


LHCb status

GRIDPP 50


ALEXANDER ROGOVSKIY



Contents

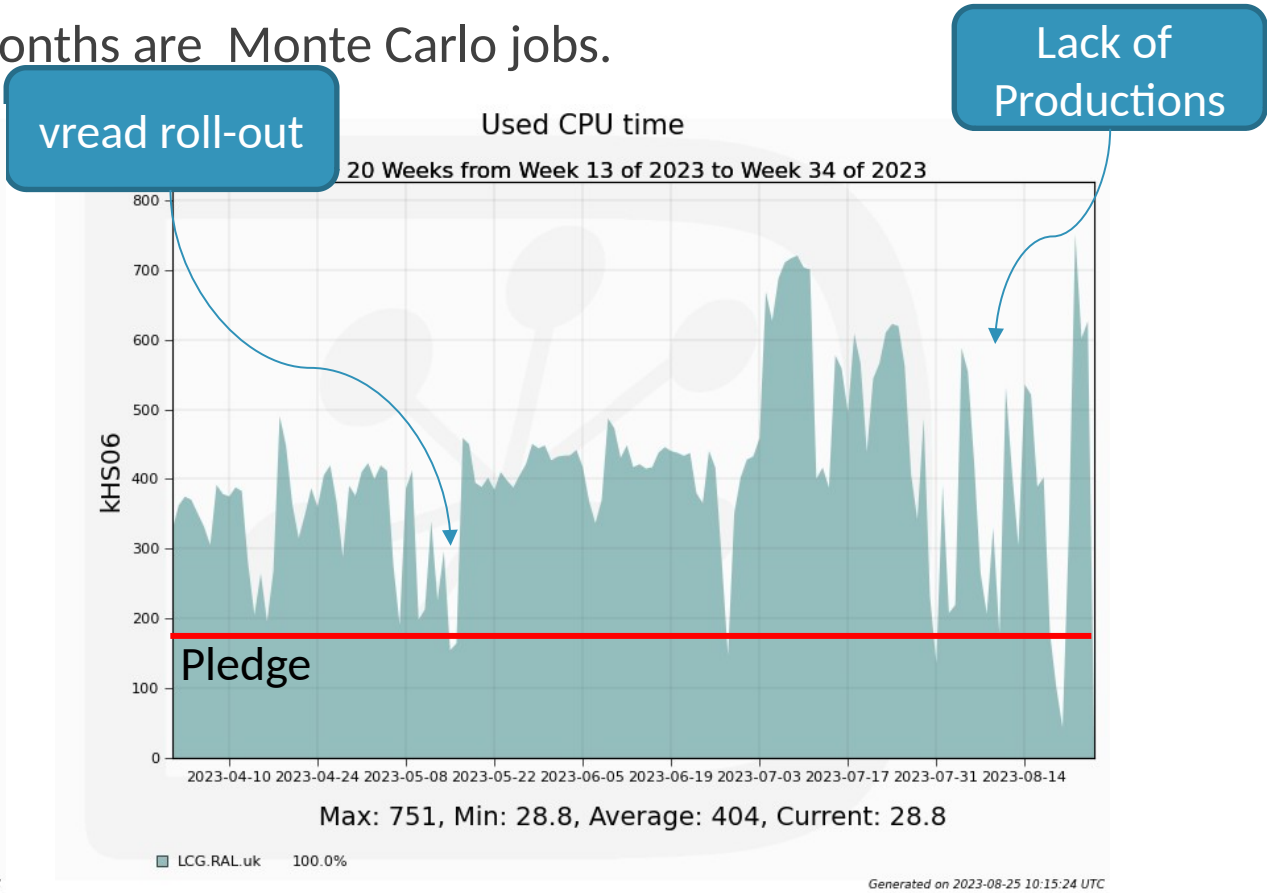

In this talk I'll try to cover the following topics:

1. LHCb T1 resource usage for the last half year.
2. RAL T1 problems and solutions.
3. LHCb T2 resource usage.
4. LHCb plans (backup).




LHCb Jobs at RAL T1

The majority of LHCb jobs at RAL for the last 6 months are Monte Carlo jobs.




Comparisson




Everything is relative. How these results are compared with other T1s/CERN?

RAL T1 has provided for LHCb the most computational resources among T1s, in terms of executed jobs and CPU time. Even more than CERN.


CPU time consumed by site



Wall time consumed by site




Executed jobs by Site




Generated on 2023-08-25 10:21:14 UTC

Tape usage



Tape usage is significantly below the pledge. This is because of various reasons.




Tape Space - TBytes

Search:

Tape Space - TBytes	Apr 2023	May 2023	Jun 2023	Jul 2023	Total	% MoU
LHCB	14,380	14,510	14,490	14,780	58,160	
Total	14,380	14,510	14,490	14,780	58,160	36%
installed capacity	0	0	0	0	0	
MoU pledge	39,890	39,890	39,890	39,890	159,560	


Disk usage



About 60% of allocated disk space is used. Interesting that there is a discrepancy between WLCG accounting and LHCb accounting – dark data is still there (not removed yet).

PFN space usage by StorageElement

20 Weeks from Week 13 of 2023 to Week 34 of 2023



Max: 8.71, Min: 7.68, Average: 8.13, Current: 8.27


■ RAL-DST 49.1% ■ RAL-BUFFER 15.8% ■ RAL-FAILOVER 0.0%
 ■ RAL_MC-DST 32.7% ■ RAL-USER 2.4%

Generated on 2023-08-25 10:25:35 UTC

Disk Space - TBytes

Search:


Disk Space - TBytes	Apr 2023	May 2023	Jun 2023	Jul 2023	Total	% MoU
LHCB allocated	12,474	12,474	12,474	14,368	51,790	
LHCB used	8,594	8,921	9,026	9,181	35,721	
Total allocated	12,474	12,474	12,474	14,368	51,790	82%
Total used	8,594	8,921	9,026	9,181	35,721	57%
installed capacity	0	0	0	0	0	
MoU pledge	15,724	15,724	15,724	15,724	62,896	



Problems: long term issues

There are/were several long-lasting issues for LHCb:

1. Vector read ([ticket](#)).
2. Slow stat calls ([ticket](#)).
3. Redirector setup ([ticket](#), backup, covered by Jyothish).




Vector read issue

Vector read (aka readv) is an xrootd request that reads multiple chunks of file identified by offset and length.




These requests sometimes fail on ECHO causing job failure. This is a serious issue for LHCb since user jobs usually process a lot of files (up to 50). A single failure during the processing of the last file causes the whole job to fail, and results for the first successfully processed files are discarded.



Vector read issue: problem


The error happens when readv operations takes too long.

- There is a “stream timeout” in xrootd – if nothing is transferred in the data channel for the given amount of time, failure will be declared.
- It is possible to increase this timeout via environment variable.
 - It was tried multiple times, without success.
- Vector read implementation in xrootd-ceph plugin was inefficient – every request was executed sequentially.




Vector read issue: current configuration

To mitigate the problem, RAL's WNs have different configuration (from external gateways). There is an additional caching instance of xrootd, which tries to read data from ceph using huge blocks.

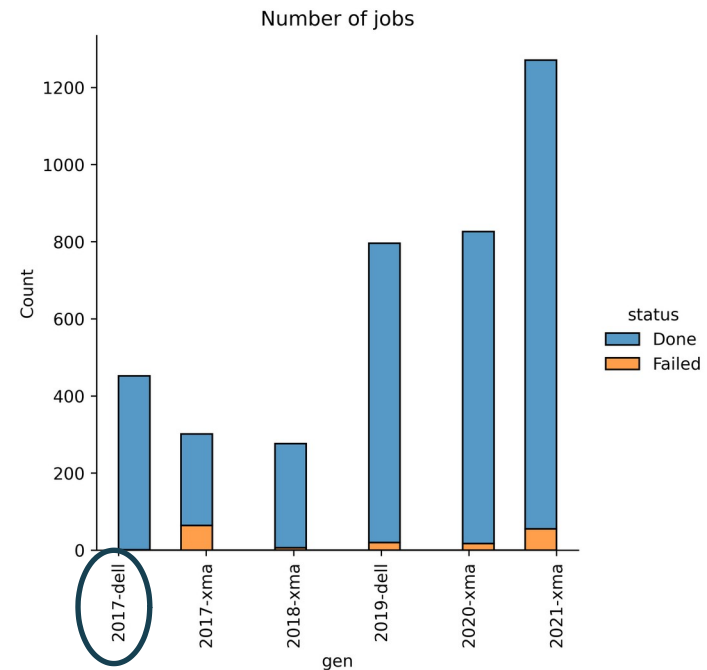
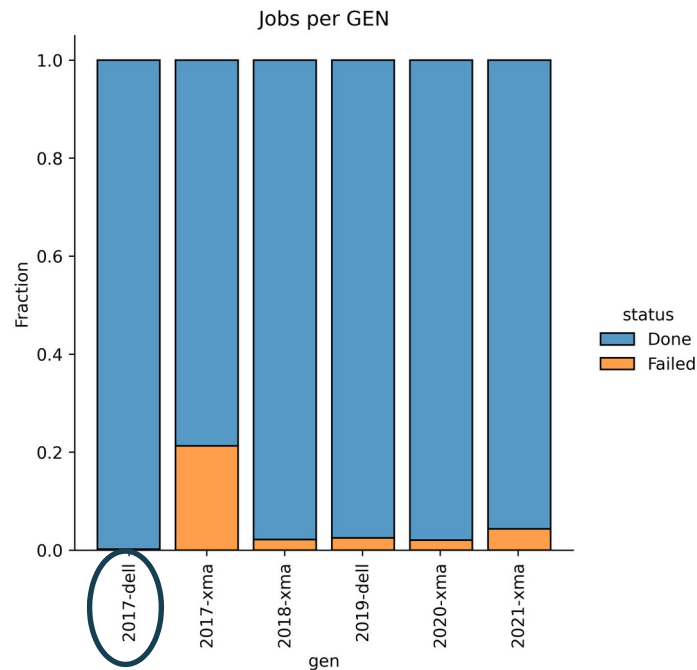


This mitigates the problem but does not resolve it completely (timeouts still happen sometimes). Furthermore, this configuration introduces additional problems.



Vector read issue: story

- A patch that improves read performance was developed.
- The initial plan was to deploy this patch (without buffering layer) **and** remove xcache.
- The patch was rolled out to one generation of WNs, it turned out to be successful.

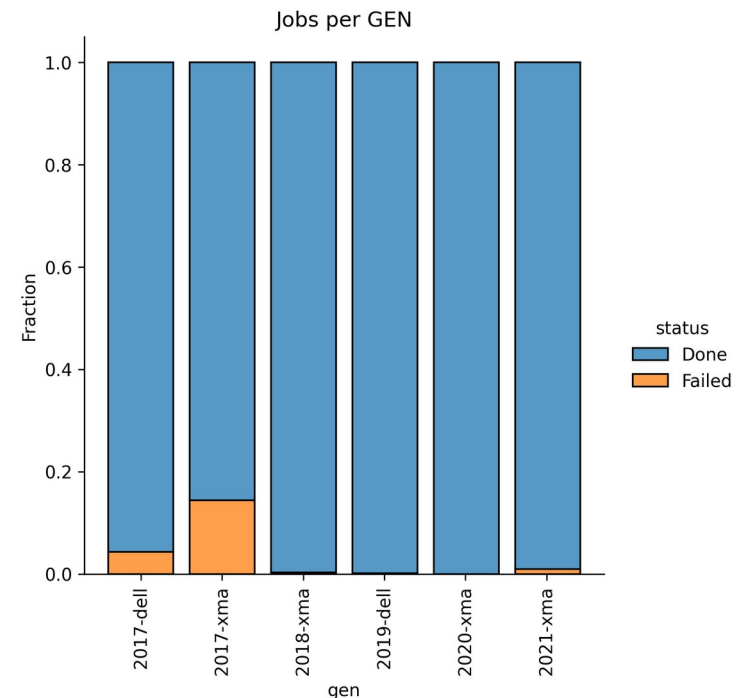
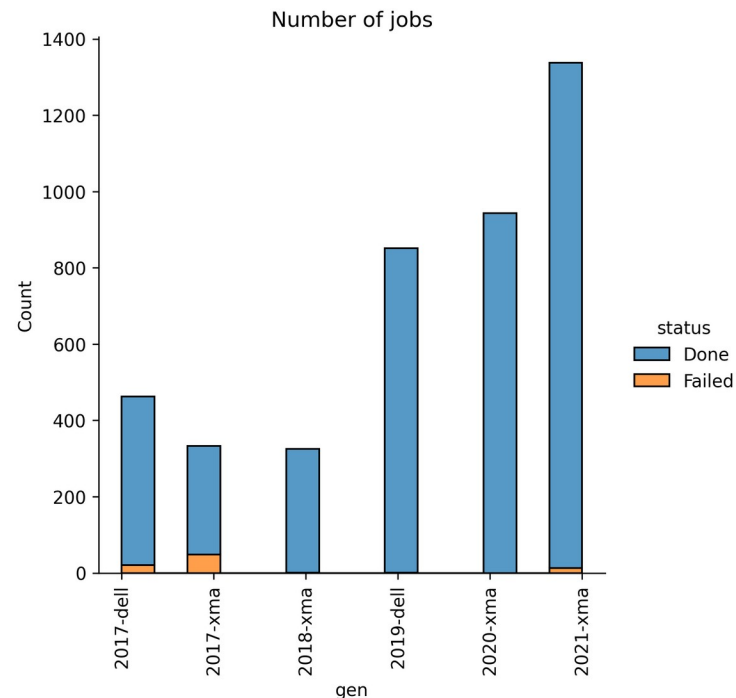


Vector read issue: story

- After this test the decision was taken to roll-out the patch to the whole farm.
- Roll-out was unsuccessful – the number of IOPs increased significantly, and ECHO was not able to cope with it.
 - Buffering layer could have helped.
 - However, it was not included in the code, unfortunately.
- Eventually the changes were rolled back.
- After some time decision was taken to keep xcache and roll-out the patch again.
 - In this configuration xcache is present in order to reduce the number of IOPs coming to echo.

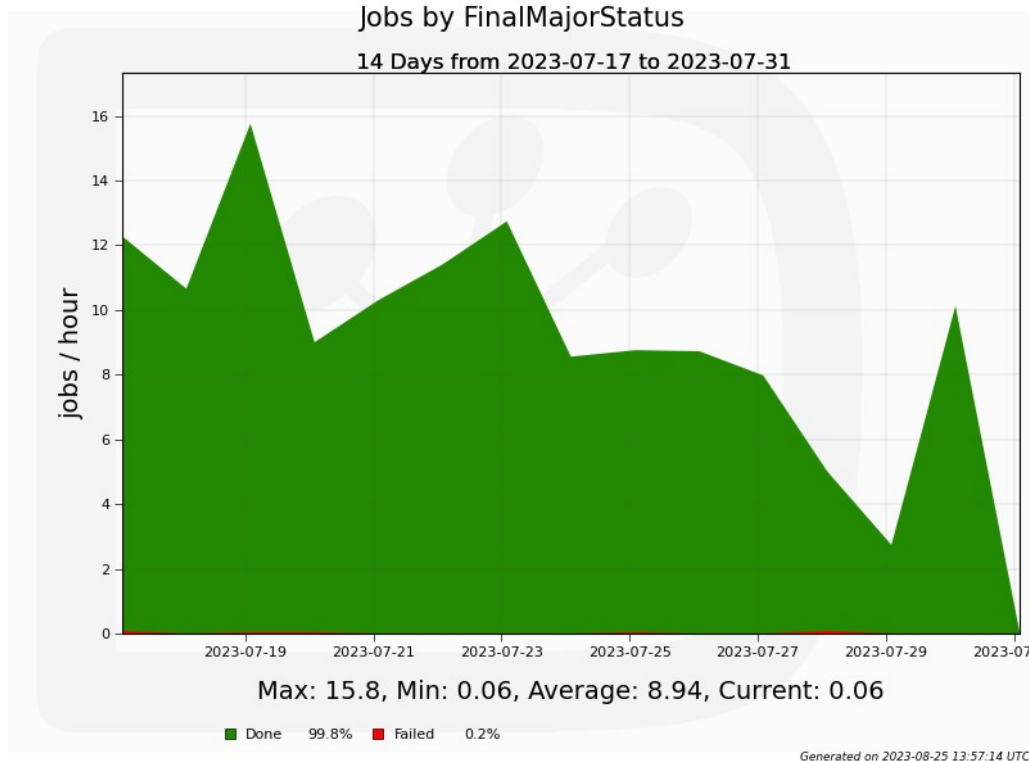
Vector read issue: story

- Second roll-out was successful, however, job failure rate was still high.
- It turned out that most of the failures come from 2017 generations, some – from 2021 and almost none from 2019, 2020.



Vector read issue: story

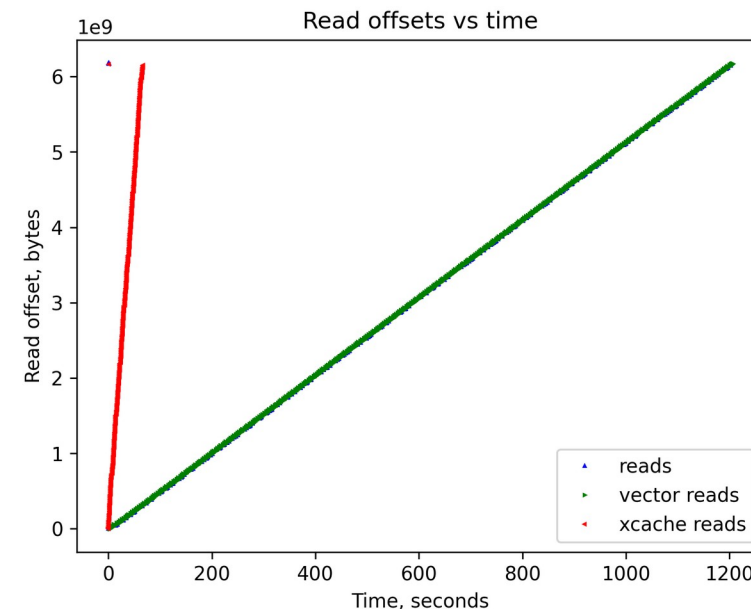

- Then 2017 gens were retired for LHCb.
- The failure rate stayed low for quite some time, and after that the ticket was closed.



update#164 Alexander Rogovskiy	2023-07-19 09:52	Int. Diary: Added attachment jobs_all.png https://ggus.eu/index.php?mode=download&attid=ATT117910
update#165 Alexander Rogovskiy	2023-07-19 09:52	Public Diary: Hi, So, after the retirement of 2017 gens for LHCb, it looks like the failure rate is decreased. Here are two plots of my test jobs: the first for the whole period of testing, and the second for the last week. Failure rate looks lower on the second one. Best Regards, Alexander
update#166 Alexander Rogovskiy	2023-07-19 09:53	Int. Diary: Added attachment jobs_last_week.png https://ggus.eu/index.php?mode=download&attid=ATT117911
update#167 Brian Davies	2023-07-20 08:37	waiting for reply (NGL_UK) Public Diary: Site understands this problem to be effectively solved and are aware VO is inclined to keep ticket open until after User Analysis jobs are re-run at site.
update#168 GGUS SYSTEM	2023-07-27 10:37	Int. Diary: Sent 1st reminder to ticket submitter (raja.nandakumar@sific.ac.uk) requesting input.
update#169 christophe.denis.haen	2023-07-31 09:42	Public Diary: I have no obvious reason to believe that it is indeed not fixed. So I believe we may indeed close the ticket! We should throw a party :-)
update#170 andrew.mcnaab	2023-07-31 09:45	Public Diary: Well done everyone! (Fingers crossed)
update#171 Alexander Rogovskiy	2023-07-31 10:06	Public Diary: Hi Chris, Ok, thanks for your feedback! And sorry for all the inconvenience this issue caused.. Best Regards, Alexander

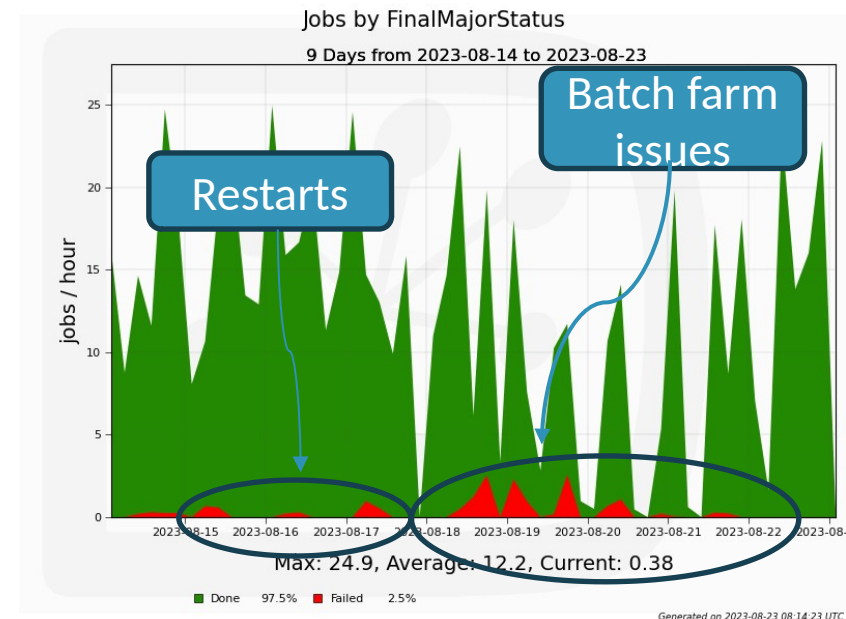
Vector read issue: story

- Following Murphy's law, almost immediately after the ticket was closed the issue reappeared.
 - The errors were coming from 2021 gen.
 - Most probably it was caused by insufficient IOps speed from local HDD.
 - Very aggressive prefetch configuration was used, it made the problem worse.




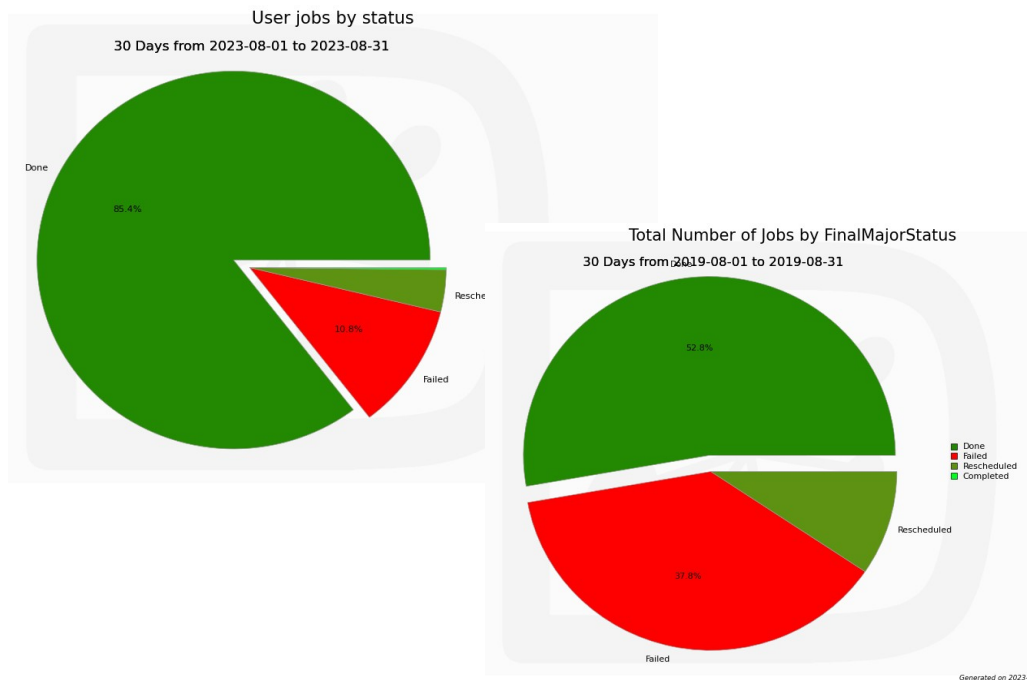
Vector read issue: story


- Test showed that turning prefetching off should reduce failure rate.
- This configuration (with prefetch off) was rolled-out to 2021 generation.
- Errors were still present after the roll-out, however they were only (well, almost) related to gateway restarts (happens daily).
- Timeout increase should make restarts more transparent.
- Approx. two weeks ago this timeout increase was applied to my test jobs on the client side.
- So far looks OK.



Vector read issue: story

The whole story in 3 pictures.







Slow “stat” calls

There is a ticket with a complaint about “Slow stats” at RAL’s ECHO.

It turned out that it is not the stats that are slow, but the checksum requests (different request in terms of xrootd).

A patch was developed to improve checksum retrieval speed. It should be deployed after redirector setup.






Tier-2 statistics

Normalized CPU usage by Site

21 Weeks from Week 13 of 2023 to Week 35 of 2023




Max: 617,340, Average: 309,578

LCG.Manchester.uk	22.6%	LCG.UKI-LT2-IC-HEP.uk	7.8%	LCG.ECDF.uk	2.4%
LCG.UKI-LT2-QMUL.uk	16.6%	LCG.UKI-LT2-RHUL.uk	6.0%	LCG.UKI-LT2-Brunel.uk	1.9%
LCG.Lancaster.uk	11.0%	LCG.Oxford.uk	5.9%	VAC.BHAM-HEP.uk	1.6%
LCG.Liverpool.uk	9.8%	LCG.Bristol.uk	3.4%	LCG.Sheffield.uk	0.2%
LCG.RAL-HEP.uk	8.2%	LCG.DURHAM.uk	2.7%		

Generated on 2023-08-25 14:39:45 UTC

Job CPU efficiency by Site

21 Weeks from Week 13 of 2023 to Week 35 of 2023




Generated on 2023-08-25 14:49:05 UTC

Tier-2 statistics


Total Number of Jobs by Country

21 Weeks from Week 13 of 2023 to Week 35 of 2023




CPU days used by Country

21 Weeks from Week 13 of 2023 to Week 35 of 2023




Wall time days used by Country

21 Weeks from Week 13 of 2023 to Week 35 of 2023




UK	5798744.9
CERN	2846862.8
IT	2307822.4
FR	2036445.4
DE	1391016.4
PL	1057842.1
CH	546552.7
RU	457680.1
NL	430304.6
US	364311.4
ES	270301.3
BR	268714.4
CN	208423.8
RO	138905.7
AU	28490.2
IL	10817.3
MULTIPLE	0.0

Generated on 2023-08-25 14:40:18 UTC



Summary

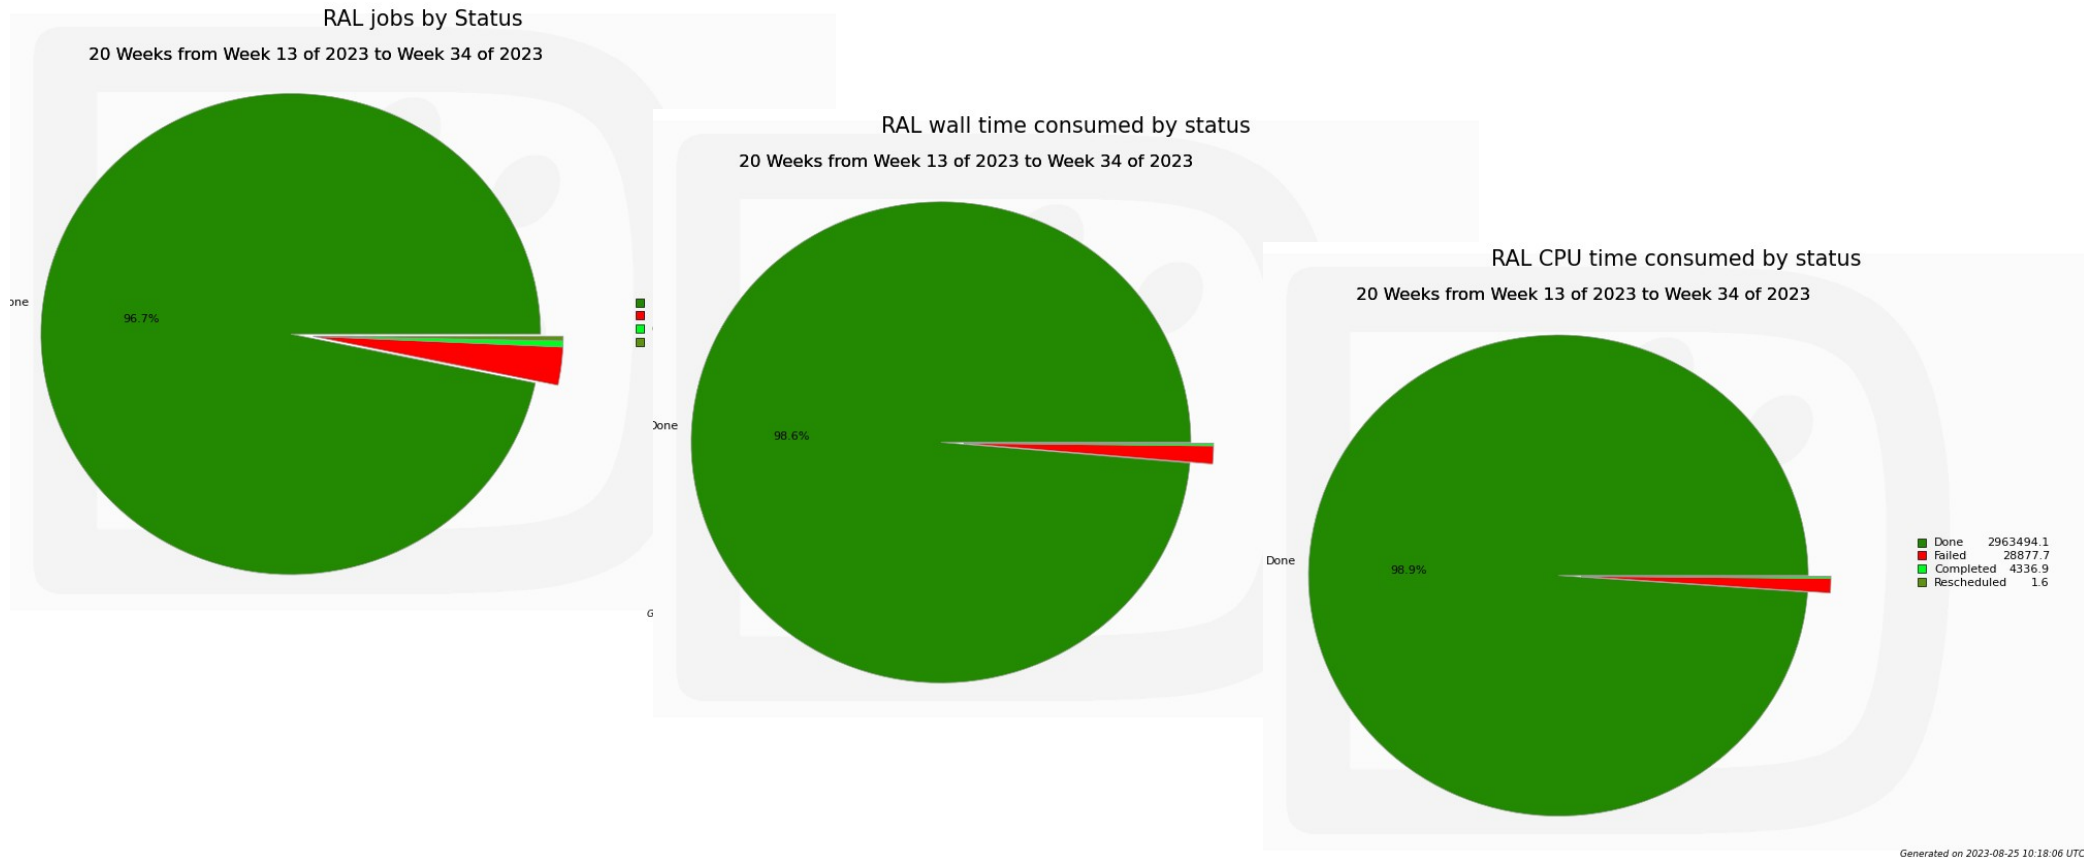
- UK sites provided a lot of resources to LHCb during last 6 months.
- In general operations were smooth.
- There are some issues, but their solutions are progressing.




Ticket summary (backup)

Site	Tickets opened since Oct 2022	Open tickets
RAL-T1	11	2
UKI-LT2-Brunel	1	1
UKI-LT2-QMUL	2	0
UKI-NORTHGRID-LANCS-HEP	1	0
UKI-NORTHGRID-LIV-HEP	1	0
UKI-NORTHGRID-MAN-HEP	2	0
UKI-NORTHGRID-SHEF-HEP	2	0
UKI-SCOTGRID-DURHAM	1	0
UKI-SCOTGRID-ECDF	2	0
UKI-SOUTHGRID-RALPP	10	0
UKI-SCOTGRID-GLASGOW	1	0
UKI-LT2-IC-HEP	1	0
All	35	3

RAL T1 Job success rate (backup)




RAL T1 CPU efficiency (backup)




Transfers to ECHO (backup)

Low-rate failed uploads from RAL WNs to ECHO are always present.



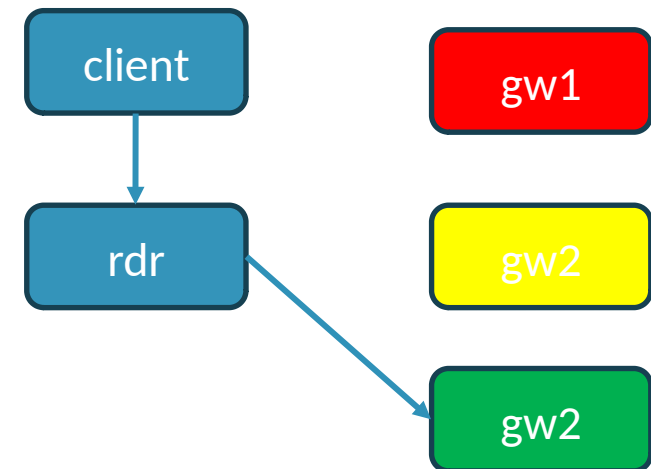
Generated on 2023-08-25 10:25:24 UTC



Generated on 2023-08-28 12:36:57 UTC

Transfers to ECHO (backup)

- These problems usually happen because of the gateway overload.
- One way to mitigate this is to add more gateways (pending).
 - We may face the same problem when the load increases.
- The other solution is to set up an xrootd redirector.
 - Redirector should select “less busy” gateway for every transfer.
 - Redirector deployment at RAL is in progress.
 - LHCb switched their https transfers to redirector in July.
 - There is no significant improvement so far, probably because of the lack of gateways.
- Seems like a combination of both are needed.



Plans (backup)

- LHC is back, more data to come!
- Switched from xrootd to https for tape recalls at RAL
 - To be tested
- Pledges to be reduced(!), because of underusage
- Switch from gsiftp to arex for ARC CE
 - Ongoing, some sites have already switched
- Token migration
 - HTCondor CEs have already migrated
 - ARC CEs to follow
 - For SEs it is still work in progress (close to the beginning)
 - New ETF tests to be developed