

Containers info mesh

Alessandra Forti

TIM @ CERN

21 September 2017



OS and Version

- To run containers in the pilots installing singularity on the WNs will be a requirement for sites moving to CentOS7.
- Latest singularity version 2.3.1 has become the baseline in WLCG.
 - <https://wlcg-mw-readiness.cern.ch/baseline/current/>
 - For now only available from the singularity site.
 - OSG and WLCG repositories are under discussion.
- Most of the features work only on CentOS7 family of OS
 - In particular the latest is the possibility to run singularity without setuid in the pilot which works only on **CentOS7.4**.
 - It will require some changes to the /etc/default/grub files that can be added at installation time.



Site or pilot

- Some sites are running containers as part of the WN infrastructure.
 - More normally Docker but also singularity
- Running containers in containers is possible but for now only if the external container is docker.
- Site singularity and pilot singularity are incompatible at least for now.



Images

- Images are distributed by ATLAS via CVMFS.
- ATM single file images
 - </cvmfs/atlas.cern.ch/repo/images/singularity/>
 - Too big too scale
 - Work ongoing to unpack them in cvmfs to access them as normal directories and let CVMFS work as usual.
- CVMFS must be binded inside the image in this model
- Sites without CVMFS (HPC) can use 'fat images' (See Wei's presentation)



Binding and setuid

- Some WNs directories have to be binded inside the image.
- Images can be created with a set of standard directories (/cvmfs, /scratch,...) but sites are varied not all dirs can be put in the image.
- Bind points can be passed to the command line as parameters. That's how currently pilot is running
 - Set of singularity parameters in AGIS
 - singularity_options="-B bind_path[:bind_target_path], bind_path[:bind_target_path] -contain"



Binding and setuid

- This requires allowing users to bind any directory in their images.
 - ATLAS(/CMS) run singularity in the infrastructure but this is general WN setup and other users can access singularity.
- Having only system bind paths though using images that don't contain the target requires setuid.
 - Not all sites want that (infact was major objection to glexec)
- Singularity can run in non setuid mode even allowing users to mount bind points.
 - Hence recommendation for CentOS7.4 and singularity 2.3.1 for this.



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/>

