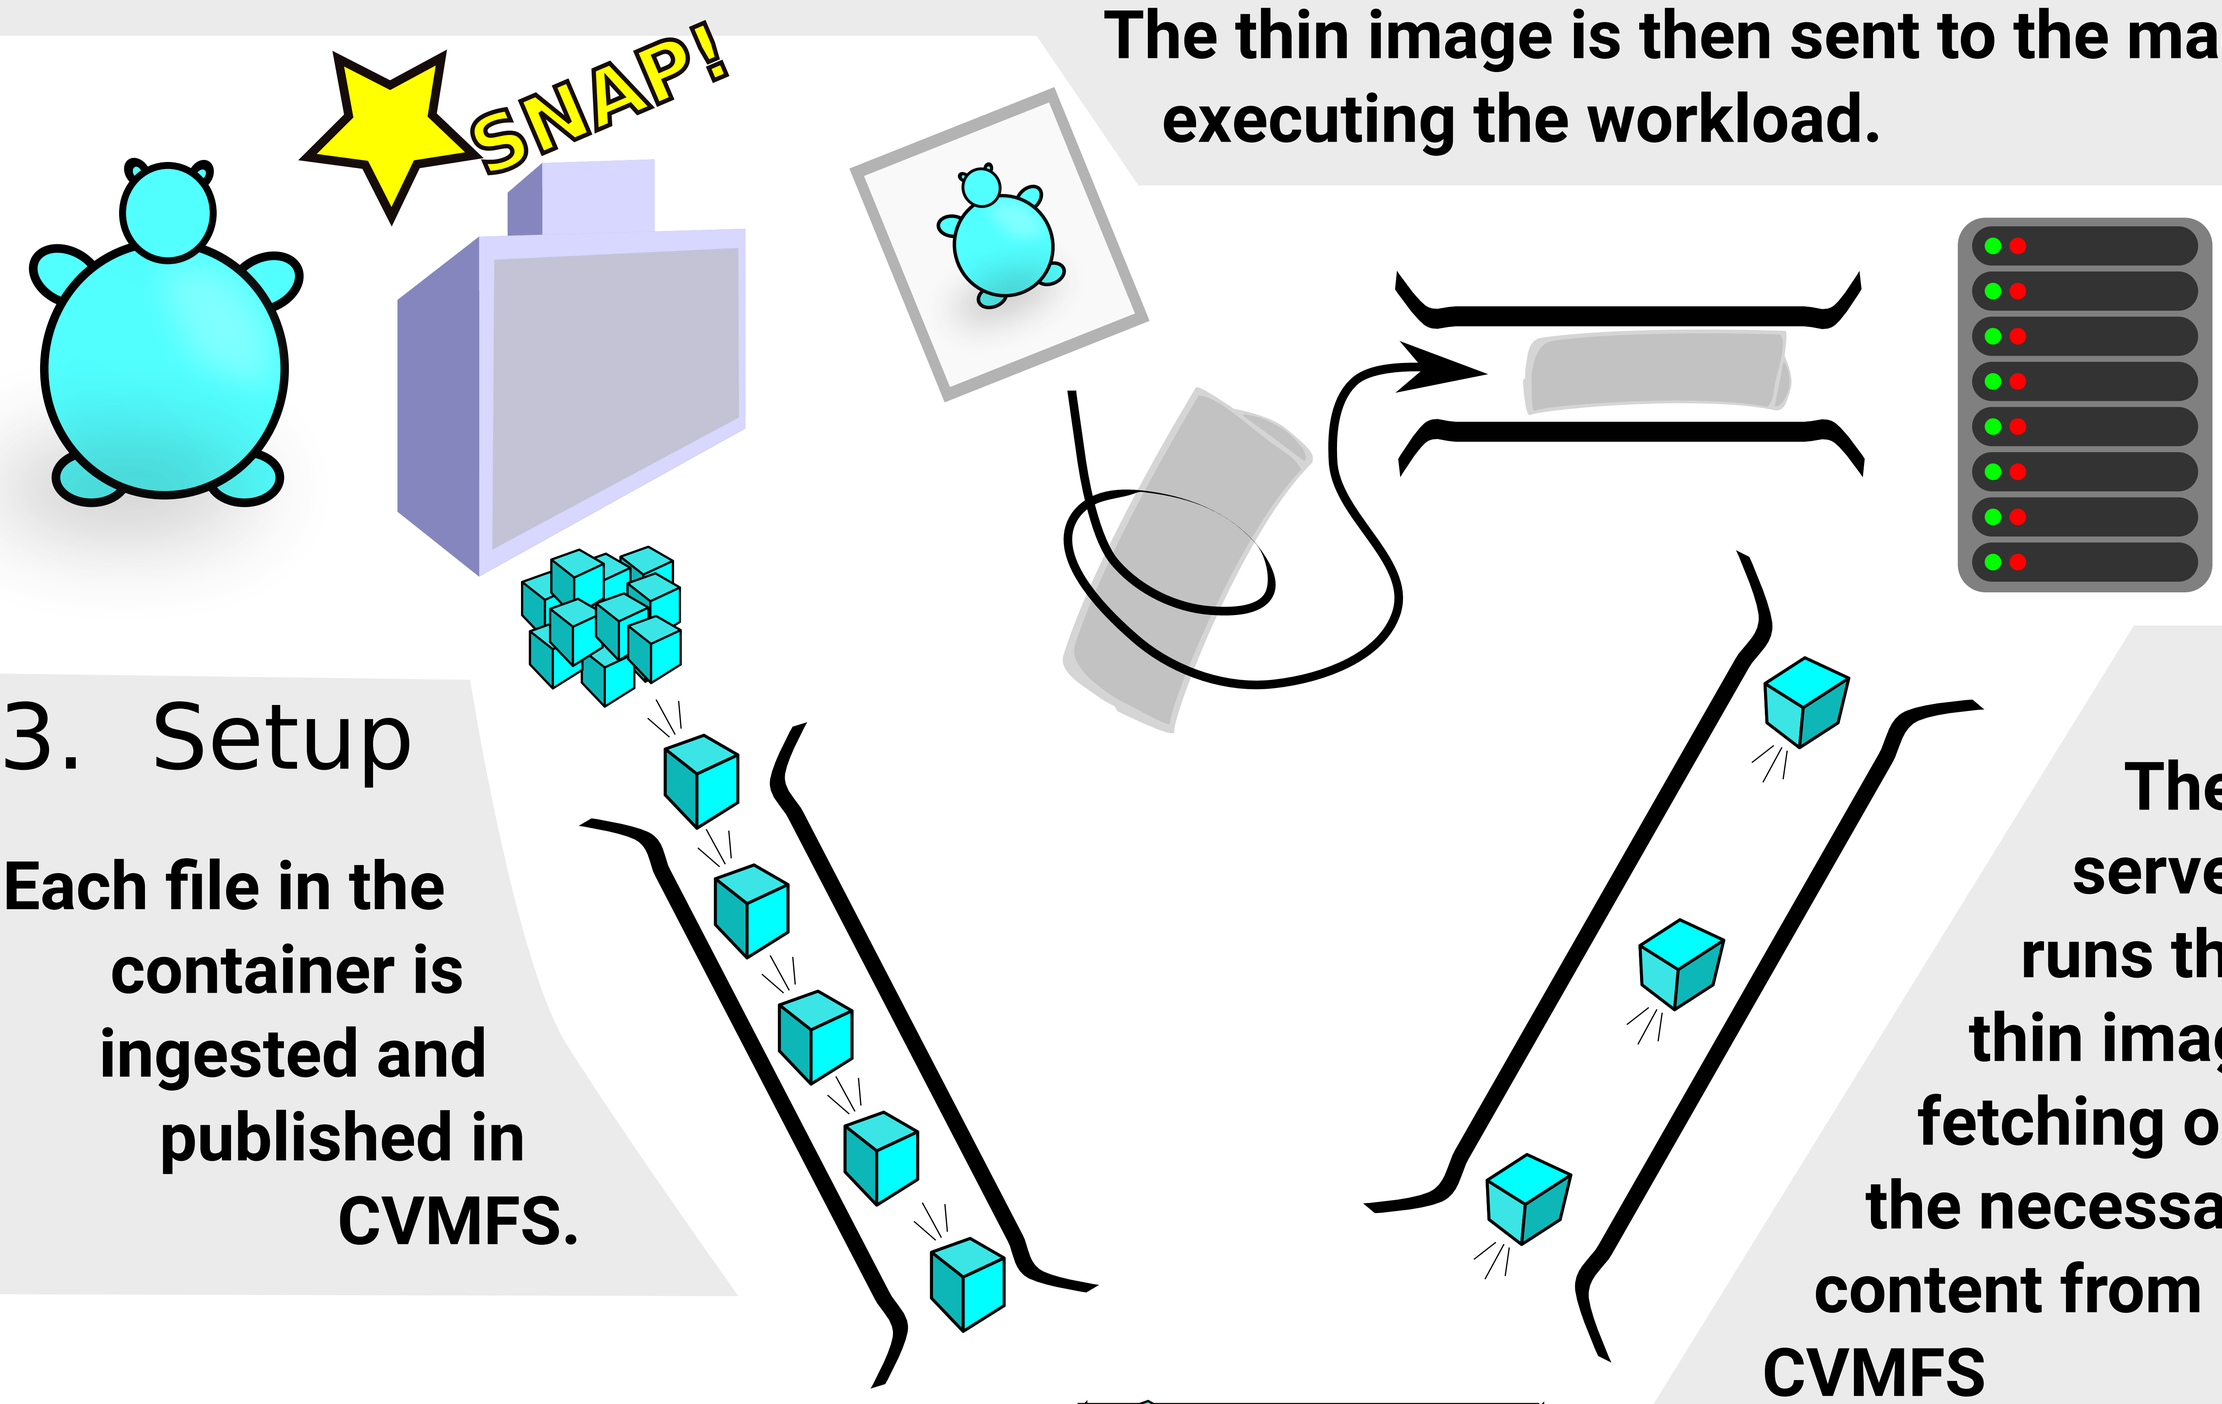


## 3. Setup

**Each file in the container is ingested and published in CVMFS.**

```
/cvmfs/unpacked.cern.ch
├── hub.docker.com
│   └── atlas
│       ├── analysisbase:21.2.60
│       └── athena:21.0.31
└── .layers
    ├── 00
    ├── a4
    └...
```

## 4. Runtime

**The server runs the thin image fetching only the necessary content from CVMFS**

## 5. Run anywhere

**We publish the images as a flat root filesystem along with every layers to support Singularity, Docker and containerd engines.**

## 6. Maintenance

**When an image is not necessary anymore, it is garbage collected after 30 days.**

**Now in CVMFS 2.6**

1) N Hardi, J Blomer, G Ganis, and R Popescu. “Making containers lazy with Docker and CernVM-FS.” In: Journal of Physics: Conference Series 1085.3 (2018), p. 032019. url: http://stacks.iop.org/1742-6596/1085/i=3/a=032019
2) Open Science Grid: https://github.com/opensciencegrid/cvmfs-singularity-sync