
CERN

HTCondor Issues When Running SixTrack

A.Mereghetti

Many thanks to all who shared with me their issues with HTCondor, i.e. A. Alekou, R. De Maria, N. Karastathis, J. Molson, D. Pellegrini, R. Rossi, J. Wagner

NB: talk a bit re-scoped, since there were some updates in HTCondor lately…
hence, review of problems, lessons learnt and next steps – SixTrack-wise…

SixTrack

- Tracking code simulating single-particle dynamics in circular, ultra-relativistic machines (R. De Maria, this WG, meeting held 13th Oct 2016);
- DA:
  - Scan particle phase space (amplitude / angle in x-y) for different machine configurations (tune / chroma / octupoles / xing scheme) to capture the onset of chaotic motion;
  - Possibility to include statistical variations on multipolar errors and misalignments;
  - Very large use of CPU, depending on extension of phase space to be spanned, machine configurations, use of multipolar errors, max number of revolutions;
  - Management of input files and simulations via SixDesk;
  - 2k-50k jobs per single machine configuration, 60 p/job, 10^4-10^6 turns;
- Collimation:
  - Assess the performance of cleaning systems, combining single-particle beam dynamics and interactions with beam-intercepting devices;
  - Quite large use of CPU, especially for LHC (thanks to high cleaning efficiency);
  - Fluka-SixTrack coupling: variation on the theme, but same aim;
  - Management of simulations through very basic scripts (not SixDesk, yet);
  - 1k-2k jobs (4k-5k for coupling) per LM, 3.2k-6.4k p/job (1k-2k for coupling), 200 turns;

Typical running situation: many “short” jobs running in parallel!
→ Each 5m-20m for collimation/coupling, 40m-10h DA;

SixDesk

- Management of input files and runs for DA studies:
  - *mad6t.sh* (bash):
    - MADX jobs, to generate fort.2 (lattice and *k*-values), fort.8 (misalignments) and fort.16 (multipolar errors);
    - Up to 60 jobs in parallel: statistical variations on multipolar errors and misalignments;
    - Submission to LSF and HTCondor, or interactively (i.e. local lxplus node or own pc);
  - *run_six.sh* (bash):
    - actual SixTrack jobs;
    - 2k-50k jobs in parallel for a single machine configuration ;
    - Submission to LSF, HTCondor, BOINC (soon HTBoinc);
  - *sixdb* (python):
    - Storage of results and analysis;
- IT-related problems make every-day life of user painful, e.g.:
  - Shortage of namespace in user spooldir on work.boinc volume on AFS;
  - work.boinc volume hanging due to high I/O to disk – file editing, metadata, …
  - lxplus-releated: seg-faults when using sed (input file manipulation);
  - lxplus-releated: expiration of Kerberos/AFS token;
  - HTCondor (see later);

Issues with HTCondor – Tokens

Mainly related to AFS/Kerberos tokens lost / no longer valid at a certain point of the simulation (including AFS not mounted correctly on working node):

- Jobs lasting longer than expected, with no apparent progress, and then put on hold (R. De Maria, 27th Oct 2017);
- `SHADOW failed to receive file ..` (R. De Maria, 27th Oct 2017);
- Jobs mysteriously disappearing from `condor_q` without returning results (R. De Maria, 12th Oct 2017, D. Pellegrini, 26th Oct 2017);
- `cp: cannot stat `/afs/cern.ch/project/sixtrack/build/4630/SixTrack_4630_libarchive_bignblz_crlibm_fast_tilt_cmake_Linux_gfortran_static_x86_64_64bit': No such file or directory` (N. Karastathis, 18th Oct 2017);
- `cp: accessing `/afs/cern.ch/work/r/rrossi/private/SixTrack/scan_6p5_v/AM/ang_+7.5/run0001/': Permission denied` […] (R. Rossi, 7th Nov 2017, J. Wagner, 6th Nov 2017);
- `klist: No credentials cache found (ticket cache FILE:/tmp/krb5cc_102642)` (M. D'Andrea, 4th July 2017);

Issues with HTCondor – Tokens (II)

B. Jones – constantly there for us!

- If you're interested, the problem is this:
  1. On condor_submit you acquire an "ap_req" token, which is passed to the schedd
  2. The schedd uses the ap_req to acquire a Kerberos tgt
  3. The Kerberos tgt is used to acquire an afs token
  4. The tokens are used to write log files / out / err on the schedd
  5. The tokens are copied with the job to the execute node
  6. The job executes then passes back its results to the schedd.
- store_credd failed means 2. failed. The other error you've seen (silent failure) means 5. failed.
- There needs to be a fix for both of these. For 2 people understand the problem, for 5 people understand the symptom.
- My workaround is to send your Kerberos token along with the ap_req. This gets both fixes off the critical path, as renewal failures you won't see.

- Workaround implemented by B. Jones:
  - tested on lxplus-testing – submission automatically goes to 2k nodes with patched soft;
  - New software actually ported to all nodes – available since 8th Nov 2017;
  - Occurrence of problems seems to be lower… – hard to check with HL-LHC annual meeting on going…

Issues with HTCondor – Misc.

Problems on user's side, triggered by a change in default settings/values

- `RemoveReason = "Job removed by SYSTEM_PERIODIC_REMOVE due to Remove Reason unknown." DiskUsage = 37500000` (D. Kodjaandreev, 24th Oct 2017);
  - In collimation version we request: `request_disk = 50000000` - max given by htcondor: 20-25GB;
  - Never had an issue with jobs – cannot believe never passed 20-25GB…
  - No check at submission level → now introduced check during run time → jobs killed;
- `-- Failed to fetch ads from: <128.142.194.108:9618?addrs=128.142.194.108-9618+[2001-1458-301-e1--100-66]-9618&noUDP&sock=16152_81cb_17> : bigbird01.cern.ch - SECMAN:2007:Failed to end classad message.` as reply to a `condor_q` (A. Mereghetti, 15th July 2017)
  - Schedulers are not reachable from time to time → be patient!

Issues with HTCondor – Misc. (II)

Problems on user's side, triggered by a change in default settings/values

- Low number of jobs being processed by HTCondor – e.g. 100 instead of 5k (N. Karastathis, 12th Oct 2017);
  - problem with accounting group: `group_u_BE.ABP.SLAP` vs `group_u_BE.UNIX.u_pz` (lower priority);
  - Can be fixed with `+AccountingGroup = "group_u_BE.ABP.SLAP"` in .sub;
  - Then, issue with scheduler → `bigbird01` ran out of memory and was rebooted (no announcement);
- `Job removed by SYSTEM_PERIODIC_REMOVE due to wall time exceeded allowed max.` (N. Karastathis, 30th May 2017)
  - Problem with `+JobFlavour = "nextweek"` (use double quotes);

Lessons Learnt

1. Always use double quotes " when specifying strings as content of variables / parameters in `.sub` file;
2. Always add `$(ClusterId)` in `.sub` file when specifying `.out/.err/.log`, otherwise subsequent submission of the same job will overwrite files, and debugging by IT won't be possible;
3. Always add `$(ProcId)` in `.sub` file when specifying `.out/.err/.log`, to have a single `.out/.err/.log` file for each job, clearly identifiable;
4. How to change scheduler on the current terminal:
   - Temporary change: you have to remember the selected scheduler;
   - We all have a default scheduler – hence, I think that to manually change the scheduler is a last-resort ;

- Useful commands:
  - `condor_rm -forcex jobID` when `condor_rm` is not sufficient to remove clusters/jobs from `condor_q`;
  - `condor_ssh_to_job jobID` when want to connect to running job (node and tmp dir where job is being run);

Upcoming Changes in SixDesk

HTCondor gives the possibility of remote submission:

- you control everything from your own machine – must be in CERN net;
- No need to log-in to lxplus in order to submit!!!
- No need of a shared filesystem between submitting machine and actual processing node!!

  → …at least for inp / out / err / log and results, though necessary for storing the executable compiled with the proper libraries (and in the proper places);
- `-spool` at `condor_submit`: files in `transfer_input_files` (`.sub` file) are sent to the scheduler at `condor_submit`;
- Retrieval of out / err / log and results at `condor_transfer_data`:
  - when the user wants;
  - Drawback: user gets result not automatically → if many results have to be downloaded at the same time, this action can take long;
  - Key ingredients to be installed on your own pc: Kerberos + HTCondor (+openAFS) → guidelines (Ubuntu);

This requires to re-work the management of the jobs!

mad6t.sh

- submission command: use -spool option:

```
condor_submit -spool -batch-name "mad/$workspace/$LHCDescrip"
${sixtrack_input}/mad6t.sub
```

- mad6t.sub

```
universe = vanilla
filejob = %filejob%
executable = mad6t_$(seedID).sh
output = $(filejob).$(ClusterId).$(ProcId).out
transfer_input_files = $(filejob).$(seedID)
transfer_output_files =
$(filejob).out.$(seedID),fort.3.mad_$(seedID).gz,fort.3.aux_$(seedID).gz,fo
rt.2_$(seedID).gz,fort.8_$(seedID).gz,fort.16_$(seedID).gz
+JobFlavour = "microcentury"
queue seedID from jobs.list
```

Just a user variable (with query-replace by sed)

Similarly for .err / .log

List of seeds in file – flexible mechanism for re-submission of missing points

- Retrieval of results:

```
local __treatIDs=`condor_q ${__clusterID} -l -const 'JobStatus == 4' | grep
'^ProcId' | awk '{print ($NF)}'`
condor_transfer_data ${__clusterID}.${__treatID}
if [ $? -eq 0 ] ; then
   condor_rm ${__clusterID}.${__treatID}
fi
```

Get ID of completed jobs and download data one by one – long and painful…

run_six.sh

- All steps are identical to mad6t.sh;
- The only difference is in the .sub file:

```
universe = vanilla
executable = /afs/cern.ch/project/sixtrack/build/sixtrack
output = $(dirname)/htcondor.$(ClusterId).$(ProcId).out
transfer_input_files = $(dirname)/SixIn.zip
transfer_output_remaps = "fort.10=$(dirname)/fort.10"
ShouldTransferFiles = YES
WhenToTransferOutput = ON_EXIT_OR_EVICT
+JobFlavour = "tomorrow"
queue dirname from <fullPathToWork>/htcondorjobs/<study_name>.list
```

Similarly for .err / .log

Remap result file, to save it in correct sub-folder

List of dirs for input / output files in file – flexible mechanism for re-submission of missing points