

Containers WG report

November 2019 GDB

Recent meetings

- Last GDB update was February this year
- One meeting in summer and then focused on writing a container baseline
- Baseline doc now agreed by WG:
 - https://twiki.cern.ch/twiki/pub/LCG/WLCGContainers/WLCG_Containers_Working_Group_Baseline.pdf

Job containers

- Simple command line to run your job inside a containerised environment
- Isolation for multiple payloads of the same pilot
- Separation of job's OS from system one. We typically mount the container OS from an unpacked directory in CVMFS
- Currently one tool, "Singularity", is in use, providing a common experience across SLC6 / CC7

Singularity Baseline

- WLCG sites must either:
 1. Enable unprivileged user namespaces (\geq EL7.6), or
 2. Install a local setuid Singularity enabling underlay and/or overlayfs support

Option 1. Enable unprivileged user namespaces

- In this case, no local Singularity need be installed, (though feel free if you also need it for other non-WLCG users)
- All WLCG experiments will use their own Singularity binary from CVMFS directly, even if it's also installed locally
- This is the easiest option, assuming your site is happy to enable unprivileged namespaces in the kernel
- <https://opensciencegrid.org/docs/worker-node/install-singularity/#enabling-unprivileged-singularity> describes the process rather well

Option 2. Install Singularity setuid locally

- Available in EPEL
 - 3.4.2-1.1 current, as of writing
- Ensure overlayfs (\geq EL7) or underlay is enabled to allow mounting of directory images hosted in CVMFS (needed in cases where you want to bind-mount to a dir that doesn't exist in the image)
 - These are on by default in Singularity 3

Image distribution

- D1. Experiments are encouraged, where possible, to distribute their unpacked images via CVMFS using the **unpacked.cern.ch** or **singularity.opensciencegrid.org** CVMFS repos and associated services to ensure efficient access and de-duplication
- D2. Sites are recommended to provide CVMFS access to these repos on their worker nodes

Working group

- Done its main job
 - WLCG operations should track the status of support on WLCG sites with respect to the two options
- Still some discussions wanted on how we distribute (many) user images in the most efficient way
- Mailing list:
 - wlcg-containers@cern.ch