


CMS Tier-1 Experiment sign off for Q4 2022

Katy Ellis, 22nd Feb 2023

Overview

- Good performance
- Isolated incidents caused some interruptions
- Quite similar to Q3

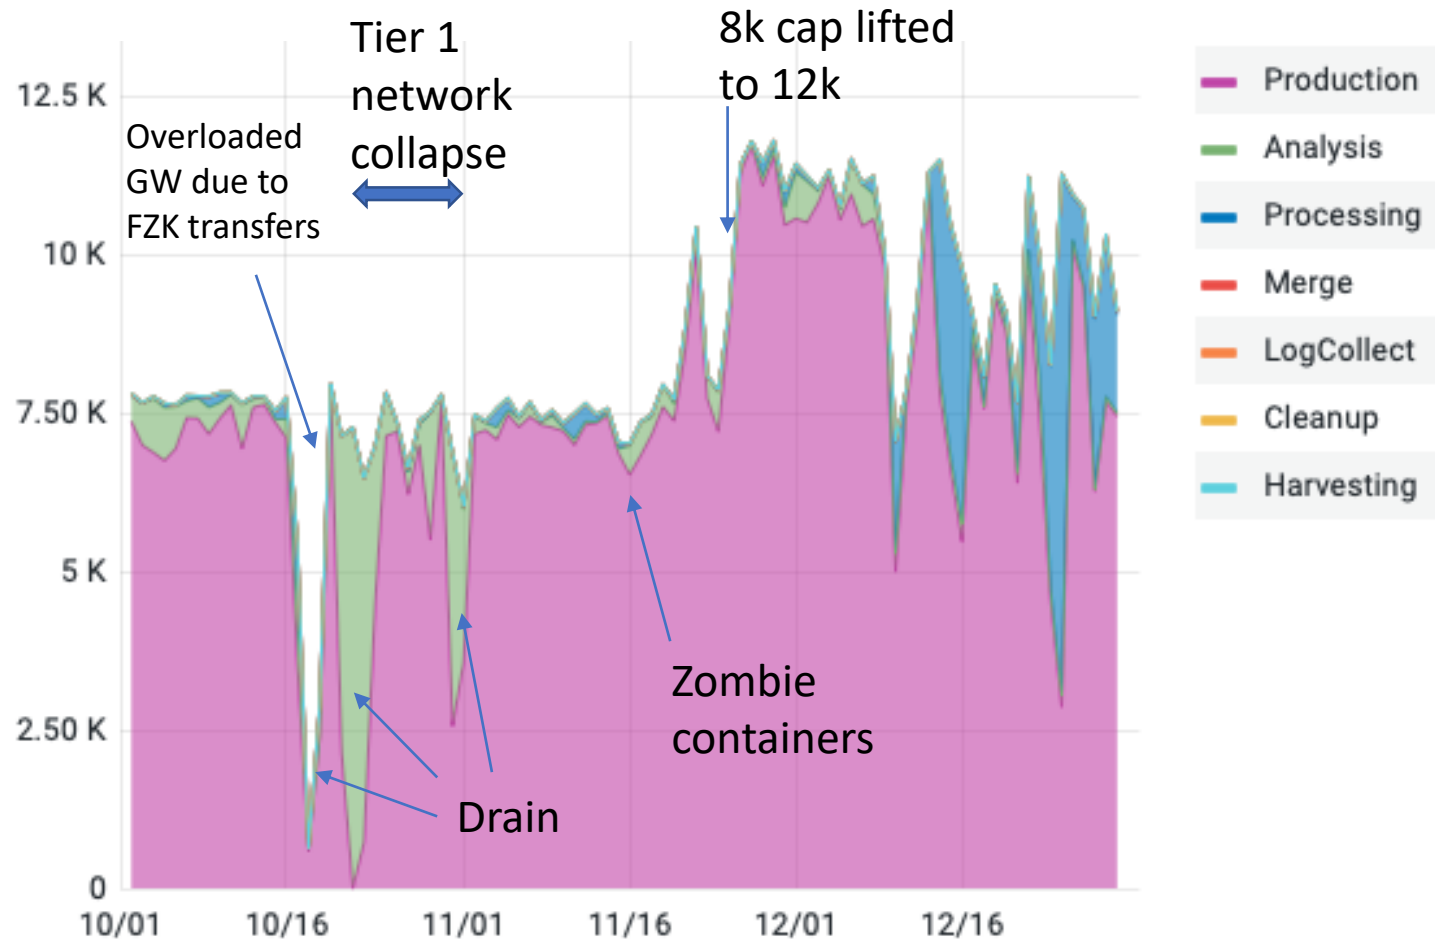
Completed jobs at Tier 1



| | total ▾ | percentage ▾ |
|---------------|---------|--------------|
| T1_US_FNAL | 5984574 | 38.3% |
| T1_RU_JINR | 2966445 | 19.0% |
| T1_DE_KIT | 1791796 | 11.5% |
| T1_UK_RAL | 1619812 | 10.4% |
| T1_IT_CNAF | 1611138 | 10.3% |
| T1_FR_CCIN2P3 | 1154709 | 7.4% |
| T1_ES_PIC | 507261 | 3.2% |

Failed jobs are included
in all metrics

Running cores at RAL



A few periods of production job drain in second half of October due to failing webdav tests (T1 legacy network collapse etc.)

The farm coped well going to a higher number of CMS jobs, and performance is generally good.

Job mix still dominated by MC Production, but significant Processing is starting.

'Zombie containers' issue had Dockers continue to run on RAL WNs and show up in RAL monitoring but not VO monitoring. Small effect for CMS but larger for other VOs

Running cores was easily above pledge throughout the quarter.

SAM tests

Nov

Webdav failures match up with periods of drain
– much better since the start of November

T1_UK_RAL webdav.echo.stfc.ac.uk org.cms.SE-WebDAV-1connection:



T1_UK_RAL webdav.echo.stfc.ac.uk org.cms.SE-WebDAV-2ssl:



T1_UK_RAL webdav.echo.stfc.ac.uk org.cms.SE-WebDAV-3extension:



T1_UK_RAL webdav.echo.stfc.ac.uk org.cms.SE-WebDAV-4crt-read:



T1_UK_RAL webdav.echo.stfc.ac.uk org.cms.SE-WebDAV-5open-access:



T1_UK_RAL webdav.echo.stfc.ac.uk org.cms.SE-WebDAV-6crt-write:



T1_UK_RAL webdav.echo.stfc.ac.uk org.cms.SE-WebDAV-7macaroon:



T1_UK_RAL webdav.echo.stfc.ac.uk org.cms.SE-WebDAV-9summary:



T1_UK_RAL arc-ce-test01.gridpp.rl.ac.uk org.cms.WN-analysis-/cms/Role=lcgadmin:



T1_UK_RAL arc-ce-test02.gridpp.rl.ac.uk org.cms.WN-analysis-/cms/Role=lcgadmin:



T1_UK_RAL arc-ce01.gridpp.rl.ac.uk org.cms.WN-analysis-/cms/Role=lcgadmin:



T1_UK_RAL arc-ce02.gridpp.rl.ac.uk org.cms.WN-analysis-/cms/Role=lcgadmin:



T1_UK_RAL arc-ce03.gridpp.rl.ac.uk org.cms.WN-analysis-/cms/Role=lcgadmin:



T1_UK_RAL arc-ce04.gridpp.rl.ac.uk org.cms.WN-analysis-/cms/Role=lcgadmin:



T1_UK_RAL arc-ce05.gridpp.rl.ac.uk org.cms.WN-analysis-/cms/Role=lcgadmin:



In previous quarters I reported on CE test problems when other VOs start draining out of the farm – again this was still apparent in Oct but definite improvement in Nov

Summary table of jobs

Efficiency a little below other T1s in Q4, likely due to low-efficiency analysis jobs when in drain.

| Tier ^ | Site | Job Count | Failed jobs | CPU Eff | | CpuTimeHr | CoreHr | Avg Queue time |
|--------|---------------|-----------|-------------|---------|--|-------------|-------------|----------------|
| T1 | T1_US_FNAL | 5886200 | 369655 | 80.0% | | 50966194.26 | 63735043.22 | 07:49:53 |
| T1 | T1_UK_RAL | 1630327 | 258660 | 70.7% | | 13741797.20 | 19433534.14 | 06:06:44 |
| T1 | T1_RU_JINR | 2998034 | 179688 | 80.1% | | 33628769.19 | 41981888.95 | 08:13:55 |
| T1 | T1_IT_CNAF | 1609954 | 129583 | 81.2% | | 21802568.06 | 26845727.11 | 12:08:23 |
| T1 | T1_FR_CCIN2P3 | 1153695 | 88064 | 79.9% | | 11358175.17 | 14215662.92 | 07:10:02 |
| T1 | T1_ES_PIC | 506790 | 108859 | 76.9% | | 4114622.35 | 5347145.84 | 05:23:39 |
| T1 | T1_DE_KIT | 1778365 | 98270 | 74.7% | | 18128797.43 | 24253695.36 | 06:01:34 |

16% failure rate, worse than last quarter


Similar to previous quarters (74%, 72% in Q2/3)

I don't trust this number

Failed jobs analysis


Spikes in prod job failures
(not just at RAL)

- This quarter 16% of jobs failed. If you discount LogCollect, it's just under 10%
- News on LogCollect:
 - Since December, not all LogCollect jobs are failing.
 - Why?

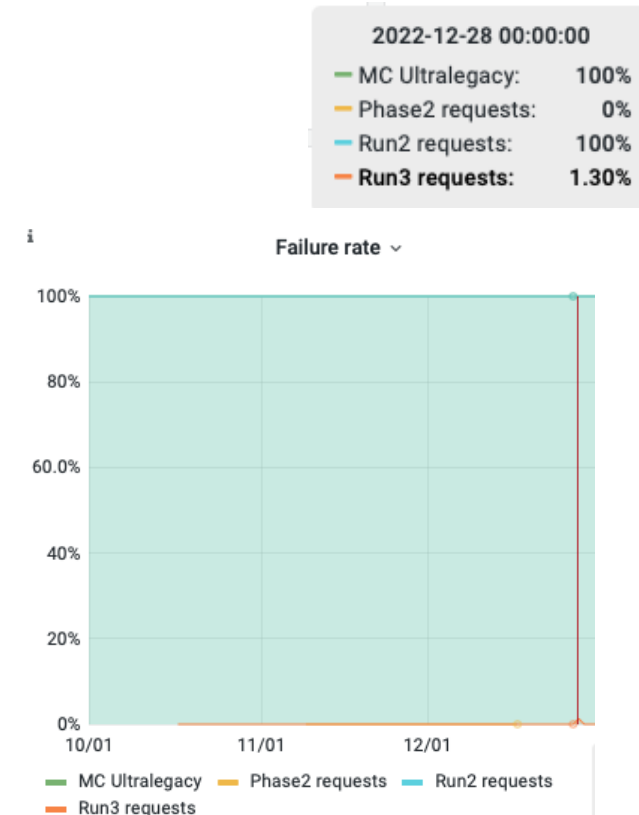


LogCollect jobs, Split by Campaign Type


Removed jobs of CampaignType 'UNKNOWN'



New campaigns using EL8?




CPU efficiency



As last quarter, high numbers of Analysis jobs are sometimes picked up when going into drain for prod

Disk usage (as of end of Q4)



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

RSE

No changes in disk usage

Consistency checking

- CMS are performing weekly consistency checking on RAL disk and tape
 - Only 'enforcing' consistency on disk at the moment
- For disk
 - 'Missing' data is being handled well by the end of Q4
 - 'Dark' data is still causing confusion – further investigation required
- For tape
 - Zero missing files; 604 dark data; situation seems good

Tape and mini-tape challenge

- Tape allocation increased above pledge to 20.6PB in July (end of Q4, "usedsize": 17729686000000000) to help with lost resources in Russia
- Known issues with tape recalls (CTA/Rucio/FTS) largely fixed – proven by additional recall test in October
 - Started recall of 114TB/9k files from 4 different tape pools on 13th Oct
 - Much of the same data used in the March challenge
 - ATLAS also added traffic
 - Antares -> Echo -> IC

Mini-tape challenge

- Goals:

- Sustained throughput to at least the RAL disk endpoint of 1.5GB/s as required by the March challenge document
- 100% transfer of the container
- Low failure rate
- Absence of the particular issues identified in the past on RAL and CERN CTA endpoints

- Side-effects:

- Found several issues in Rucio/monitoring to follow up

Summary

- CPU:
 - Number of cores in use is above pledge
 - Failure rate a bit high – due to CMS rather than RAL. LogCollect improving
 - Job efficiency is on the lower side of T1s – average would be better if we didn't go into drain, this would gain us a few percent improvement
- Disk usage is high and being managed as usual; consistency checking in progress
- Tape allocation increased; mini-challenge showed good performance
- Still instabilities with Webdav causing test failures
 - Not obviously a problem for production except when it pushes us into drain
 - Further hardware required