



FIFE Meeting

Lisa Goodenough

17th January 2025

Agenda

- **FIFE News**
- **Experiment Production/Operations Updates**
- **Open Discussion/AOB**

News: jobsub

- Next release will be sometime in Spring
- This release will consist of a restructuring of the jobsub libraries so that users can utilize the jobsub libraries without calling *jobsub_submit*, *jobsub_q*, etc
 - So, if a submission script currently does a whole bunch of stuff in Python and then calls out to *jobsub_submit*, the script can be modified to just import the relevant jobsub libraries and run the submit directly
- Also, the new release will support a command syntax that will be like ‘jobsub submit <args>’, where the top level command is ‘jobsub’, while ‘submit’ and ‘q’ etc. are sub-commands

News: NAS Issue

- Wednesday January 15th (between 2 a.m. and 6 a.m. CT) the Information Technology Division performed network attached storage (NAS) maintenance that impacted various systems that mount central NFS storage.
- All EL9 systems had problems with kerberized mounts after the maintenance. User home directories on GPVMs, and the /web and /publicweb areas, were affected.
- Plan was to implement a workaround to restore access to the home directories, if the vendor was not able to solve the issue by Thursday morning.
- This indeed was the case, and the **workaround was put in place and access to home areas on gpvms was restored Thursday morning**. Production central websites were restored Thursday afternoon. /web and /publicweb remained unavailable (workaround not put in place for these).

News: NAS Issue

- Redhat (Linux vendor) and Hitachi (NAS vendor) continued to work together Thursday at high priority to come up with a permanent fix.
- As of Friday morning, Hitachi HNAS Engineering has a remedy for the problem. We will likely be able to implement it this morning. No NAS reboot is required; no Client reboot is required. The fix involves running a script that changes a configuration parameter. **This fixes the issue for kerberized NFS, so the /web and /publicweb areas will become available immediately.**
- The workaround was to revert to using un-kerberized NFS3, rather than kerberized NFS4. So, computing still needs to make the change back to NFS4. This will happen next week, and will require reboots of gpvms.

Memory Issues in HTCondor

- Since the transition to AL9 in the Spring, jobs run in HTCondor have been unable to get the same RSS (Resident Set Size) information they were able to get pre-transition
- The issue arises from the change in the version of cgroup from v1 (in SL7) to v2 in AL9. (cgroup is a mechanism to organize processes hierarchically and distribute system resources along the hierarchy in a controlled and configurable manner)
 - cgroup v2 reports memory usage by user jobs differently
 - v1 reported only RSS memory usage, which is why that metric was a good one to use for job management
 - v2 reports a total that includes RSS plus additional memory usage for caching, etc
- HTCondor uses whatever cgroup reports as it's measure of memory usage

Memory Issues in HTCondor

- Because cgroup v2 reports a higher memory value than cgroup v1, HTCondor records a higher value and uses this higher value to determine if a job is going over the requested memory
- A user assuming the same memory usage as before the transition to AL9 (cgroup v2), may find that their jobs are now being held for going over memory.
- Many users use the RSS memory usage that art reports at the end of the job to determine how much memory their job requires, and they may find that their jobs are now being held for going over memory.
- HTCondor developers are working on this issue. We will deploy it when it becomes available.
- If you see your jobs being held for going over memory, this is likely the reason.

News

- **Retirement of CILogon User Certificates**
 - **CILogon is set to retire X.509 Certificates at the end of May 2025 (<https://ca.cilogon.org/retirement>)**
 - **We need to understand where experiments (and services stand on this)**
 - **RUCIO (relevant for DUNE, Mu2e, ICARUS)**
 - **old versions in xrootd that are not token-enabled**

Experiment Production/Operations Update

- **Open Discussion**

Open Discussion and AOB