
Containers on the grid

— Andrej Filipčič —

Motivations to use containers

- Similar concept to virtual machines, but much more flexible and “zero” performance loss
 - chroot done properly
- Independence of execution environment vs host OS
- Can provide custom development, testing environment on grid
- Much easier and flexible to use
- Can make site maintenance and central image/software management much more transparent

Benefits for ATLAS and sites

- Sites can use any host OS of their preference
 - Minimal OS (eg CoreOS), bleeding edge (Fedora), latest enterprise OS
- Site OS major upgrades not affecting ATLAS production
 - OS upgrades can be done on the fly
- Many images can be used simultaneously on the same site, eg
 - SL6 for rel 21 production
 - SL7 for rel 22 testing and validation
 - SL5,4,3 for analysis of old data, data preservation
- Sites can only provide basic OS on the nodes
- Much more flexible and secure from site point of view
 - Isolation and traceability - replacing glexec functionality
- Common approach for execution, software distribution for all sites, including HPCs and ATLAS@Home

Technologies

Many, but the two most viable ones are

- Docker:
 - Suitable for general purpose (service) containers
 - Designed to for full privileges inside container
 - Execution of custom workloads in custom environment, eg preparing and checkpointing the development environment and subsequent large scale execution
 - Setup: docker daemon + approved host OS users (docker group)
- Singularity:
 - Designed for payload execution in batch systems (eg HPCs)
 - Users can execute a payload inside “any” OS image
 - No UID switching
 - root privileges disabled by design (capabilities dropped at launchtime)
- Both are complementary, docker images can be used by singularity

Containers strategy (supported by ICB)

- Enable containers at all ATLAS sites
 - Singularity -- the software to execute production and analysis containers
 - Simple to install, one package, minimal configuration
 - For sites with SL7 or modern OS, although it works with SL6 host OS as well
 - Some sites are executing all ATLAS jobs already in singularity containers
- Enable docker on few sites to support custom SW activities
 - More difficult to configure
 - root access inside the container - more expertise needed at sites
- ATLAS production system will adapt to use containers automatically
 - AGIS configuration, PanDA job image specification, pilot execution
 - Image distribution - cvmfs (directories, image files), or remote http
 - Relatively low effort

Technical details will be discussed today

Some sites already support containers

- RAL - docker with kubernetes, mesos, will try singularity
- MWT2 - singularity already used
- OSG/CMS sites - some sites in production with singularity
- SiGNET (gentoo) and ARNES (SLC6) execute all ATLAS jobs with singularity

More to test and deploy in the future

Further actions

- Discuss the strategy within WLCG
 - Workshop in Manchester
 - CMS already decided singularity is the way to go
 - Discuss with the sites
- Discuss the implications with other organizations (eg EGI, OSG)
- Discuss support, recommendations and security procedures
- Define isolation and traceability policies and requirements