



# Container Working Group: Kick-off Meeting

Gavin McCance, CERN IT

***No rubbish here, please  
use containers outside.***



***Pas de déchets ici, mettez  
les dans les containers à  
l'extérieur svp.***

# Mandate

Discussed and agreed:

“””

To understand and track the use of containers within WLCG for compute deployment, and to make recommendations to sites on potential deployment patterns that they may wish to adopt.

To ensure that Singularity and its support services are deployable for the LHC experiments.

“””

# How do we know when we're done?

We agreed 2 main goals:

1. Singularity is sensibly deployable(\*) in WLCG and used by all experiments, and we've documented and understand the operational aspects.

(\*) actual deployment tracked by WLCG ops

# How do we know when we're done?

2. We should have surveyed and made recommendations on potential deployment patterns (and technologies) for container-based **compute deployment and provisioning**, with an aim to reduce operational effort on the sites, and to meet the needs of the ongoing analysis reproducibility work going on in the experiments.

At the end, we should have clearly documented a few reasonable paths for a site (should it wish to) to migrate their compute to these technologies.

# What were not doing in the WG

- Forcing all sites to move their compute to containers
- We're focussing on containers for *compute*. While we may inevitably touch a bit on containers for *services*, it won't be the main focus.
- Actual deployment tracking of Singularity belongs with WLCG ops.

# This pre-GDB

- Focus was on Singularity
  - Relationship to Traceability & Isolation WG
  - Open operational issues, concerns
  - Experiment adoption status, concerns, open questions
  - Are we aligned on how we want to create/distribute the images/chroots? Are we aligned on how we want to configure it?
  - Make a plan of things we need to address



# Language!

- Ian noted a language point:
- Words like “container” and “image” are used in many contexts wrt. Docker, Singularity and all the various deployment patterns we might imagine
- Early task of WG is to carefully define what we mean by these terms...

# Relationship to Traceability and Isolation WG

- Singularity handed over to Containers WG as solution to payload Isolation
  - Offers namespace isolation and OS portability
  - Containers WG will get on and look at the deployment aspects
- Traceability remains an issue for T&I WG
  - How to suspend bad user (VO?)
  - Need for Argus?
    - Some sites rely on the job's UID on the WN also for storage operations → can't use “file” protocol without Argus.
      - Jobs can use GridFTP, Xrootd, ... instead.
- Desire to go to non-SUID mode in RHEL7.4

# Singularity Issues 1

- A few irritations for now
  - Some automount workarounds are needed
  - ...
- Practical issues
  - RH7 batch environment vars gets handed to RH6 container
- Singularity still moving fast - as a community, we need to engage with upstream!

# Singularity Issues 2

- Biggest one is bind-mounts paths
  - Need to pre-exist inside singularity container
  - overlayFS fixes it on RHEL7, but unlikely to be able to use this non-SUID
  - Lots of discussions on this and whether we can achieve a common solution given variety of site configurations
  - Site config vs. experiment site config (e.g. AGIS)

# How to distribute images

- General agreement on unpacked CMVFS “chroots”
  - Don't drag-in the entire image file
  - Benefit from de-dup
- ... but some sites, notably HPC, need single img files, and some want those to checksum in advance
  - Ideally would have common workflow that produces both

# What are these images anyway?

- Mostly standard “experiment” pilot images at the moment
  - Few, curated
- Some desire for specials (e.g. final analysis preservation)
  - To be understood
- So far, each experiment has prepared their own ones

# Common singularity images?

- Lots of discussion on this
- A few approaches discussed
- Can we use CERNVM as the basis?
  - Eases the lifecycle/update issue
  - Really non-experiment specific? Maybe (LHCb says yes)
  - Failing exact same images, a common starting point and workflow would be good

# Singularity activities

- Collect pointers to current deployments within WLCG and experiments
  - e.g. investigate mount point issues
  - collect current experiences
- Encourage and track deployment on volunteer T1 sites (notably multi-VO ones), and collect pain points
- Start looking at how we might converge on a common image (base)



# Join

- Join up at:

[wlcg-containers@cern.ch](mailto:wlcg-containers@cern.ch)

(sign up at <https://egroups.cern.ch>)

- Will coordinate activities on the list
- Next meeting will be in Autumn
  - Updates on Singularity and CVMFS
  - More general container activities

