GDB Issues

John Gordon, STFC-RAL

MB meeting May 25th 2009

Science & Technology Facilities Council

May GDB Issues

- SL5 (separate presentation)
- gLExec
- Experiment Data Flows

gLExec (1)

- Now that SCAS has been deployed at a few sites, the testing of the combination gLExec/SCAS has moved forward. It has revealed:-
- gLExec is required to be installed on every Worker Node. Many sites make middleware available by a shared filesystem. They are reluctant to make an exception for gLExec.
- Comment by Antonio Retico
  - This is a serious issue (in my opinion the most serious). The deployment of glexec introduces a non-trivial constraint on the sites adopting the re-locatable solution (basically an intervention of 'root' user would be required for every WN installation, which is not the case now) and forces these sites to review their internal deployment procedures. This is a violation of the gLite WN feature “any non-privileged users can install a worker node” . Unfortunately according to the developers this behaviour cannot be worked around, because ultimately for the installation of a setuid program an intervention of root has to be required for security reasons. Even the extreme solution of re-compiling glexec locally (to adjust the linking of libraries) would be equivalent to the creation of a rootkit . So in my opinion the implementation doesn’t look like having been following mystical fumes.
  - Apparently we have to accept the implication
    glexec deployed => re-locatable WN distribution corrupted => some big sites need to adapt

gLExec (2)

- The process created under the new identity for a new payload does not inherit the environment
- While there are security exposures in inheriting too much, too little makes the user jump through hoops to recreate both the site and experiment parts of the environment.

Antonio's comments

- I remember that the issue with the environment was reported by Atlas since the early days of the pilot (already at the March GDB, slide 14 of my presentation), so not as new as suggested at GDB.
- However, during the following discussion with the developers and one of the sites two interesting points emerged:
- 1) the setuid programs like glexec should not inherit too much of the original environment for security reasons. In this respect that was a design choice of glexec suggested by common sense. Even had the developers made a different decision they couldn't however have included key environment variables (e.g. LD_LIBRARY_PATH) which setuid automatically discards for security reasons. Therefore a partial back-up/restore of the environment by the client would always have been needed
- 2) Pierre Girard observed that "the client" of glexec hasn't to be necessarily identified with the Experiment's frameworks. The necessary back-up and restore of the environment variables could be managed centrally by a wrapper distributed with the WN, which would be in turn called by the pilot frameworks. This would avoid the risk of independent solutions being implemented by the experiments, which would make debug analysis much more difficult.
- In conclusion there is actually something that can be done about the issue b) the question is: is the number of framework concerned worth the development of a central solution (and the consequent delays)? I think this is a good question for the TMB

Experiment Dataflows

- Tier1s have been asking for a long time for details of the experiment dataflows in order to estimate the required total tape bandwidths.
- At the May GDB, ATLAS, CMS, LHCb presented numbers
- Are they detailed enough for site planning?
- Negotiations on the site specific requirements are between the site and experiment
  - But the MB would welcome an overview. Are all sites happy/unhappy/relaxed?

A Common File Access Protocol

GDB

- Dirk Duellmann asked if there was support for xrootd as a common file access protocol.
- Support was lukewarm
- Direct filesystem access is likely to become pervasive in the next few years
  - Lustre etc
- Probably doesn't justify major developments for xrootd.

CREAM

- Improving each release
- Still not many sites testing.
- Reminded sites of MB Milestones

- **June 1st:**
  - All European T1 + TRIUMF and CERN; at least 1 CREAM CE each.
  - 5 T2s supporting ALICE with 1 CREAM CE each
- **August 1st:**
  - At least 2 T2s for each experiment providing 1 CREAM CE each
- **October 1st:**
  - 50 sites providing CREAM CEs in addition to the ones above

June Pre-GDB

- Storage with emphasis on T2s and Support
- Storm, DPM, dCache