

# Singularity deployment

---

A. Forti, P. Nilsson, A. De Silva, A. De Salvo,  
P. Love, A. Filipcic  
ATLAS Jamboree  
7 March 2018



# Pilot 2 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



# Pilot 2 testing cont

- Stage-in/out in container delayed because Pilot 2 AGIS interface (info service component) not delivered until end of February
  - Pilot 2 data component now being updated to use info service
- Schedconfig field (container\_type) can be used to let pilot know if stage-in/out should be done inside container
  - Not necessary on standard grid site since middleware is available, but useful option if middleware is not locally available



# Pilot2 and ALRB

- 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
  - ALRB containerization
    - 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
      - export
 

```
ATLAS_LOCAL_ROOT_BASE=/cvmfs/atlas.cern.ch/repo/ATLASLocalRootBase;source $
{ATLAS_LOCAL_ROOT_BASE}/user/atlasLocalSetup.sh --quiet;source
$AtlasSetup/scripts/asetup.sh AtlasProduction,20.1.4.14,notest --platform
x86_64-slc6-gcc48-opt --makeflags="$MAKEFLAGS";Reco_tf.py
--maxEvents=1 --inputHITSFile
mc15_13TeV/HITS.06828093._000096.pool.root.1 --outputRDOFile
RDO_7933006f-49e2-4739-889b-099aed850610.root
```



# What happened to the wrappers

- There were 3 wrappers when we started
  - European, US and aCT
- In the future there will be one
  - which can call both pilot2 on most sites
  - or run the container for aCT sites
- European wrapper can already do both.
  - Was also successfully tested on most of the US sites with aCT/condor and harvester/condor.
- We don't have to worry about dealing with 3 different implementations in the future



# 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
    - .....
- For now image decided from the cmtconfig
  - Platform='x86\_64-slc6-gcc48-opt' ==> slc6 image
  - ALRB by default takes the unpacked but it can be overridden
  - This needs to be discussed more though



# 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 /pnfs/usatlas.bnl.gov -B /cvmfs -B /home/condor`
- Parameters can be left empty and standard payload will run



# 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 security vulnerability was announced by OSG this week for the old version (2.3.2)



# 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/>

