

CMS Tier-1 Experiment sign off for Q3 and Q4 2020

Katy Ellis, 27 Jan 2021


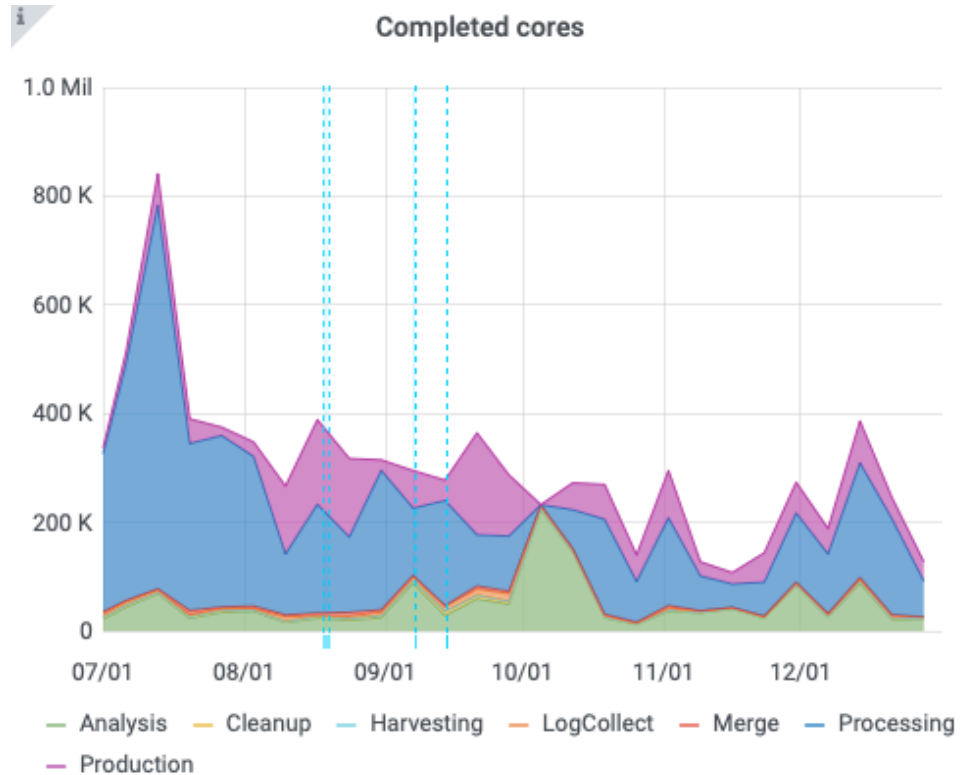
Completed jobs at T1s

Total completed jobs




	total ▾	percentage ▾
T1_US_FNAL	8574803	33.3%
T1_RU_JINR	5067013	19.7%
T1_IT_CNAF	3499839	13.6%
T1_UK_RAL	2765846	10.8%
T1_DE_KIT	2656634	10.3%
T1_FR_CCIN2P3	1973210	7.7%
T1_ES_PIC	1179167	4.6%

Completed and running cores at RAL



Number of completed cores appears to fall over Q3/4



Whole of 2020

Running more than pledged number of cores throughout, except for 4th – 12th October when RAL T1 was in drain for CMS due to failing SAM tests for the AAA service. This period coincided with when Katy was on holiday, so there was a delay in fixing the problems. Analysis jobs took up some of the slack.

Summary table of jobs, Q3

- <https://monit-grafana.cern.ch/d/C8ewaCrWk/hs06-report?orgId=11&from=1593558000000&to=1601506799000>

Site	Job Count	Failed jobs	CPU Eff	HS06CoreHr	CpuTimeHr	CoreHr
T1_US_FNAL	4298964	318243	78.6%		29139208.96	37077022.39
T1_UK_RAL	1408505	326238	46.0%		7364954.80	16002380.40
T1_RU_JINR	2550200	247442	74.4%		15584823.65	20951025.27
T1_IT_CNAF	1831999	218394	82.1%		17169298.14	20924392.13
T1_FR_CCIN2P3	1034812	153699	78.7%		9321127.05	11843862.36
T1_ES_PIC	616087	70921	74.0%		3362685.39	4546926.00
T1_DE_KIT	1094449	221587	73.7%		8583540.97	11639597.25

23% failure rate,
(18% in Q2)

Down from 58% last
quarter

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

Summary table of jobs, Q4

- <https://monit-grafana.cern.ch/d/C8ewaCrWk/hs06-report?orgId=11&from=1601506800000&to=1609459199000>


Site	Job Count	Failed jobs	CPU Eff	HS06CoreHr	CpuTimeHr	CoreHr
T1_US_FNAL	4236889	1056921	74.5%		26999662.83	36254979.99
T1_UK_RAL	1325081	384798	45.8%		5651029.29	12342931.18
T1_RU_JINR	2525518	213168	72.7%		16655443.88	22902439.40
T1_IT_CNAF	1659985	154825	68.6%		10573788.43	15417579.80
T1_FR_CCIN2P3	916922	110672	79.8%		6480047.11	8123770.26
T1_ES_PIC	566991	43751	74.5%		3204021.66	4298984.46
T1_DE_KIT	1568380	289177	74.4%		16808810.17	22605424.62

29%
failure
rate


Same as Q3

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

Failed jobs




All T1s have a spike here and there, but looks like RAL has a higher 'base rate'.




Spikes in RAL failures are mostly from Analysis jobs (green). LogCollect (orange) still fail at a rate of ~100% due to known issue with xrootd version.


CPU efficiency – including failed jobs




All Tier 1s




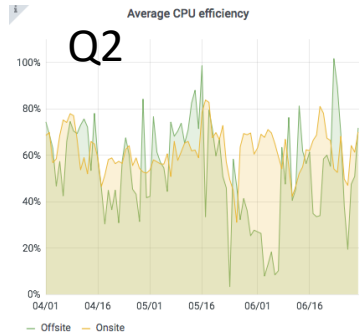
At RAL, split by job type



Input data and read time (all jobs)



Offsite reads at RAL – a problem for CMS



- The problem with the Onsite/Offsite flag here is that it only applies to the Primary dataset.
- Many jobs use significant secondaries.
- However we can still see here a much worse efficiency in jobs with the Primary input being offsite, than it was at the start of Q2.
- In July, I started measuring directly the data rate to the batch farm subnet from offsite, and found it to be $\sim 1\text{MB/s}$ for many sites.
- No improvements have been made since reporting this.

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 of quarter	Echo usage (TB)	Allocation (TB)	RAL %	FNAL %	IN2P3 %	PIC %	CNAF %	KIT %	JINR %
Q3	4498	5440	83	74	75	66	64	71	71
Q4	4386	5440	81	75	77	34	68	75	71

* <http://cmsmonitoring.web.cern.ch/cmsmonitoring/storageoverview/2020/09/30/meeting.html>

Q4 numbers are from 01/12/20 as this is the last date for which numbers are available. I will chase up why this is.

Tape usage

End of quarter	Castor usage (TB)	Allocation (TB)	RAL %	FNAL %	IN2P3 %	PIC %	CNAF %	KIT %	JINR %
Q3	15738	17600	89	98	101	89	98	90	96
Q4	15800	17600	90	98	101	91	98	90	97

Q4 numbers are from 01/12/20 as this is the last date for which numbers are available. I will chase up why this is.

I know that tape usage at RAL is now ~100%, and I asked for tape transfers to stop in December. They have not yet restarted.

Summary

- CPU usage:
 - Number of cores in use is well over pledge.
 - Failure rate has increased and becoming more concerning.
 - Efficiency is low, and believed to be at least partially related to offsite reads. No progress has been made on the network involved, which is shown to be extremely slow.
- Disk usage is high and being managed.
- Tape is full at RAL and we have not accepted new writes since Dec.