

CMS Tier-1 Experiment sign off for Q4 2024

Katy Ellis, 29th January 2025

Overview

- Another good quarter for CMS at T1
- Echo and gateways have performed well
- Remote reads have performed well
- Issues with job submission similar to ATLAS
- Mini-DC for CMS T2s in December

Completed jobs at Tier 1

Total completed jobs ⓘ



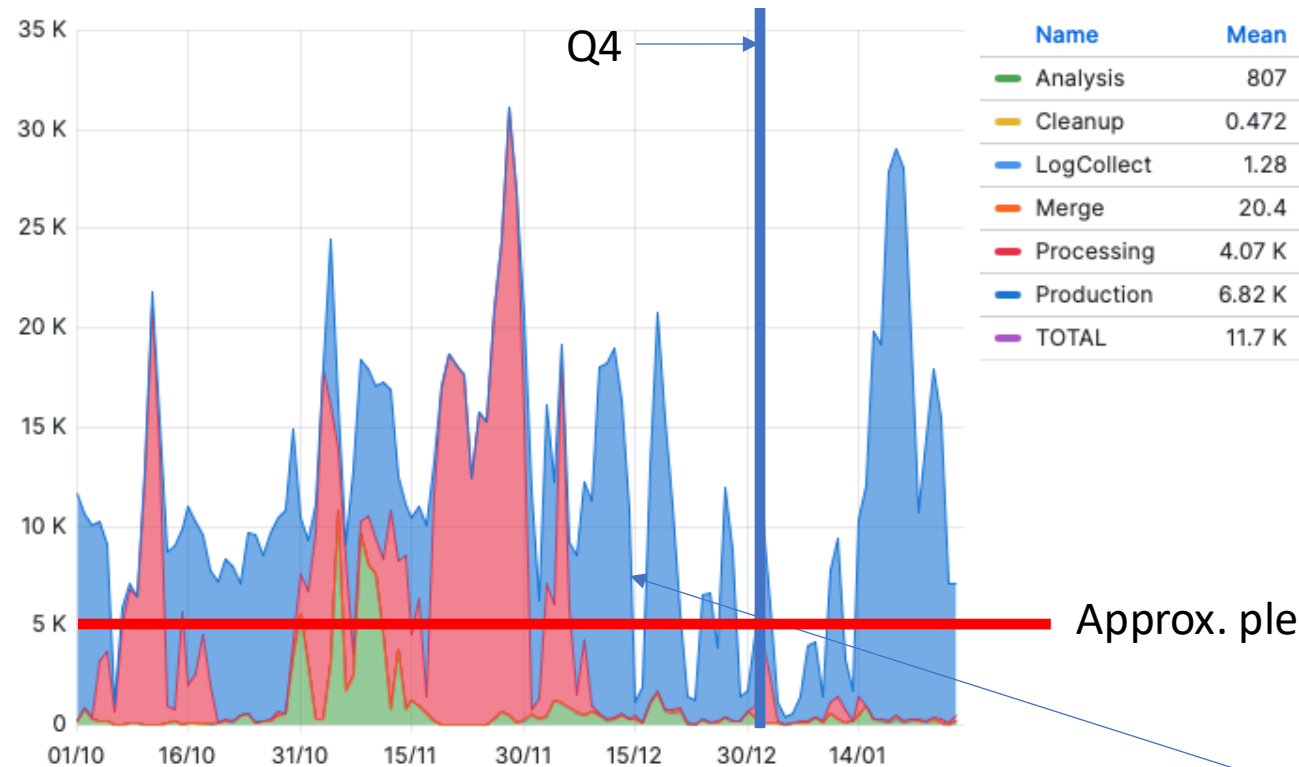
	total ▾	percentage ▾
T1_US_FNAL	18582983	58.5%
T1_RU_JINR	3351971	10.5%
T1_FR_CCIN2P3	2500831	7.9%
T1_DE_KIT	2347829	7.4%
T1_UK_RAL	2149438	6.8%
T1_IT_CNAF	1876600	5.9%
T1_ES_PIC	973724	3.1%



Failed jobs are included
in all metrics

Running cores at RAL

Running cores ⓘ



Still plenty of data
Processing jobs, but MC
Production starting to
dominate again


Approx. pledge

Starting to see periods of
variable number of cores
running even though not
'in drain'

Running cores was easily above pledge for the quarter.

SAM tests on ARC-CEs

- Tests for one CE (arc-ce01) but they all show similar behaviour




- Occasional errors were expected due to a problem in ARC
- Wasn't considered a problem as the tests on all 5 CEs never failed simultaneously
- Until the Christmas break...

Network issue prevented job submission?

- The SAM tests failing on previous slide were 'JobSubmit'
- This was coincident with variable numbers of jobs being run
- Also connected this with low numbers of ATLAS jobs running
 - Even when there were free cores on the batch farm
 - Even when CMS and ATLAS had work that could easily be matched to RAL
- Tom Birkett chased DI over a possible network/firewall change they made in November
 - The fix they implemented on 14th Jan improved things immediately
 - As well as numerous other 'slowness' observations across Tier 1 services

Job efficiency

Average CPU efficiency ⓘ



Despite some dips (experienced by all T1s) and issues with job submission, RAL remains the highest efficiency CMS T1 on average for the quarter

Name	Mean	Max	Min
T1_DE_KIT	73.6%	89.9%	16.3%
T1_ES_PIC	69.9%	89.8%	17.5%
T1_FR_CCIN2P3	74.6%	92.3%	23.1%
T1_IT_CNAF	72.8%	85.6%	24.5%
T1_RU_JINR	71.1%	96.7%	11.2%
T1_UK_RAL	76.3%	93.4%	13.6%
T1_US_FNAL	72.8%	84.9%	47.3%
OverallCpuEff	73.3%	84.6%	42.2%

Summary table of jobs

Site	Job Count	Failed jobs	CPU Eff	CpuTimeHr	CoreHr	Avg Queue time	Failure rate
T1_US_FNAL	18641989	4680850	72.4%	84107191.23	116209218.50	05:34:00	25.11%
T1_UK_RAL	2147346	245408	76.9%	20038989.76	26061683.52	08:57:11	11.43%
T1_RU_JINR	3323302	535325	72.1%	22540357.61	31275368.65	06:59:49	16.11%
T1_IT_CNAF	1880753	273753	73.4%	9814345.86	13377569.11	07:04:11	14.56%
T1_FR_CCIN2P3	2476379	395150	75.0%	10363248.96	13822073.10	06:33:22	15.96%
T1_ES_PIC	972544	380535	69.2%	3350889.70	4845643.61	06:10:08	39.13%
T1_DE_KIT	2343225	369886	72.2%	16466493.43	22811849.34	08:07:12	15.79%

Best of the T1s with all sites
a bit down on previous
quarters (RAL 69/83% in
Q2/3)

11% failure rate,
not as good as Q3
but still best of the
T1s

AAA reads from offsite

- In December we noticed an increase in the traffic on AAA
 - Still, this was only in the range of serving 100-200MB/s rate per gateway
 - But it was causing SAM test failures and complaints from other sites
- The reason for increase in traffic was premix library files on Echo
 - Policy states these files should not be sited on RAL disk...
 - They were there because of multihopping through Echo to Antares
- The AAA gateways were quite severely throttled in the past, to protect Echo
 - No longer necessary; Jyothish released the throttling
 - Then we saw some OOM errors at high load – still being investigated

Token status

- CMS continuing to use tokens at sites that support them in production/FTS data transfers since September
 - Number of transfers using tokens is low
 - CERN recently disabled due to deletion failures
 - No token transfers to tape (not yet supported with tokens in FTS)
 - Token only used if *all possible sources* and destination are token-enabled
- Tokens for tape – CMS still asking for green SAM tests but complicated situation with CERN (CTA)
- Long saga of token tests on ARC-CEs may be coming to an end soon..?
 - ARC7 nearly ready for deployment

Tape

- Using the pledge at RAL
- Not much write traffic in recent months
 - Hit 4GB/s in October
- Not much read traffic
 - Hit 2GB/s in December
- I hear plans for a deletion campaign, but no dates confirmed for the start yet

Mini-DC tests for Tier 2s

- [Test notes](#)
- [Presentation](#) at DOMA-General
- No problems with Tier 1 as source and sink (there was no intention to stress T1)
- Tier 1 tests planned for week of 3rd March

Summary

- Good performance for CMS at RAL this quarter, best of T1s again
 - Some issues with job submission at the end of the year caused by a network transceiver fixed in Jan
- Possible improvements for AAA
- Tokens still causing headaches!
- Mini-DCs have started