Containers for users @CC-IN2P3

FJPPL Workshop

12/02/2019

Sébastien Gadrat

Outline

- Containers… means Singularity
  - As I will focus on computing
- Use cases
  - And why Singularity is well suited for several of our users' use cases
- Singularity installations
- Images storage and repository
- Building Singularity images
- Perspectives and conclusions

Containers @CC-IN2P3

- Docker can’t be used (on any of our computing clusters)
  - No Docker daemon available (required by Docker to run)

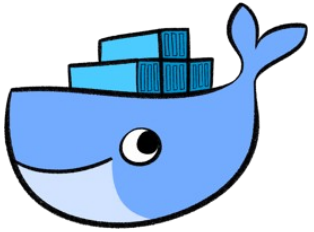

- Singularity
  - Initially provided to CMS, but available for any user
  - Easy to install and configure, no daemon required

Users' use cases

- Main use cases (~100 different experimental groups)
  - Old softwares, not compatible with CentOS7 (CERNlib for instance)
    - But that people are still using
  - Specific softwares which cannot be installed in CentOS7
    - Specific softwares running in Ubuntu only
  - On GPU: running TensorFlow (more recent) versions that we cannot install
    - Required more recent CUDA libs that the ones installed in the WNs
- 'Packed' images (.simg and .sif) are more often used
  - Easier to move the images
    - Especially when being used at several sites
  - Sandbox images only used by people only running at the CC-IN2P3
- Allow users to better manage their softwares' environment
  - Ensure reproductibility, especially when running at several sites
- Astro experiments seem quite interesting in the container technology
  - LISA is trying to build a computing model fully based on containers (Laser Interferometer Space Antenna)

Singularity installations

- Default's version currently available : 2.6.1
  - Default versions for WLCG
  - managed by Puppet
  - Running now in user namespace
- Newer versions installed in the software area
  - All 3.x versions available
  - All running in user namespace, except the latest installed one
  - Required version can be set up using the ccenv tool (home made tool to set up softwares)
- Singularity's configuration can be adjusted on users ' needs

Images storage and repositery

- Users
  - Users are storing their custom images in their group shared area (so called 'THRONG DIR') on the PBS filesystem
- CVMFS
  - Dedicated CVMFS area to store the images built by the CC-IN2P3's team
    - 200 GB repository
    - Take advantage of the cache on the Wns
    - Both packed and unpacked (sandbox) images

- GitLab Docker Registry
  - Since version 3.x Singularity is compliant with the OCI (Open Container Initiative), and images can now be stored in the GitLab Docker Registery
    - LISA is looking into that

Building Singularity images

- With Singularity v2
  - Build requires root privileges
  - Users build their images on their computer before moving them to the CC-IN2P3
    - that's why most of the images are 'packed' .sif images
  - GitLab-CI can be used to build Singularity images
    - But images have to be downloaded from GitLab, and then moved to their expected storage 'by hand'
- With Singularity v3
  - Build can be achieved in user namespace (with the 'fakeroot' option)
    - However requires a setuid Singularity installation
    -
- We provide help users to build their images

Conclusions and perspectives

- The goal is to provide the users with an user-friendly environment for running Singularity
- Further steps
  - Automatize all the steps : build → GitLab Docker Registry → CVMFS and/or PBS
  - Kubernetes can manage Singularity containers (since v3) : could it help to create and manage Singularity workflow ?

Thanks ! Any question ?