Unprivileged singularity

Dave Dykstra

WLCG Containers WG Meeting

30 January 2019

Proposal for unprivileged singularity

- Goal: convert as much as possible of WLCG to unprivileged singularity
  - Preferably running out of cvmfs instead of rpm for easier upgrades
- Suggested plan:
  - Start with more testing on RHEL 7.6 with unprivileged namespaces
    - Fermilab & Nebraska are waiting on kernel fix for bug that breaks their use of docker containers – who else can volunteer in the meantime?
    - Using rpm with 'allow setuid = no' still installs suid binary, but very low risk
  - Next, convert VOs that are using singularity in production to try running singularity out of cvmfs at least if not found in $PATH
  - Announcement from security teams to enable unprivileged namespaces and encourage removing singularity rpm (preferably) or set 'allow setuid = no'
- GlideinWMS (CMS, OSG VO, FIFE) agreed strategy to find singularity:
  - If VO configures a SINGULARITY_BIN path, it is tried first
  - Next, singularity is searched for in $PATH
  - If those fail, runs from /cvmfs/oasis.opensciencegrid.org/mis/singularity/bin
  - Each attempt tries it with VO-supplied container and bind mount points