

International Resource Integration

Sakib Rahman (BNL) and Wouter Deconinck (UManitoba)
May 14, 2024

Alliance Canada - Cedar

Problem: Uses slurm scheduler under the hood. Once job is submitted from jlab submit node, it launches slurm jobs with 24 cores each on cedar (historical reasons). The slurm logs are only present while the jobs are running. Our existing workflow was wired to bind the container to /localscratch on the end node on slurm and didn't have proper write permissions.

Solution: In the end, the solution was to switch the binding to `_CONDOR_SCRATCH_DIR`.

fix: prioritize _CONDOR_SCRATCH_DIR over SLURM_TMPDIR #7

Merged wdconinc merged 1 commit into main from _CONDOR_SCRATCH_DIR 2 weeks ago

```
@@ -72,10 +72,10 @@ echo "SLURM_ARRAY_JOB_ID=${SLURM_ARRAY_JOB_ID:-}"
72 echo "SLURM_ARRAY_TASK_ID=${SLURM_ARRAY_TASK_ID:-}"
73 + echo "_CONDOR_SCRATCH_DIR=${_CONDOR_SCRATCH_DIR:-}"
74 echo "OSG_WN_TMP=${OSG_WN_TMP:-}"
75 - if [ -n "${SLURM_TMPDIR:-}" ] ; then
76   TMPDIR=${SLURM_TMPDIR}
77 - elif [ -n "${_CONDOR_SCRATCH_DIR:-}" ] ; then
78   TMPDIR=${_CONDOR_SCRATCH_DIR}
79 else
80   if [ -d "/scratch/slurm/${SLURM_JOB_ID:-}" ] ; then
81     TMPDIR="/scratch/slurm/${SLURM_JOB_ID:-}"
72 echo "SLURM_ARRAY_TASK_ID=${SLURM_ARRAY_TASK_ID:-}"
73 echo "_CONDOR_SCRATCH_DIR=${_CONDOR_SCRATCH_DIR:-}"
74 echo "OSG_WN_TMP=${OSG_WN_TMP:-}"
75 + if [ -n "${_CONDOR_SCRATCH_DIR:-}" ] ; then
76   TMPDIR=${_CONDOR_SCRATCH_DIR}
77 + elif [ -n "${SLURM_TMPDIR:-}" ] ; then
78   TMPDIR=${SLURM_TMPDIR}
79 else
80   if [ -d "/scratch/slurm/${SLURM_JOB_ID:-}" ] ; then
81     TMPDIR="/scratch/slurm/${SLURM_JOB_ID:-}"
```

Unresolved issues

How do the condor array jobs get distributed among the glide-ins running on slurm?

- a) Uniqueness of jobs
- b) Memory and core requirements

Running tests with larger sample size in production to understand and optimize.