

HTCondorCE Longer Term Test

Ste Jones
Liverpool
28th Feb 2019

sjones@hep.ph.liv.ac.uk

Status

- The HTCondor-CE system with APEL accounting has been operational here since late January (it might have been better if the meeting was tomorrow, because then I would have a full month under operations. But I'll try to analyse the process with what we have.)
- Adequate documentation is available to guide any other site who wishes to use the same scheme.
 - <https://twiki.cern.ch/twiki/bin/view/LCG/HtCondorCeAccounting>

Setup

- The cluster that we use behind the HTCondor-CE gateway is a standard HTCondor batch system running CentOS 7.6, with 752 slots.
- There is no provision for defragmenting for mcore jobs, yet the system seems to run a good proportion of ATLAS mcore jobs. It is a mystery how this is happening. But it is.
-

Usage

- Job mix looks “OK”:

```
# ./measureUse.sh
HEPGRID6 HTC-CE TOTAL SLOTS IN USE 752
HEPGRID6 HTC-CE TOTAL MCORE SLOTS IN USE 384
HEPGRID6 HTC-CE TOTAL SCORE SLOTS IN USE 368
```

- Only ATLAS and LHCb use it at present.

```
# condor_userprio -flat
Last Priority Update: 2/28 12:22
```

User Name	Effective Priority	Res Factor	Total Usage In Use (wgthed-hrs)	Time Since Last Usage
group_u_atlas.prdatl28@ph.liv.ac.uk	7197192.50	100000.00	193	23102.23 <now>
group_u_lhcb.pillhb25@ph.liv.ac.uk	9065720.00	100000.00	68	78703.15 <now>

```
-----
Number of users: 2                261  101805.38  0+23:59
```

Space used

- As mentioned, we write both batch and blahd logs. This has not made a big impact on the space; the system still has plenty of file space after a month with no cleanup.

```
# df
```

```
Filesystem      1K-blocks    Used Available Use% Mounted on
/dev/vda1       15349760 2823900 12525860  19% /
devtmpfs        1929420     0 1929420   0% /dev
tmpfs           1940172     12 1940160   1% /dev/shm
tmpfs           1940172 115148 1825024   6% /run
tmpfs           1940172     0 1940172   0% /sys/fs/cgroup
/dev/vda5       17600512 2489480 15111032  15% /var
/dev/vda2       15349760 127252 15222508   1% /opt
tmpfs           388036     0 388036   0% /run/user/0
```

Test

- Same test setup as before, but much more data
 - https://twiki.cern.ch/twiki/bin/view/LCG/HtCondorCeAccounting#Tests_on_a_HTCondor_CE



Results

- Abysmal. I was shocked...
- The “supposed” amount of HS06 Hours (work) sent from HTCondor-CE (hepgrid6.ph.liv.ac.uk) is 4.1 million (since there are 752 slots, then $752 * 24 * 10 * 27$ would give a max theoretical of 4.9 million, so occupancy is 84% - not bad.)

```
# for f in batch-*; do ~/scripts/workDone.pl < $f; done | perl  
-ape'$s+=$F[0]}'{$_=$s' ; echo  
- 4092824.7907333
```

- Yet the portal says only 368,009 HS06 Hours; a small fraction of the work done is accounted for. What has happened?

Explanation

- It's not just our HTCondor-CE.
- All our clusters (we have 2 x ARC and 2 x VAC clusters) are showing very low values; just a part of a month has got through to the portal.
- Suspect Portal is only updated after month end. So this test shows nothing.
- So, instead, count up only first four local batch log files for Feb. (a guess)

Explanation

```
# cat batch-20190201-hepgrid6 batch-20190202-hepgrid6  
batch-20190203-hepgrid6 batch-20190204-hepgrid6 |  
~/scripts/workDone.pl  
368098.361761132
```

- Now we get 368098 via HTCondor-CE APEL accounting, versus portal value of 368009. They are within 0.025%, which is tolerable.
- So; there's a delay in the Portal. Will wait until next week and try again for the full month.
-

Further work

- Will test with full month, but it looks good.
- We use a workaround to get the site benchmark into the system. But it's a "good" workaround, documented, almost as good as a fix.
- Some solutions out there; we may make it a config option, or put in the new JSON/GOCDDB BDII replacement Don't know yet. Not a show stopper.
- Also small issues about label fields; discussed with John Gordon – they are not important and largely unused.
- Need to see the release of the new APEL parser in UMD, so sites don't need to make special downloads.

End, questions, suggestions ...