

CMS Tier-1 Experiment sign off for Q1 2022

Katy Ellis, 25 May 2022

Completed jobs at T1s

Total completed jobs ▾




	total ▾	percentage ▾
T1_US_FNAL	3483806	27.2%
T1_RU_JINR	2815744	22.0%
T1_DE_KIT	2206618	17.2%
T1_IT_CNAF	1205205	9.4%
T1_UK_RAL	1168596	9.1%
T1_FR_CCIN2P3	1005090	7.8%
T1_ES_PIC	540114	4.2%
T1_ES_PIC_BSC	397049	3.1%

Easily
meets
pledge

Failed jobs are included
in all metrics

Completed and running cores at RAL

In drain due to SAM test failures on Echo (+ HC failures on 4th Jan)



Last quarter, went into drain at least 6 times, which was clear from this plot. It LOOKS much better this quarter! However, we have been in an 'overridden' state since the Antares migration, so could not go into drain.

In drain...Katy was on holiday. AAA and HC were the joint cause.

Running cores was easily above pledge on average.

Summary table of jobs

Great improvements in Q1...

Site	Job Count	Failed jobs	CPU Eff	HS06CoreHr	CpuTimeHr	CoreHr
T1_US_FNAL	3451194	548064	71.9%		36893489.00	51276947.17
T1_UK_RAL	1166295	167675	66.3%		9855669.21	14861478.66
T1_RU_JINR	2805884	414989	75.1%		23657225.12	31486798.95
T1_IT_CNAF	1215186	71318	79.7%		11582828.18	14531591.24
T1_FR_CCIN2P3	1006034	83934	80.1%		10486947.69	13089714.34
T1_ES_PIC	542416	156652	79.7%		3590565.09	4504095.91
T1_DE_KIT	2217186	366745	72.5%		27754284.50	38295260.71

Test site, to be ignored

14% failure rate, see next slide

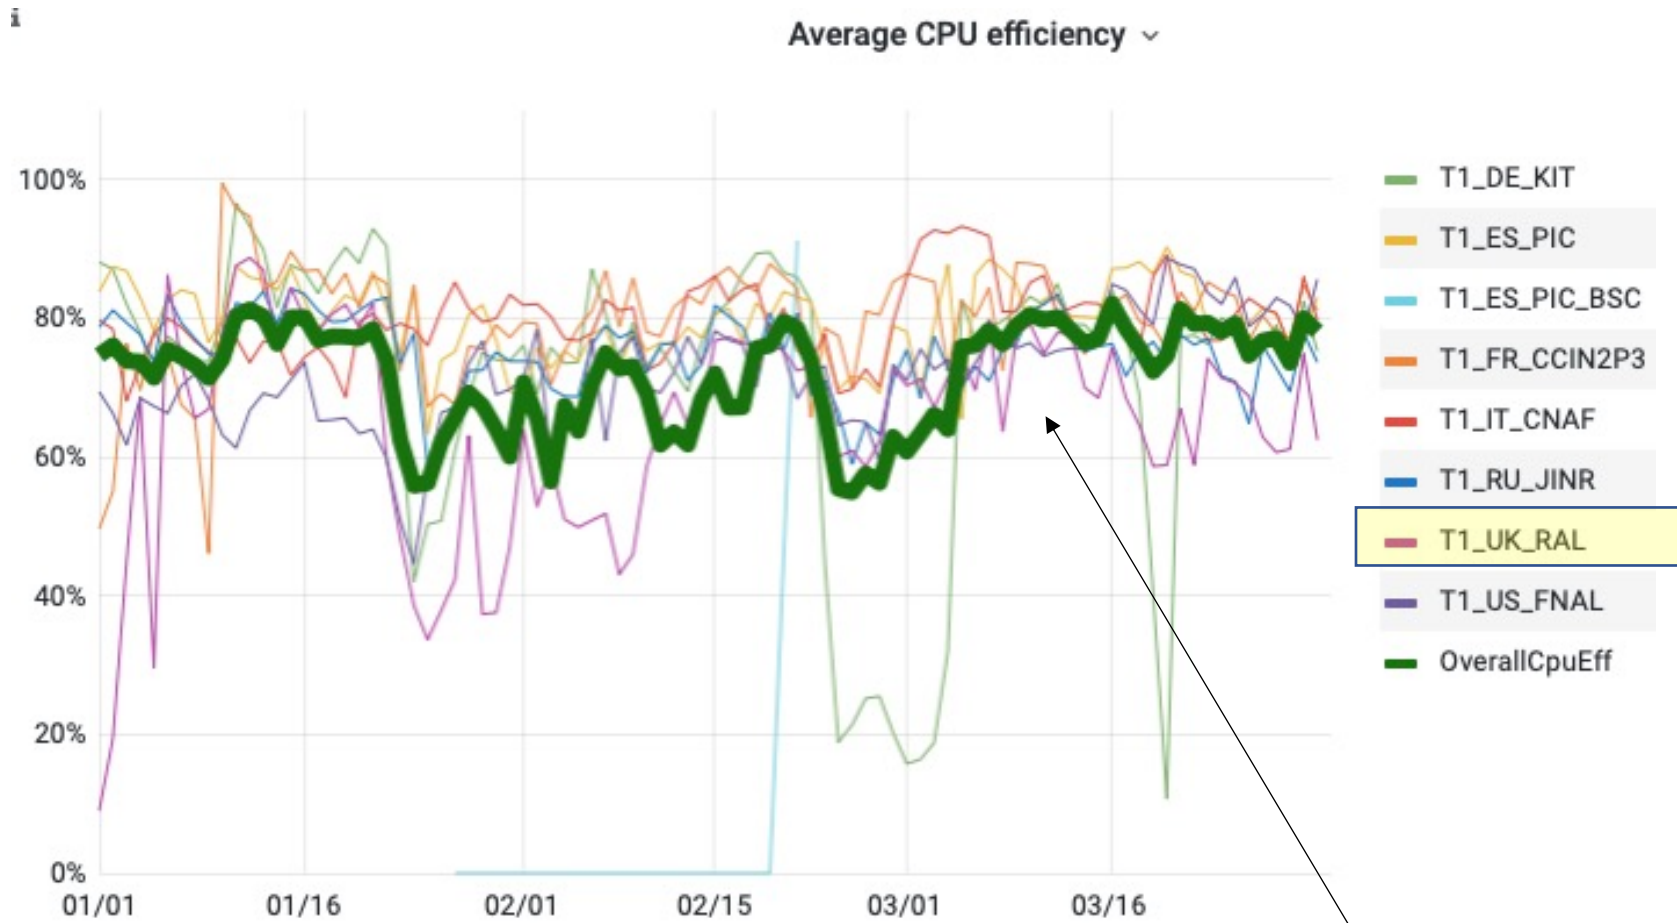
Another step up compared to previous quarters (50%, 57% in Q3/4)

I don't trust this number

Failed jobs analysis

- This quarter 14% of jobs failed – this is half of the previous quarter (large failure rate due to known reason), but still a little higher than the target.
- Jobs still failed at a higher rate in the first couple of weeks of the year due to the HammerCloud/CRAB problem – described in detail in the previous quarter sign-off.
- Also, 32% of failed jobs were of type LogCollect (which always fail at RAL..and still need to be fixed by CMS).
 - So really the ‘true’ failure rate is closer to 10%.

CPU efficiency

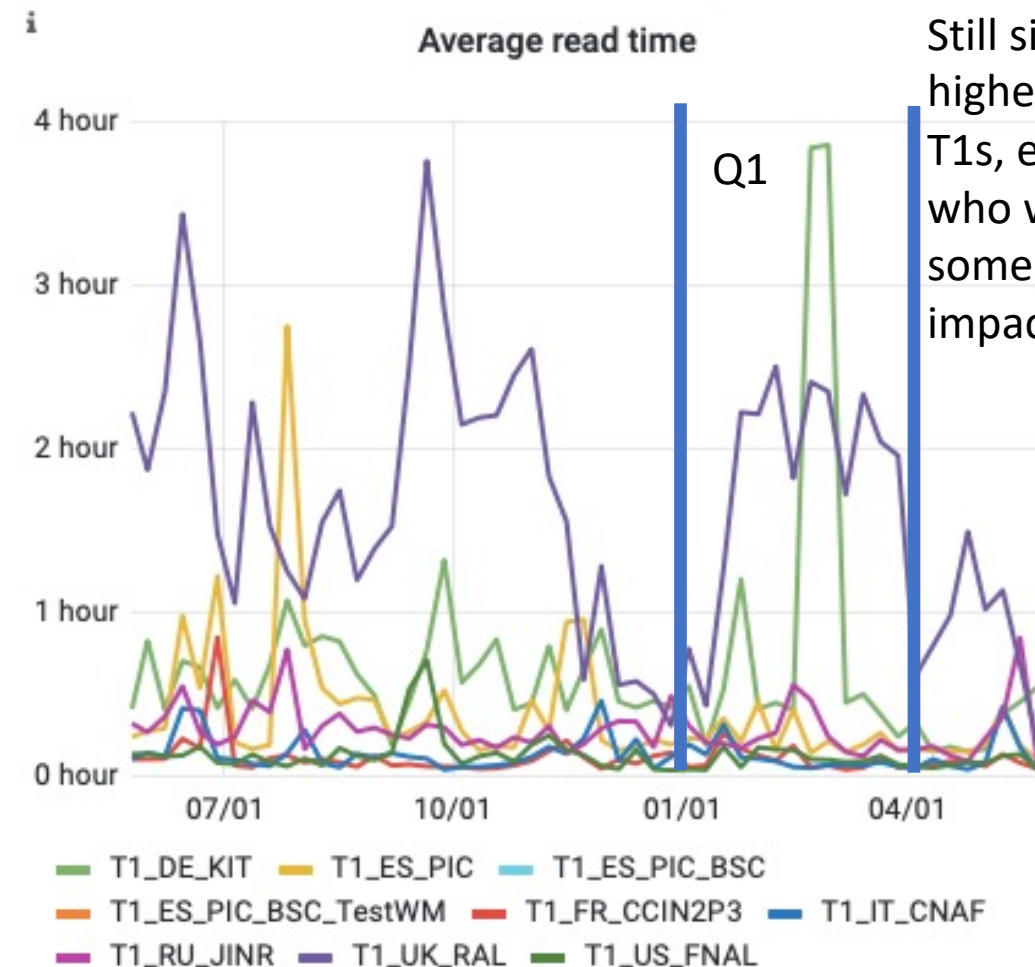
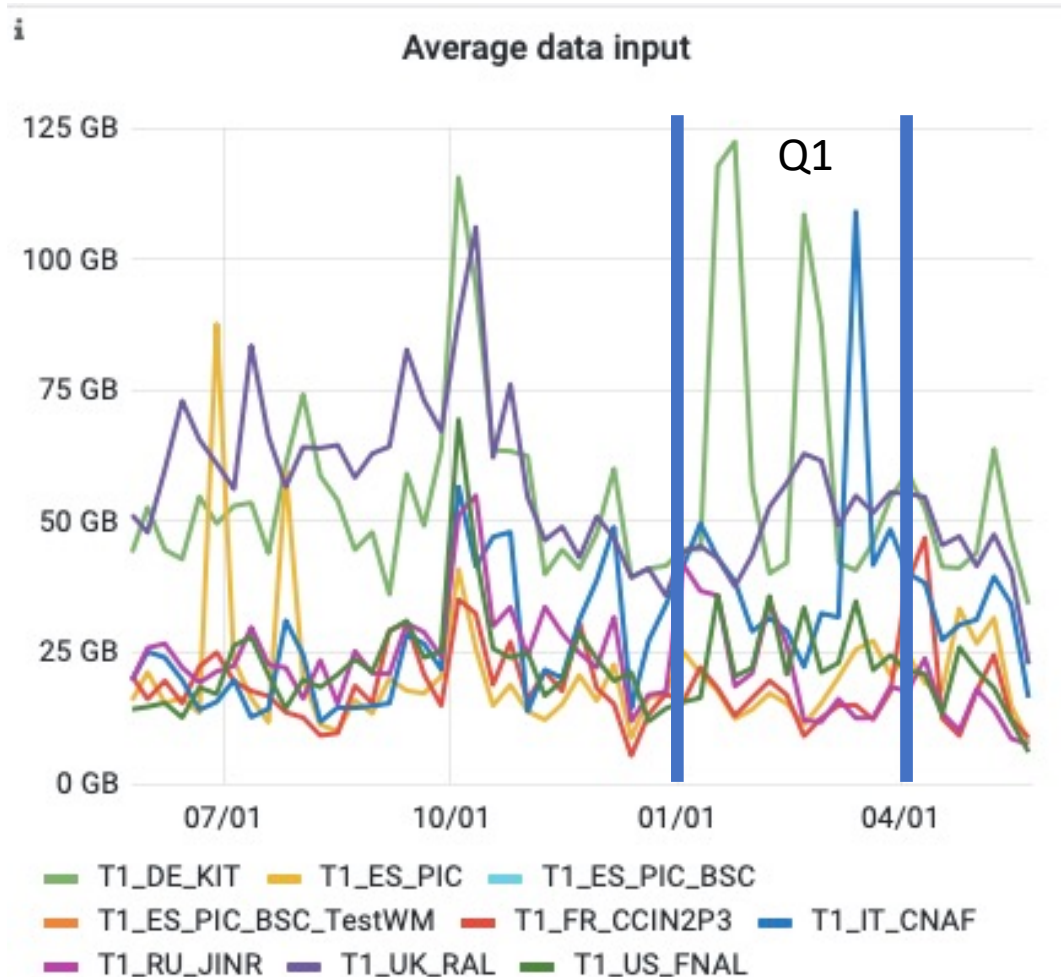


RAL average 66%, compared to ~75% for other T1s

Becoming a bit more stable, and close to the mean most of the time

Input data and read time (all jobs, Q1)

Could be reading data from onsite or offsite




Still significantly higher than all other T1s, excluding KIT who were doing some special high impact work

No changes made that should affect input quantity – still using L-D

Weekly binning, for the last 1 year

Disk usage (as of end of Q1)



↑ The usual managed usage at T1, with a balance of Locked and Dynamic

Webdav

- Tests failed regularly during the quarter
 - Would have taken us into drain a lot during Feb/March, but the status was overridden (impossible to go into drain)
- Generally, seems stable by time of writing (May)

Tape

- Data was migrated from Castor to Antares by 28th Feb 2022.
- Service started commissioning in the following week
 - A baptism of fire! (It was also the first week of the Tape Challenge for CMS)
- The tape challenge showed good performance for writes
 - Only problem observed in that week (14th March) was stalling of MGM
- Tape challenge reads were more problematic
 - This reflects CMS' experience with CERN CTA
 - Various fixes are already in place
 - Updates needed on RAL EOS/CTA version as well as CMS-Rucio
 - Then tests will be re-run

Summary

- CPU usage:
 - Number of cores in use is over pledge on average
 - Failure rate was high at the start of the quarter due to identified problem in CRAB jobs, only affecting RAL
 - Another big step forward in job efficiency - not yet as good as other sites but getting close!
- Disk usage is high and being managed
- Migration to Antares went well. Further improvements/tests to come
- Regular test failures with Webdav, but no big effect on production
- Not many times in drain, but would have been without override
 - As last Q, lots of changes happening, T1 team involved in a lot of activity.