

# Singularity deployment

---

A. Forti

GridPP40

11 April 2018



# Recap

- This is a follow up to the talk I gave at GridPP39
- Decided to wait for new pilot (Pilot2.0) to do any development in the pilot.
- Testing using the European pilot wrapper while waiting for the new pilot
  - This was very helpful to get started
  - Triggered a unification of different wrappers
- Since January the new pilot Pilot2.0 is being tested on the grid and this is what happened since then



# AGIS parameters

- Two parameters
  - `container_type` used to express two things
    - Which container to run and at which level: pilot, wrapper, batch system
    - If the middleware has to run in containers or it is available on site
    - Example: `singularity:pilot;middleware:container`
  - `container_options` used to pass options to the container if needed
    - Not always needed
    - Example: `-B /scratch -B /cvmfs -B /home/condor`
- Parameters can be left empty and standard payload will run



# Overlayfs vs unprivileged

- WLCG traceability and isolation WG is pushing for using singularity in unprivileged mode
  - Unprivileged mode while it restricts what users can do it also restricts an important feature called overlayfs
  - Overlayfs allows to bind the host directories (bind paths) in the image without the image containing those directories.
  - There are work around but they don't work everywhere and they don't work for everyone
    - Singularity configuration gets applied to everyone
- ATLAS is working to have the flexibility to work in both modes.
- WLCG policy finalization is under discussion.
  - Statement on unprivileged is a SHOULD not a MUST



# Pilot and bind paths

- Two ways of using container implemented in Pilot
  - Explicit wrapping of payload command with singularity exec
    - Bind mounts either in AGIS or in the site config
  - Containerization of the ATLAS setup (ALRB)
    - Decides the bind mounts based on existing env vars and uses \$PWD:/srv
    - No need to put any option either in the configuration or in AGIS
      - Hopefully for most sites
      - Hides most of the complication



# Pilot testing on grid sites

- Pilot factory wrapper updated for Pilot 2
  - Successfully running container test jobs at DESY-HH\_TEST and UKI-NORTHGRID-MAN-HEP\_TEST\_SL7
  - Testing at UKI-SCOTGRID-GLASGOW\_TEST, UKI-LT2-IC-HEP\_TEST\_SL7, OU\_OSCER\_ATLAS\_TEST and BNL\_PROD\_CONTR\_TEST underway
- Some site configuration issues remain to be sorted out
- Tests run so far executed payload in container



# Images

- In CVMFS
  - Both single file and unpacked
- Decided to have base images (OS+HEP\_OSlibs only)
  - Add use cases requirements on other images
    - T3 interactive
    - HPC
    - Images already loaded with pilot code
    - .....
- Several images how does the pilot chooses?
  - For now image decided from the release platform
    - Platform='x86\_64-slc6-gcc48-opt' ==> slc6 image
  - This needs to be discussed more though



# Site or pilot

- Some sites are running containerized WNs.
  - More normally Docker but also singularity
- Running containers in containers is possible but **only** if the external container is docker.
  - Containerized WNs with singularity and pilot singularity are **incompatible**.
  - The job will still run if the pilot detects it's already in singularity it just runs a normal payload.
    - But it is not what we are aiming at
    - AGIS container\_type should be set to singularity:batch
- Running singularity in the batch system is not a requirement
  - Neither from ATLAS nor from CMS





# Singularity Version

- Latest singularity version 2.4.2 has become the baseline in WLCG.
  - <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGBaselineTable>
  - RPMS available from OSG and WLCG repositories
- Sites should upgrade (or install the latest) for two reasons
  - Old executables don't read the images built with the new executables.
    - ATLAS latest images are built with the latest executable because it uses squashfs and the images are smaller
- A new version should come out soon if you haven't installed singularity yet wait for it.
  - It contains fixes for a security bug for overlayfs



# Should sites install it?

- CMS has already required all their sites to install by the 1<sup>st</sup> of March on both EL6 and EL7 like OSs.
  - Many are shared sites
- ATLAS not in a rush
  - Code works on EL7 family and sites can migrate without trouble
  - We have now enough sites to test (but the more the merrier)
- Still more a recommendation than a requirement
  - Sites moving to EL7 OS family should install singularity and familiarize with it



# Future

- Most important points for the moment are
  - Continue with pilot2 testing
    - Getting it to work on all the available testing queues
    - Add new features to pilot2 to handle different movers and middleware image
    - Add aCT sites testing
  - Discuss more about images selection and maintenance
    - Do we need a more sophisticated naming scheme?
    - Does cmtconfig satisfy all the use cases?
- Restart talking about HPC and users



# Docs

- ADC Deployment twiki
- E-group: atlas-adc-containers-deployment
- Meetings: by-weekly after ADC weekly on Tuesday.
  - ADC Ops category: <https://indico.cern.ch/category/1706/>

