

CMS Tier-1 Experiment sign off for Q2 2020



Katy Ellis, 30 Sept 2020

Running jobs at T1s


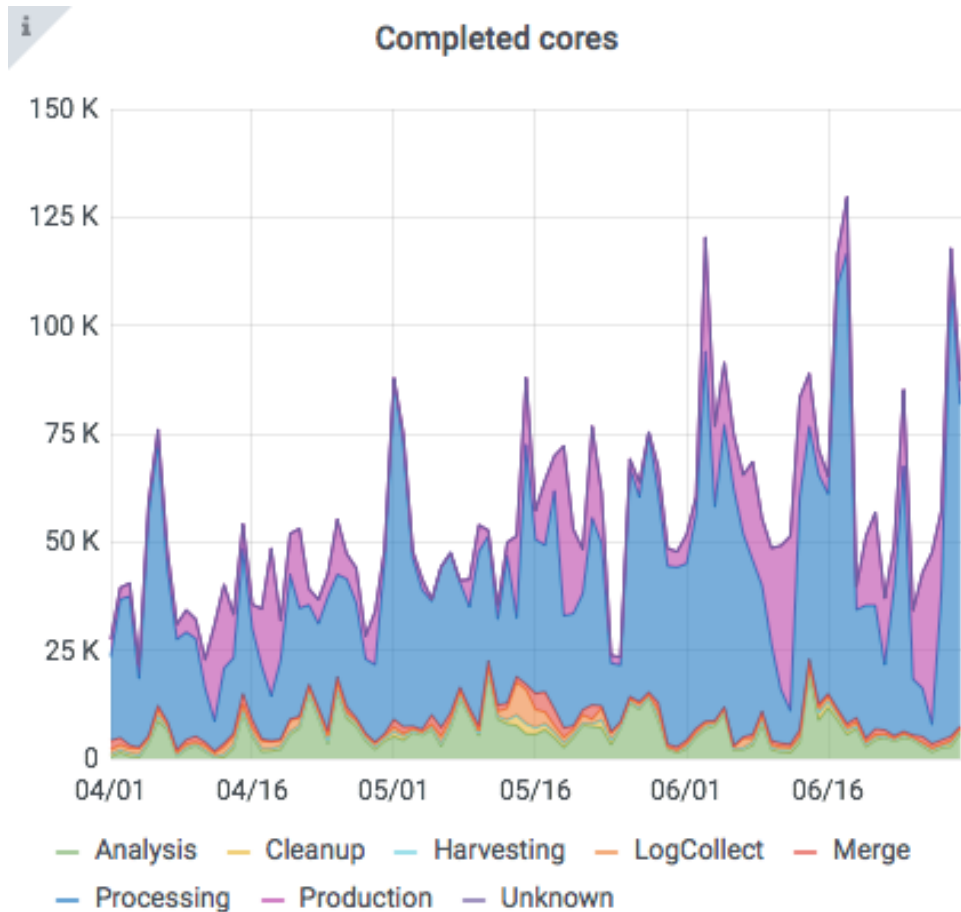


Total completed jobs

| | total | percentage |
|---------------|---------|------------|
| T1_US_FNAL | 5121852 | 32.3% |
| T1_RU_JINR | 3165267 | 20.0% |
| T1_IT_CNAF | 1776685 | 11.2% |
| T1_UK_RAL | 1773342 | 11.2% |
| T1_FR_CCIN2P3 | 1658431 | 10.5% |
| T1_DE_KIT | 1650697 | 10.4% |
| T1_ES_PIC | 709170 | 4.5% |



Running cores at RAL



Running more than pledged number of cores, except for a very brief period at the end of June.

Summary table of jobs

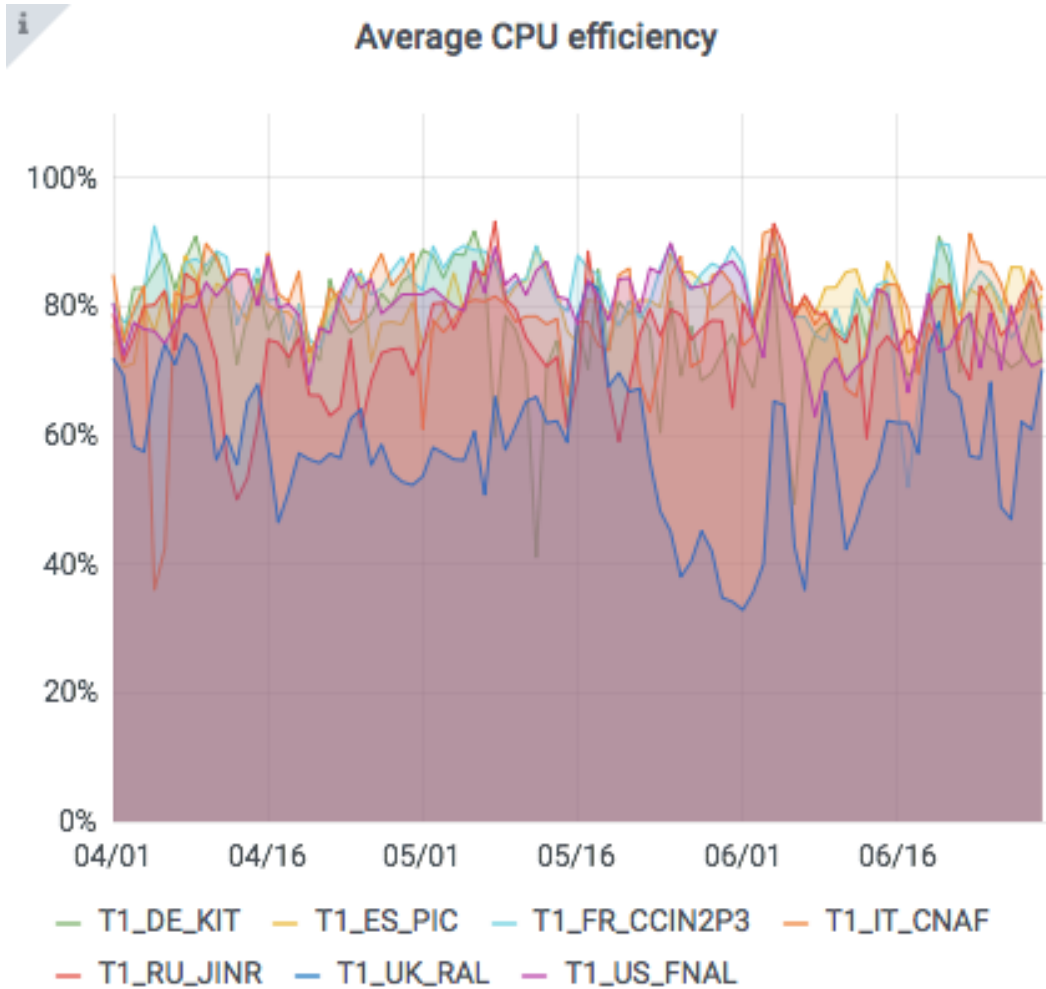
- <https://monit-grafana.cern.ch/d/C8ewaCrWk/hs06-report?from=1585695600000&orgId=11&to=1593557999000>

| Tier ^ | Site | Job Count | Failed jobs | CPU Eff | HS06CoreHr | CpuTimeHr | CoreHr | Avg Queue time |
|--------|---------------|-----------|-------------|---------|------------|-------------|-------------|----------------|
| T1 | T1_US_FNAL | 5114999 | 475294 | 79.3% | | 30835576.24 | 38871766.48 | 08:36:28 |
| T1 | T1_UK_RAL | 1761139 | 319978 | 57.6% | | 10573343.18 | 18343301.86 | 11:03:10 |
| T1 | T1_RU_JINR | 3109555 | 481180 | 75.2% | | 16172611.45 | 21518422.14 | 10:03:15 |
| T1 | T1_IT_CNAF | 1757001 | 222744 | 79.6% | | 10783073.90 | 13551748.00 | 12:33:43 |
| T1 | T1_FR_CCIN2P3 | 1635063 | 322492 | 82.1% | | 9960305.22 | 12136988.61 | 10:36:58 |
| T1 | T1_ES_PIC | 704882 | 125155 | 81.0% | | 3515196.53 | 4342241.37 | 09:48:33 |
| T1 | T1_DE_KIT | 1642443 | 355429 | 77.9% | | 12293309.52 | 15787711.20 | 12:42:41 |


Average CPU efficiency at 58% is well below other T1 sites.

I am convinced this value is reported incorrectly by jobs that are not 8-core.

CPU efficiency – including failed jobs

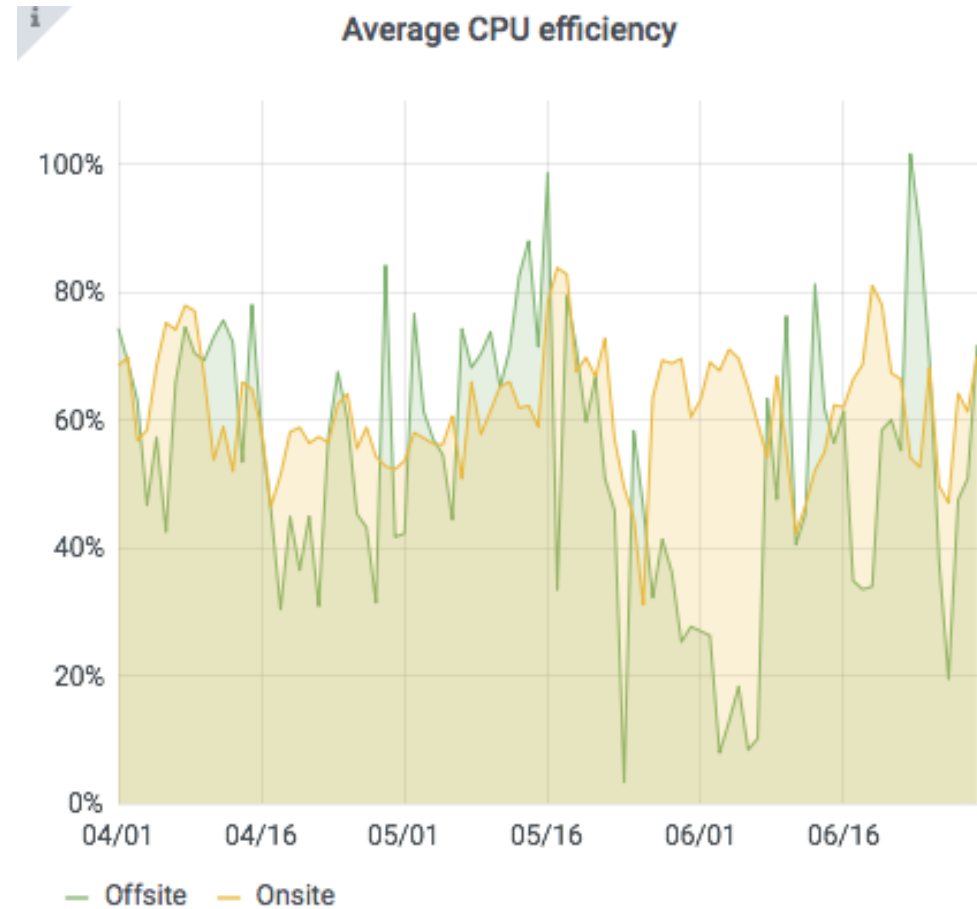



All Tier 1s




At RAL, split by job type

Offsite reads at RAL – a problem in June?




It looks like the large number of Offsite Primary reads coincides with a big dip in efficiency of those jobs. (FYI, offsite reads of 'secondary' datasets are happening even for jobs marked Onsite)

Failed jobs



Failure rate similar to other T1s during Q2



Large proportion of failing jobs at RAL are LogCollect, which still fail at a rate of ~100% due to known issue with xrootd version.

Disk usage

Numbers are taken from the webpage * on the last day of the month, disk used by Phedex as proportion of pledge for Phedex (not including small Rucio usage).

| End month | Echo usage (TB) | Allocation (TB) | RAL % | FNAL % | IN2P3 % | PIC % | CNAF % | KIT % | JINR % |
|-----------|-----------------|-----------------|-------|--------|---------|-------|--------|-------|--------|
| April | 4492 | 5440 | 83 | 76 | 81 | 73 | 64 | 80 | 75 |
| May | 4866 | 5440 | 89 | 74 | 81 | 96 | 66 | 77 | 72 |
| June | 4737 | 5440 | 87 | 74 | 75 | 86 | 58 | 79 | 70 |

* <http://cmsmonitoring.web.cern.ch/cmsmonitoring/storageoverview/2020/04/30/meeting.html>

Tape usage

| End month | Castor usage (TB) | Allocation (TB) | RAL % | FNAL % | IN2P3 % | PIC % | CNAF % | KIT % | JINR % |
|-----------|-------------------|-----------------|-------|--------|---------|-------|--------|-------|--------|
| April | 16075 | 17600 | 91 | 89 | 94 | 94 | 88 | 92 | 93 |
| May | 16205 | 17600 | 92 | 91 | 95 | 94 | 91 | 92 | 94 |
| June | 16251 | 17600 | 92 | 92 | 97 | 94 | 92 | 93 | 98 |

Summary

- CPU usage:
 - Number of cores in use is well over pledge.
 - Failure rate similar to other CMS Tier 1s.
 - Efficiency is lower, and believed to be at least partially related to offsite reads. In Q3 slow reads to the Worker Nodes are being investigated.
- Disk usage is high and being managed.
- Tape usage is also high, as are all Tier 1 tapes.