Science and Technology Facilities Council
Rutherford Appleton Laboratory

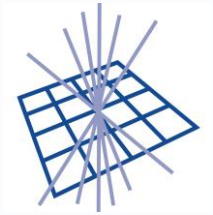

# LHCb status

GRIDPP 52

alexrg

# Contents

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

1. LHCb UK resource usage for the last half year.
2. UK sites problems and plans
   - Focusing mostly on RAL T1
3. LHCb news.

Science and Technology Facilities Council
Rutherford Appleton Laboratory

# LHCb Jobs at RAL T1

Nowadays it's a little tricky to calculate average HS23 consumption…

- DIRAC code for calculating normalization factor is being updated, so can not be trusted
- If we use static normalization factor from RAL, we have:
  - 6446*12.7 / 0.4 = 204660 HS23
  - Blue = Wallclock years; purple = Normalization factor, green = = years in the reporting period
  - Pledge is 180 kHS23

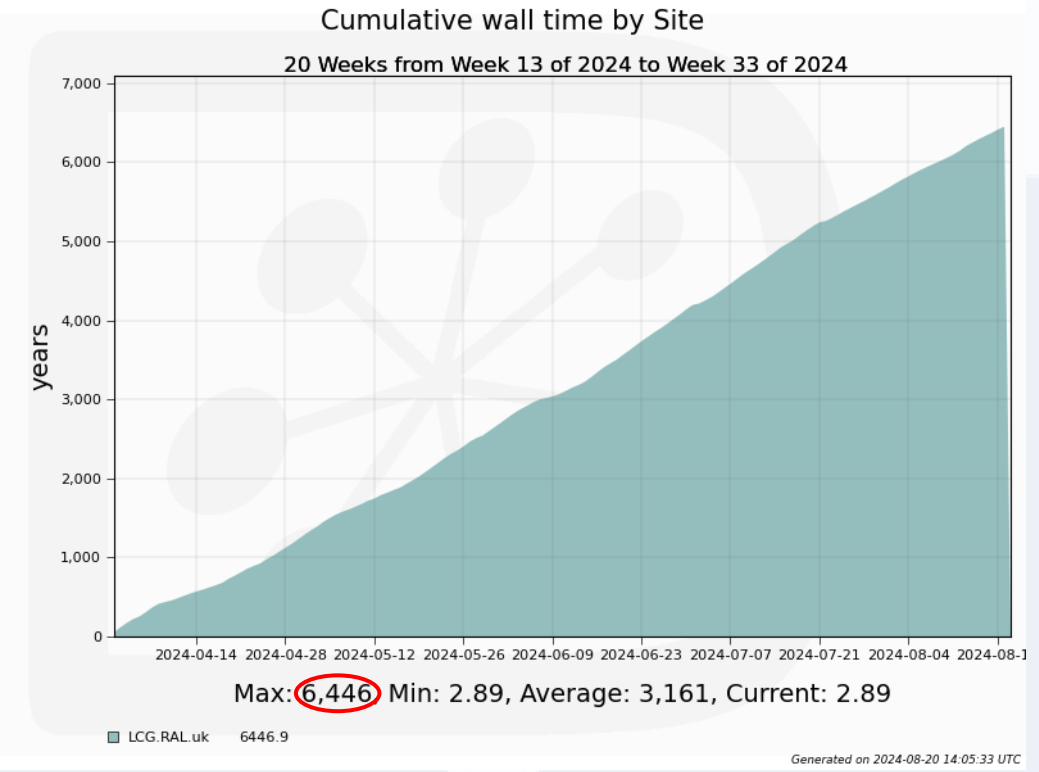

# Comparison

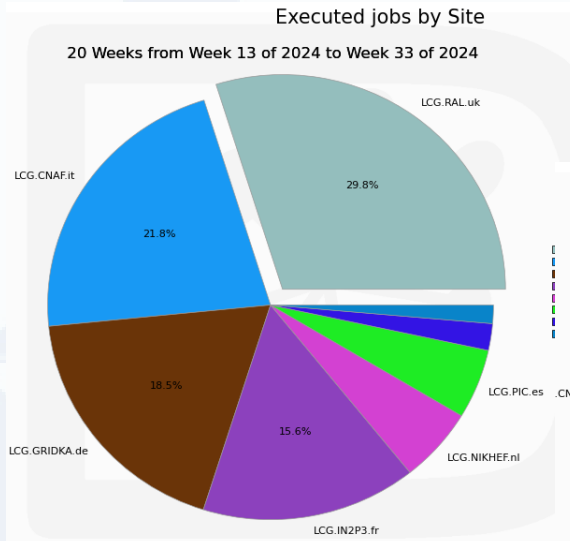


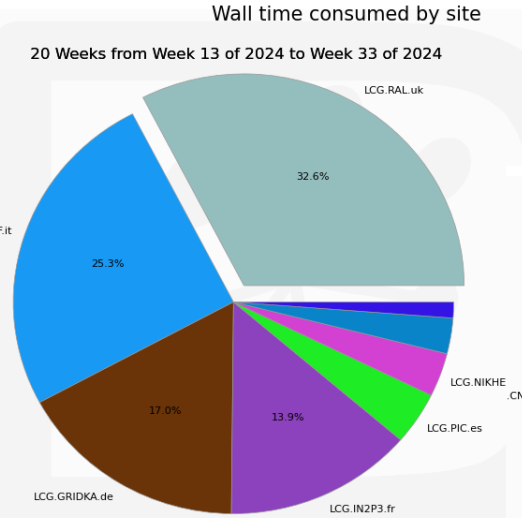


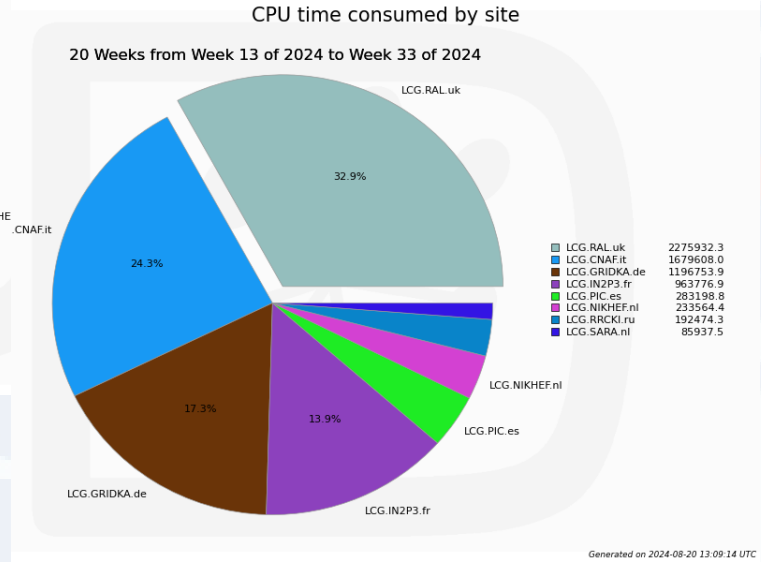


RAL has provided the biggest amount of CPU resources among all T1 sites, in terms of executed jobs (first plot), walltime (second), and cputime (third).

# Tape usage

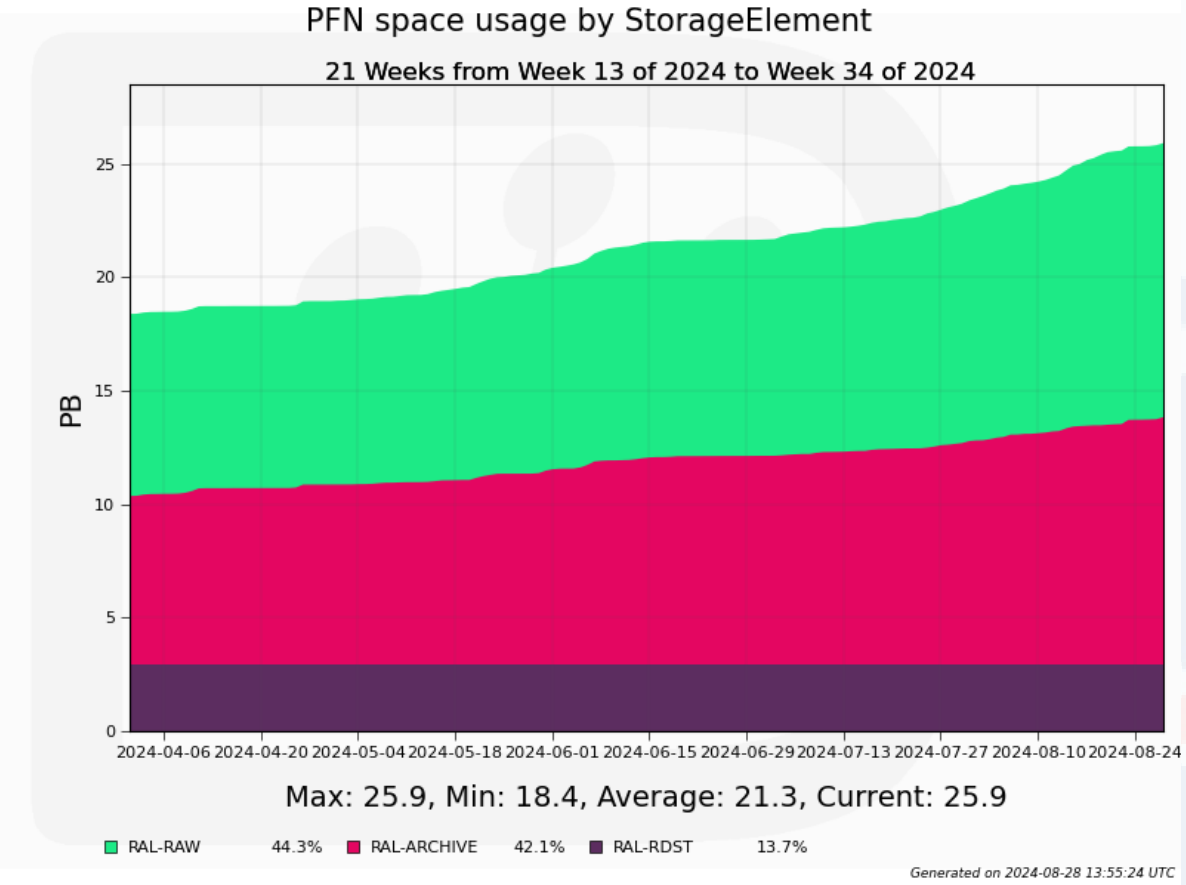


- Tape usage increased significantly
  - ~78% of the pledged space is used
    - Contrary to ~56% at the beginning of the Data Taking this year
  - Pledge for this FY is 33PB
  - Data Taking campaign is still ongoing!
    - Likely to use 7.2PB more if the rate remains the same

# Disk usage

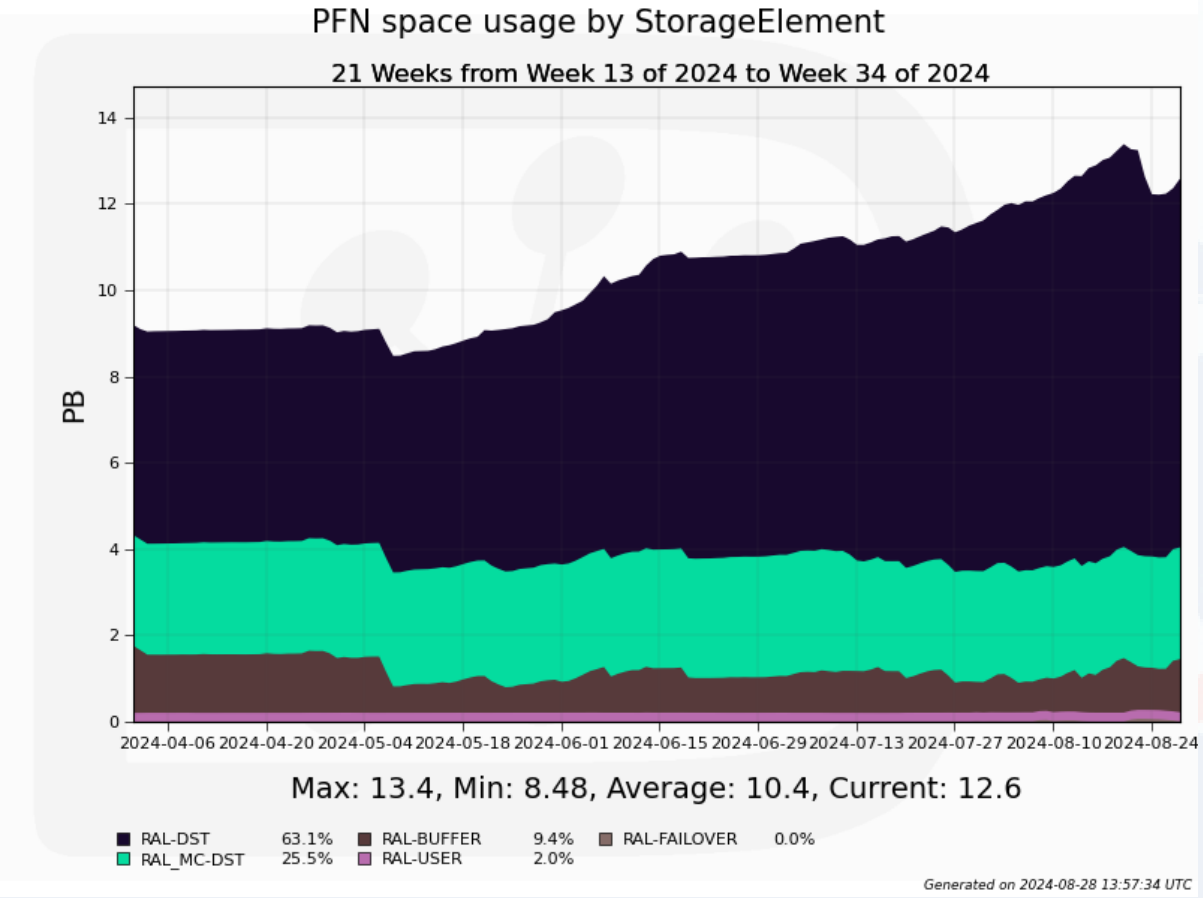


- Disk usage increased significantly
  - ~80% of the pledged space is used
  - Pledge for this FY is 15.7PB
  - Data Taking campaign, which is still ongoing!

# Tier-2 statistics

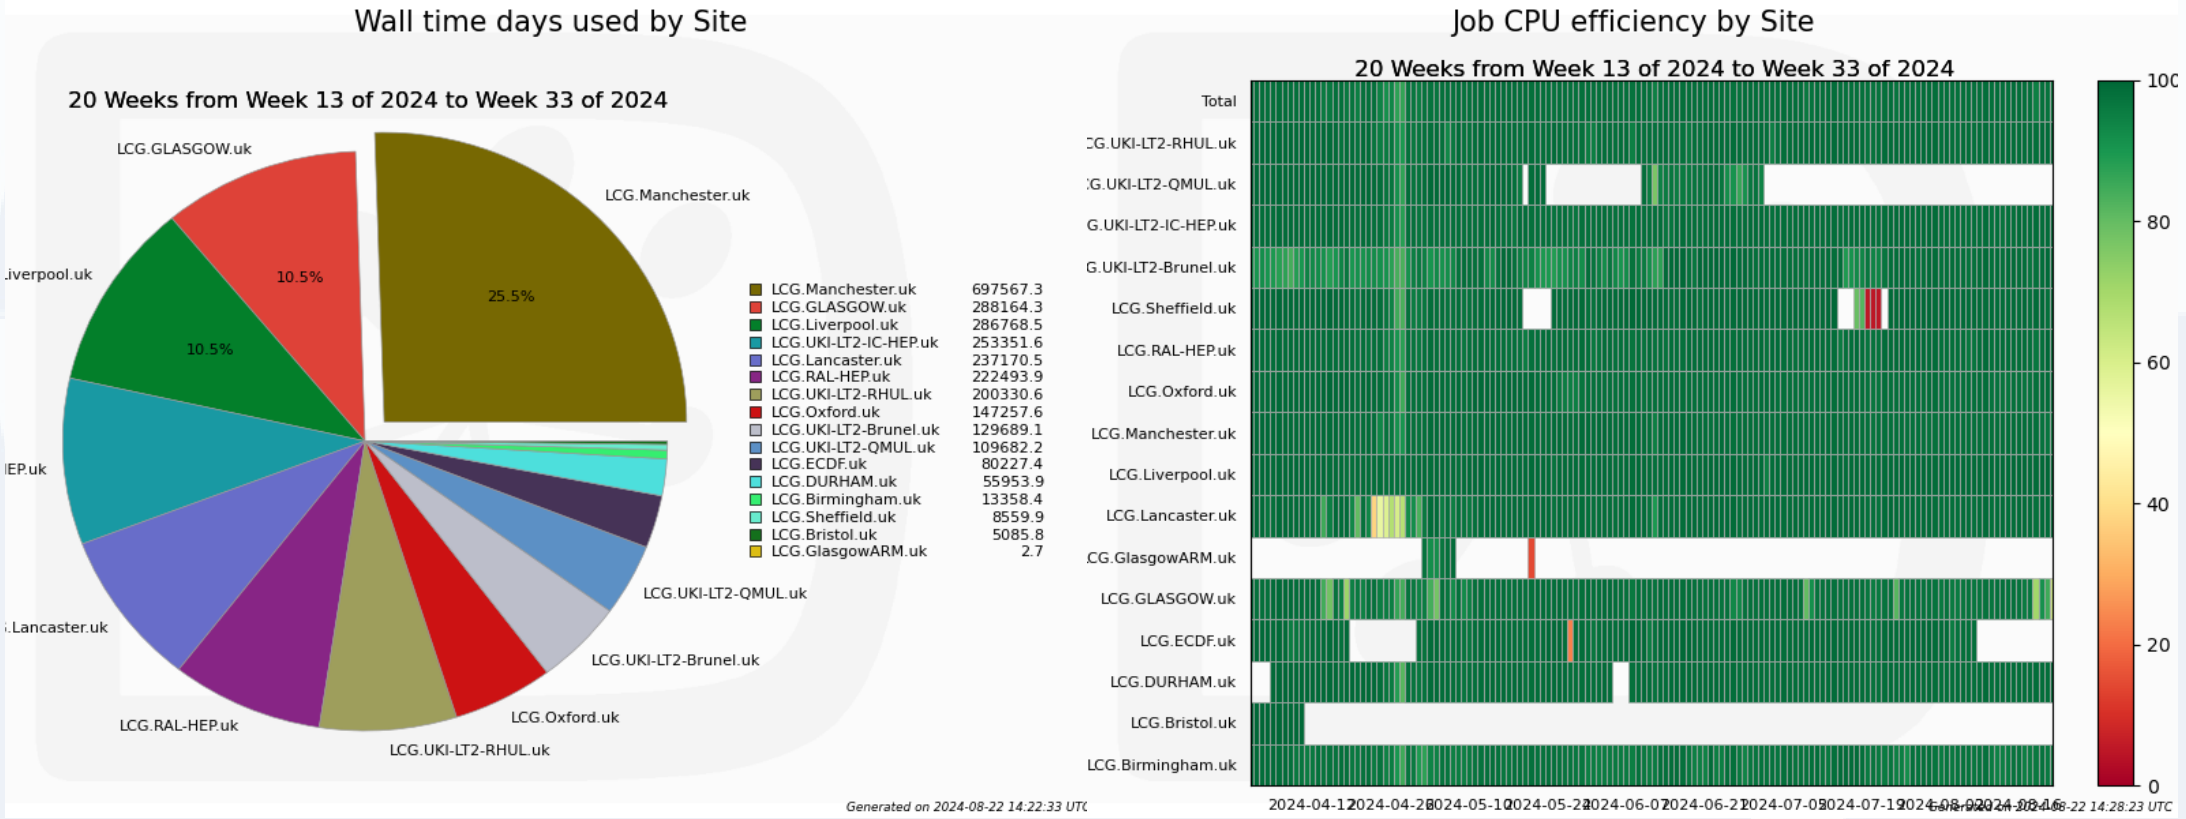

# Tier-2 statistics

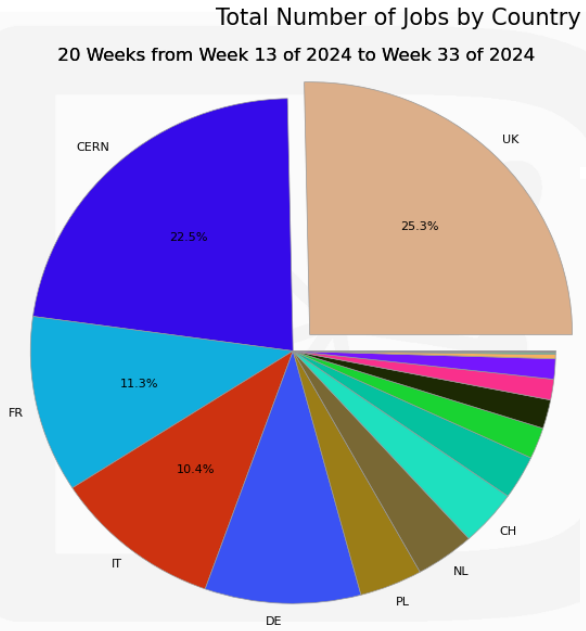


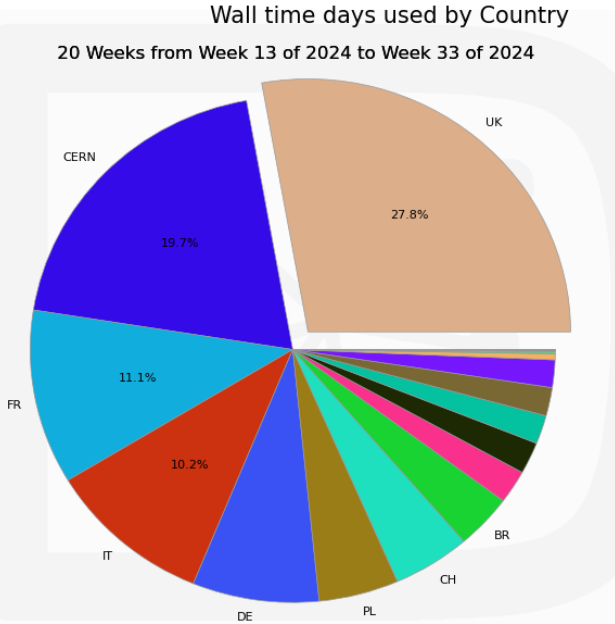


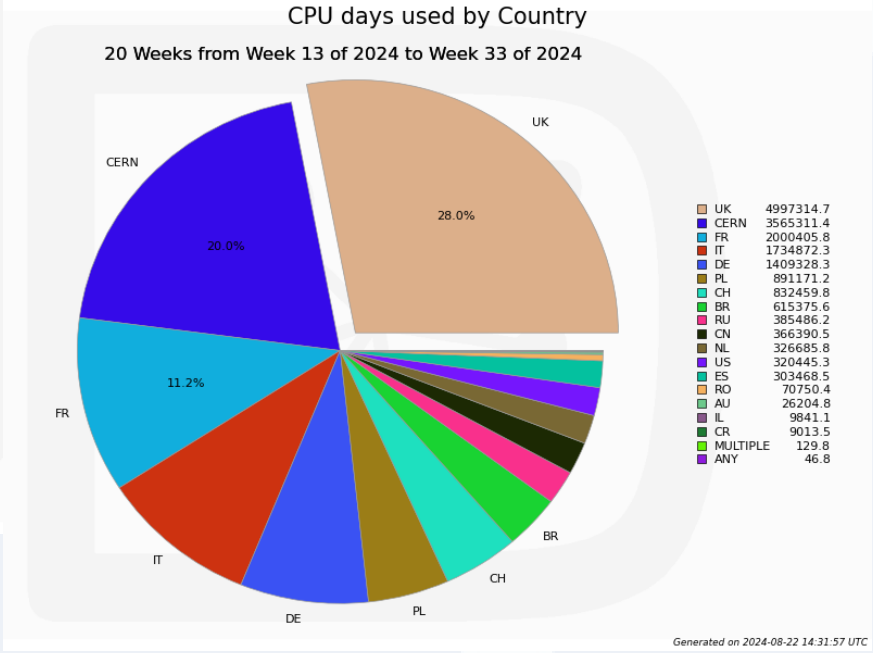


UK provides more computing resources than CERN!

- According to the pledges, though
  - Wall/CPU time fraction relation is higher than the pledge

# Problems

At the time of writing, there are 5 opened LHCb tickets against UK sites:

**5 of 5 Tickets**

| Ticket-ID | Type | VO | Site | Priority | Resp. Unit | Status | Last Update | Subject | Scope |
|---|---|---|---|---|---|---|---|---|---|
| 167910 | Team | lhcb | UKI-LT2-Brunel | very urgent | NGI_UK | in progress | 2024-08-20 | Jobs Failed at UKI-LT2-Brunel | WLCG |
| 167852 | Team | lhcb | UKI-SOUTHGRID-OX-HEP | very urgent | NGI_UK | in progress | 2024-08-12 | Pilots failing at Oxford | WLCG |
| 167682 | Team | lhcb | UKI-SCOTGRID-GLASGOW | very urgent | NGI_UK | in progress | 2024-08-21 | All FTS transfers to Glasgow are failing | WLCG |
| 167007 | Team | lhcb | UKI-SCOTGRID-ECDF | very urgent | NGI_UK | in progress | 2024-08-13 | Pilots Failed UKI-SCOTGRID-ECDF | WLCG |
| 163853 | Team | lhcb | UKI-NORTHGRID-SHEF-HEP | urgent | NGI_UK | on hold | 2024-07-24 | Failed jobs at LCG.Sheffield.uk | WLCG |

Most of them are awaiting closure, the last one (for Sheffield) is awaiting actions from LHCb.

At RAL there is still one (or “half”?) long-lasting issue for LHCb:

1. Direct access issues (ticket, closed, new ticket, closed).

# Direct access issues

- Some LHCb jobs (namely user and WGProduction) do not download data, but access it directly from the storage using xrootd's (vector) read requests
- Vector read requests were problematic for ECHO for a long time, until ~Aug 2023
- This year the number of WGProduction jobs at RAL increased, and significant failures started to appear occasionally (when large WGProds are submitted), causing also download failures (see below)
- In July 2024 a new ticket was opened

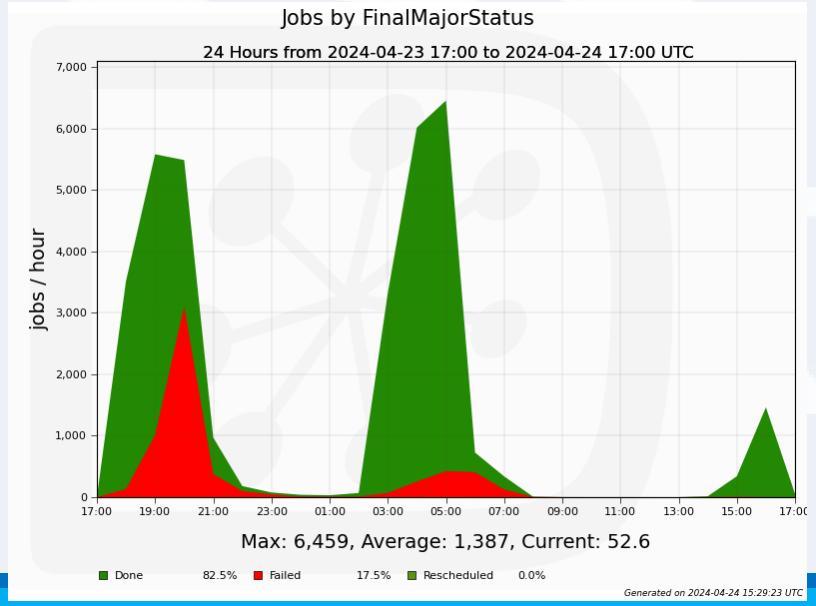


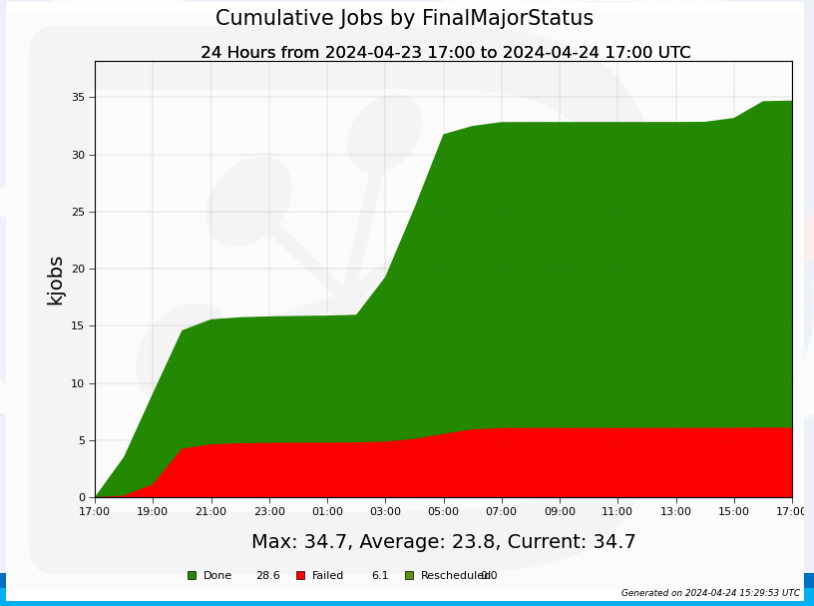


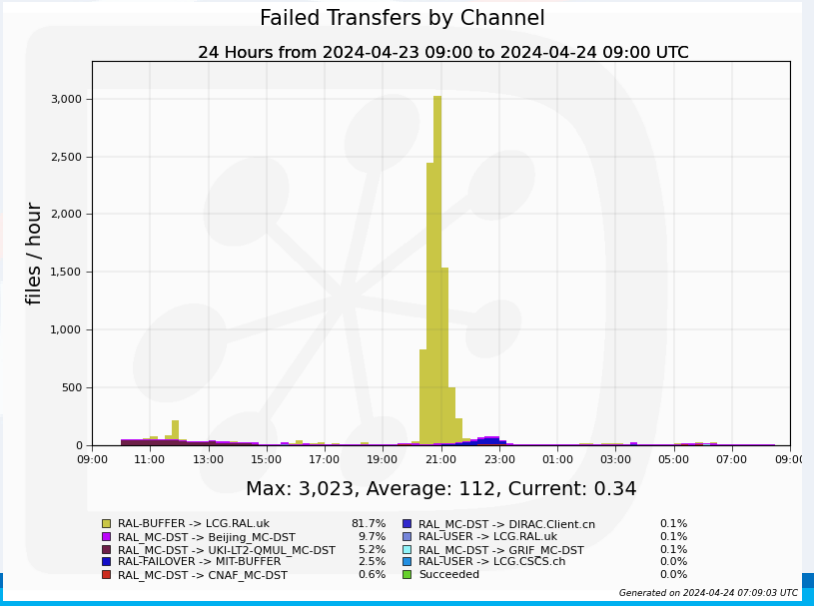

# Direct access issues: memory

- Most of the errors were happening on the 2021 Generation of WNs
  - WNs with 256 slots and SATA SSDs
- The most popular error was "`Cannot allocate memory`", reported by xrootd proxy
  - It was confirmed that Proxy's memory consumption can be close to its limit
- Failed downloads were caused by Ceph being overload with IOps
  - When proxy runs out of memory, it forwards read requests directly to the gateway
  - Ceph does not like when it is asked to do many small reads

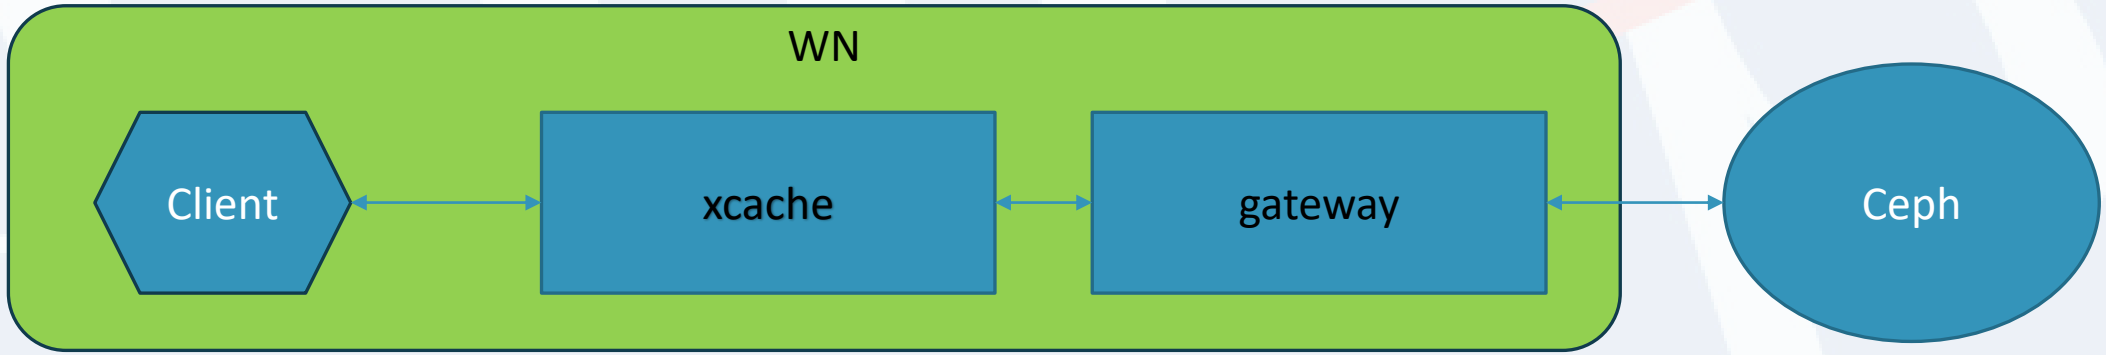

# Direct access issues: memory

- Turned out that proxy's memory limit(s) was(were) low (8GiB) and independent on the number of cores
  - Both container limit and xrootd limit
- The increase of the limit reduced ECHO IOps, but did not help to solve (Vector) read errors
- Additional changes were applied to reduce proxy's memory consumption
  - Prefetching was turned off
    - The change was planned long time ago for another reason, but could have helped to deal with memory consumption as well
  - LHCb jobs were moved from 2021 Gen to 2022
    - 2022 Gen has NVMe SSDs, allowing faster writes (and therefore faster memory release)
- All was in vain…
  - Errors still happened occasionally

# Direct access issues: problem

- Luckily, it turned out that the error can be reproduced manually
  - To do so one has to run many “jobs” doing ReadV requests + a job that is doing ordinary file downloads
- Turned out that the error was caused by a bug in xrootd and is not a simple memory excess
  - When proxy runs out of memory, it forwards all incoming requests to the gateway
  - Read requests are converted to ReadV, where read vector contains only one element
    - Conversion is very simple: Read coordinates are put into read vector, no other changes are made
  - In XrootD there is a limit on max chunk size in read vector in ReadV; for Read requests there is no limit
    - The limit is lower than proxy’s block size and XRootD’s buffer size (defines internal read lengths)
  - If a Read request is big enough, it can exceed the limit when converted to ReadV
    - When this happens, gateway sends an error “`Single readv transfer is too large`”
    - Due to XrootD creativeness, it becomes “`Cannot allocate memory`” when forwarded to the client

# Direct access issues: problem

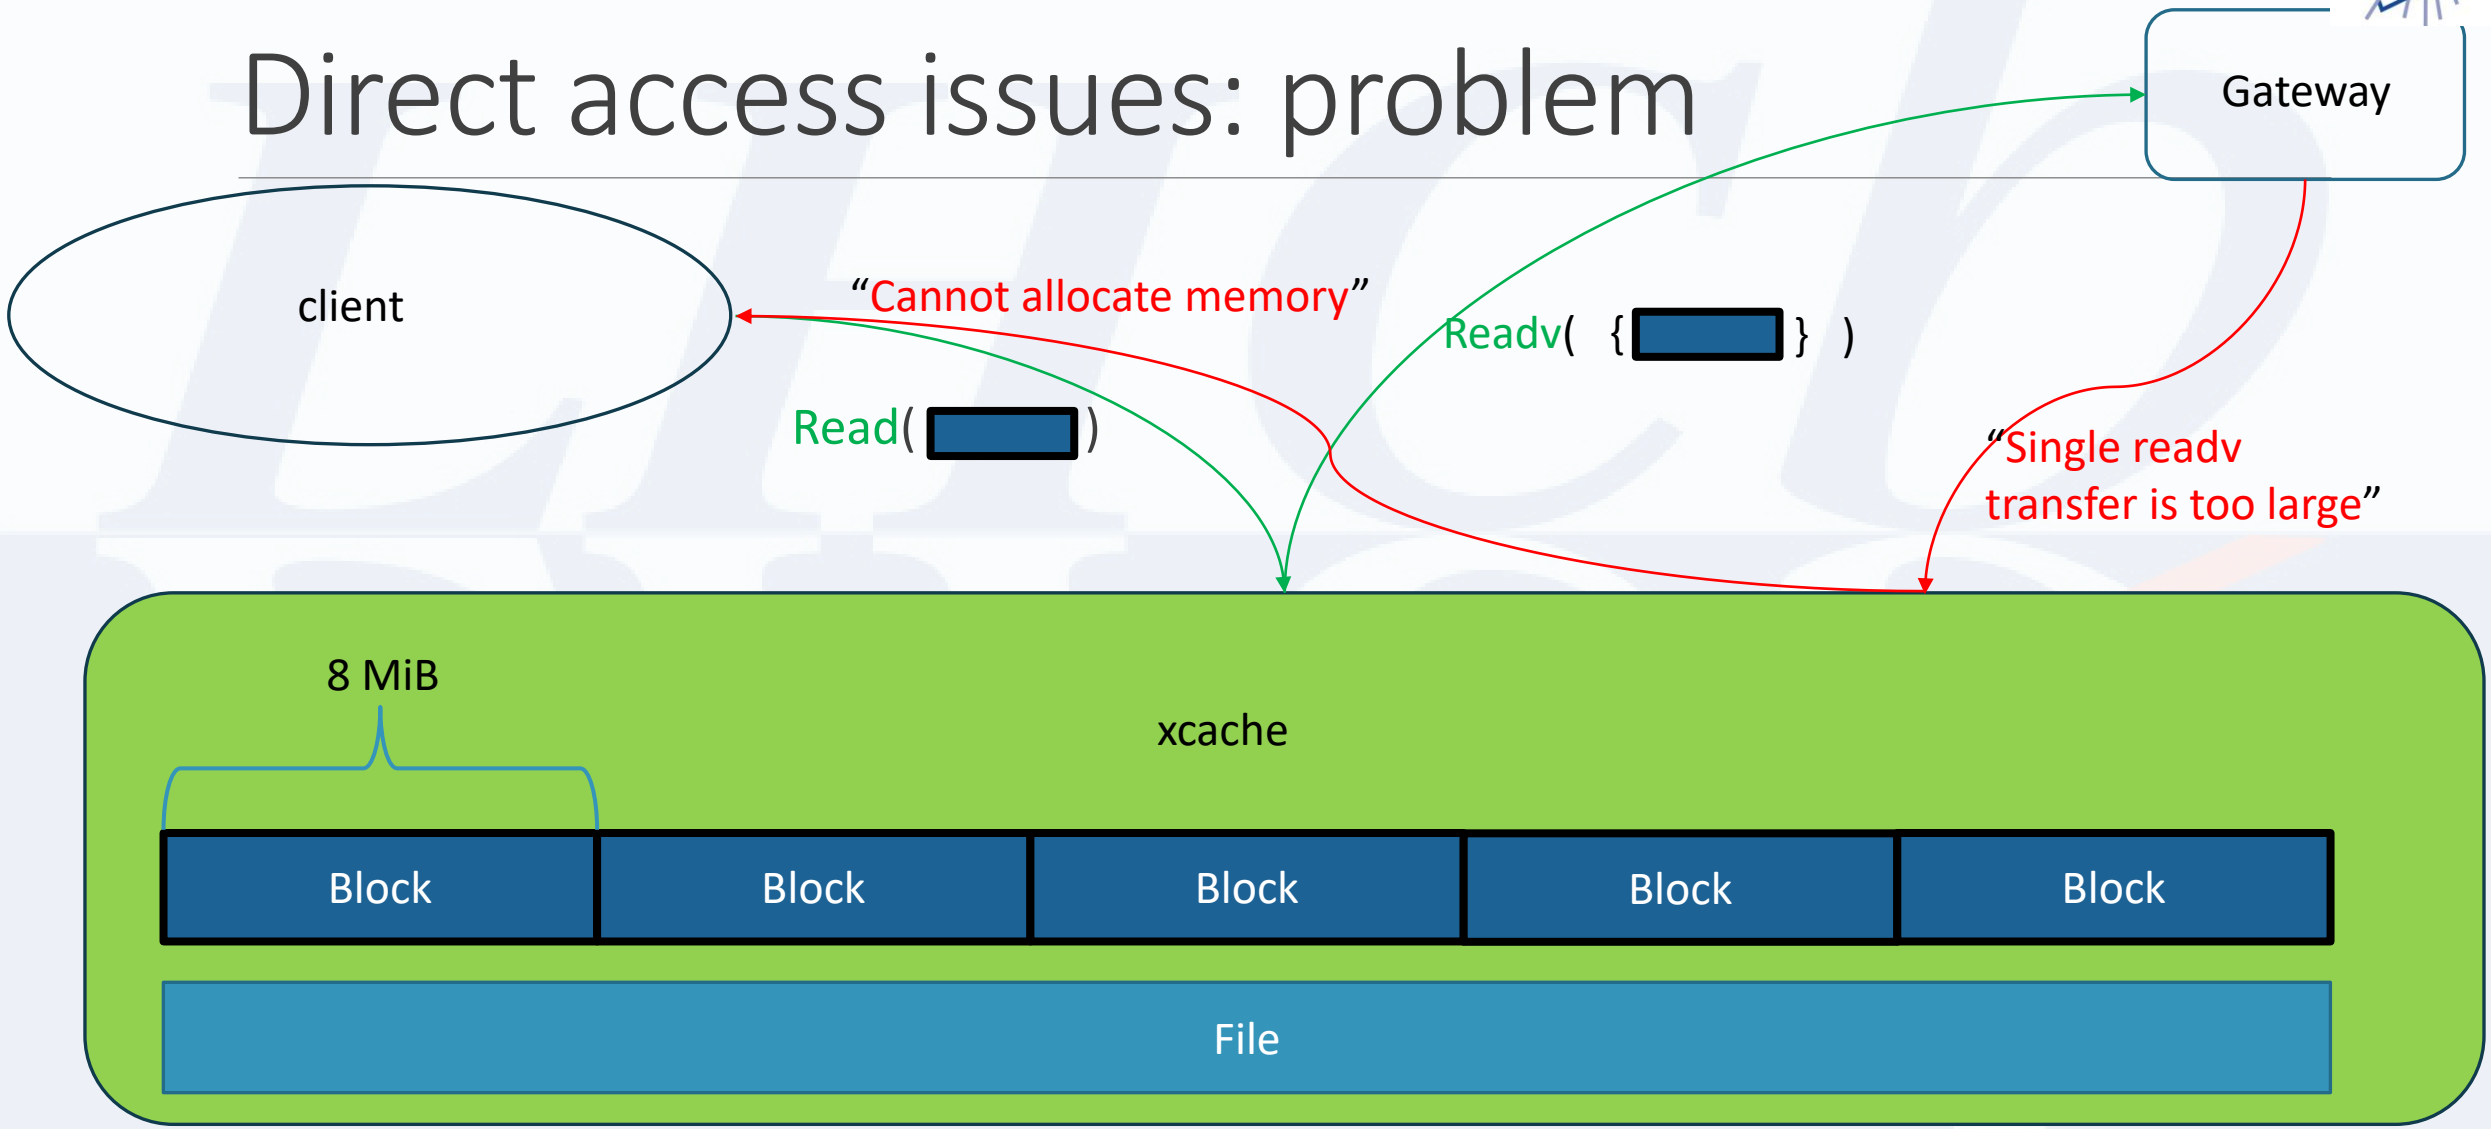

# Direct access issues: solution

- The issue can be mitigated by turning Paged Read requests on
  - In this case file downloads will use smaller block size
  - Direct access requests still can trigger the issue though
- A github issue is opened to fix the problem properly
- Other issues were affecting Vector read performance as well
  - Genuine lack of memory on WN gateways
    - Fixed by additional memory allocation
  - Daily restarts of xrootd services on WNs
    - Jobs became less tolerant to these restarts after prefetch had been turned on
    - Fixed by redesigning the restart script
- ~~Once the github issue is resolved, we can call the problem solved~~
  - Github issue is resolved, the fix should be present in v5.7.1

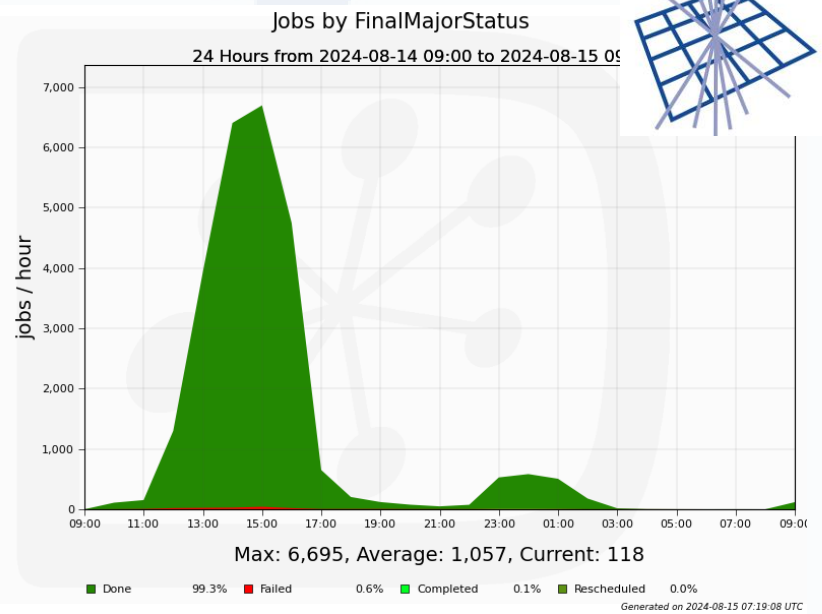


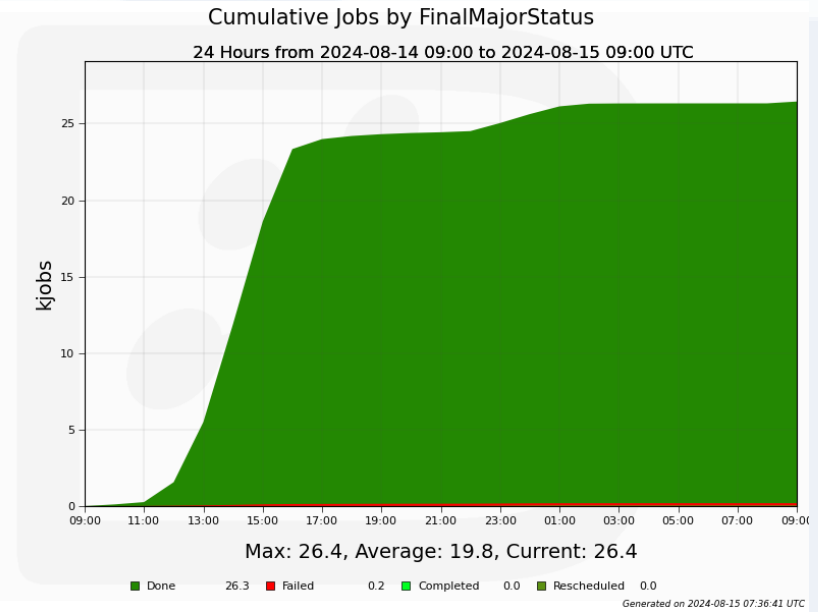

# Further WN XRootD improvements

- Multiple times we have seen external gateways using their full network throughput
  - E.g. during DC24
- These gateways are used for WN uploads as well, since proxies on WNs are read-only
  - For LHCb these WN uploads are the biggest throughput consumer (see plot)
- Can we make local WN gateways writable?
  - In principle, it is possible to patch proxy so that it forwards all write requests to the gateway
    - However, there are a few things to tune
      - Authentication: currently gateways can not authenticate clients via X509/tokens
      - Writing protocols: currently LHCb writes only via https
        - Since xrootd writes can cause file loss due to retries and our multi-gateway setup
          - Potentially can be overcome by turning retries off or setting up https on the WN gateways
    - Remove proxy completely?
      - Increases IOps on ECHO significantly, so should be tested with care

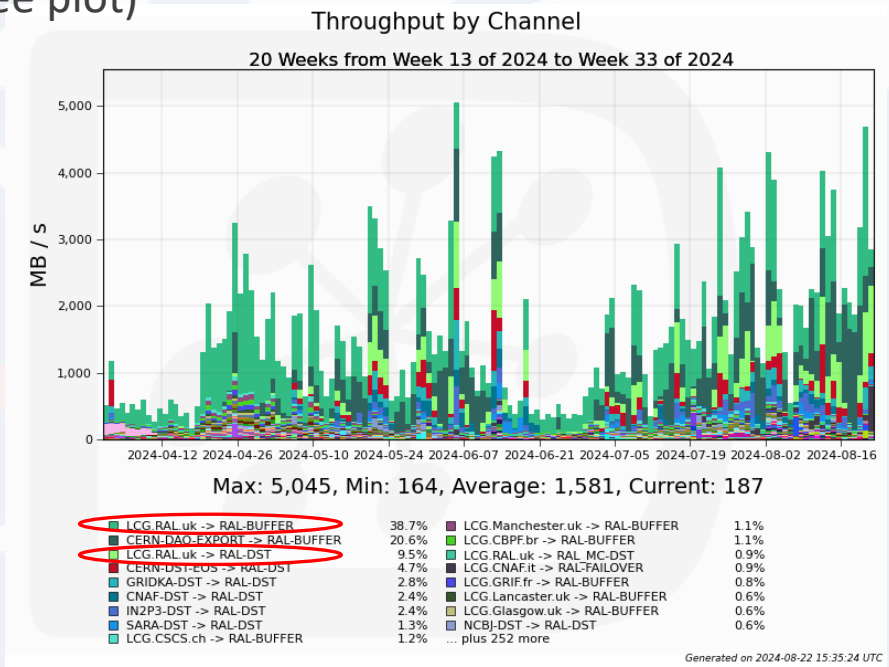

# LHCb News: future plans

- There is a significant increase in computing resources requests for the next FY [1]
  - 62% increase for CPU, 57% for Disk, 46% for Tape (relative to this FY)
  - During the LS3 there expected to be no increase at all
  - After this increase further resource requirement should stay within “Flat Cache” budget model
- Possible extension of the Run 3 introduces some uncertainty..
- Two new Tier-1 centers (Beijing and NCBJ) will be providing resources as well
- Network capacity seems to be already good enough
  - Should stay the same order of Magnitude as during DC24 until the end of run 4
- ARM usage is not ready for production yet

[1]https://indico.cern.ch/event/1389300/contributions/5840870/attachments/2833421/4950927/CERN-RRB-2024-012.pdf

# LHCb News: ETF tests

- New set of tests is used in production since May (see here and here)
  - Finally, includes storage tests!
- There are still a few issues, namely
  - ARC client bug is causing intermittent failures for CE tests
  - Container that is used for running test infrastructure is still based on CentOS7 image
    - We can not do much, since base container is maintained separately
    - All Vos are affected
- Future development plans:
  - Fix issues!
  - Add token-based SE tests (for me)
    - Currently tokens are only used for job submissions to HTCondor CEs

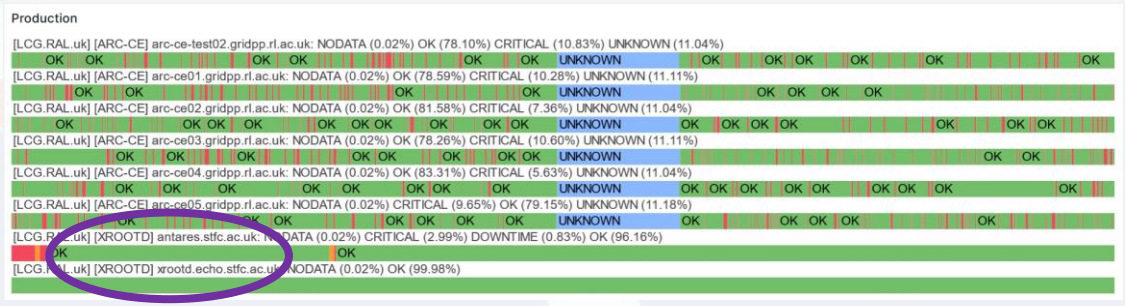

# LHCb News: UK DC

- After the Data Taking, Reprocessing Campaign starts
  - Lasts almost until the next year's DT starts
- UK-wide DC is likely to overlap with LHCb's reprocessing campaign
  - Tape usage is undesirable
    - Since it will interfere with reprocessing staging
  - Disk usage is OK as long as
    - We have enough space to accommodate test and production data
    - Testing does not interfere with production activities

# Conclusion

- UK provided a lot of resources to LHCb during the last 5 months
  - Around 30%
- Relatively smooth operations
- Long lasting issues are (almost completely) resolved
- New developments are ongoing

# RAL T1 Job success rate (backup)

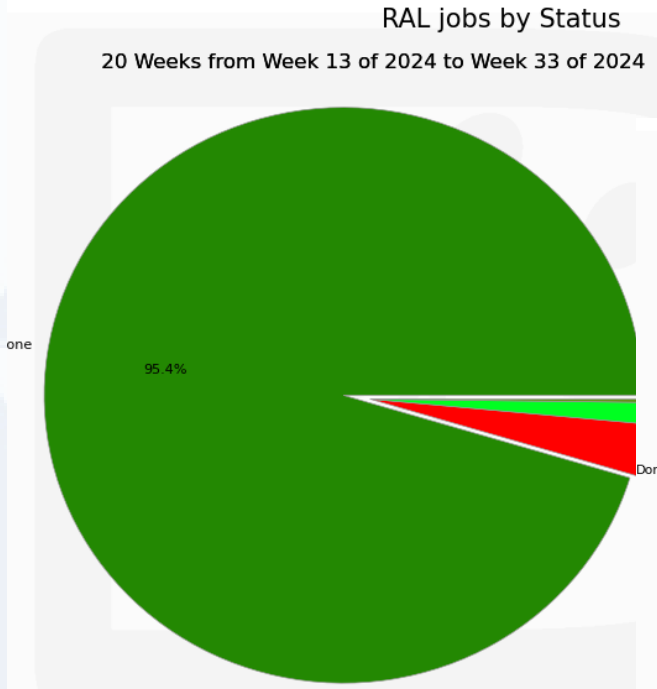


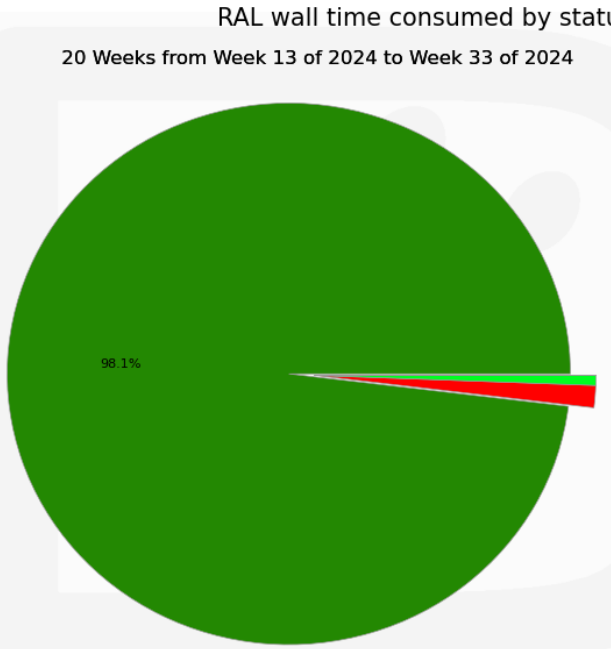


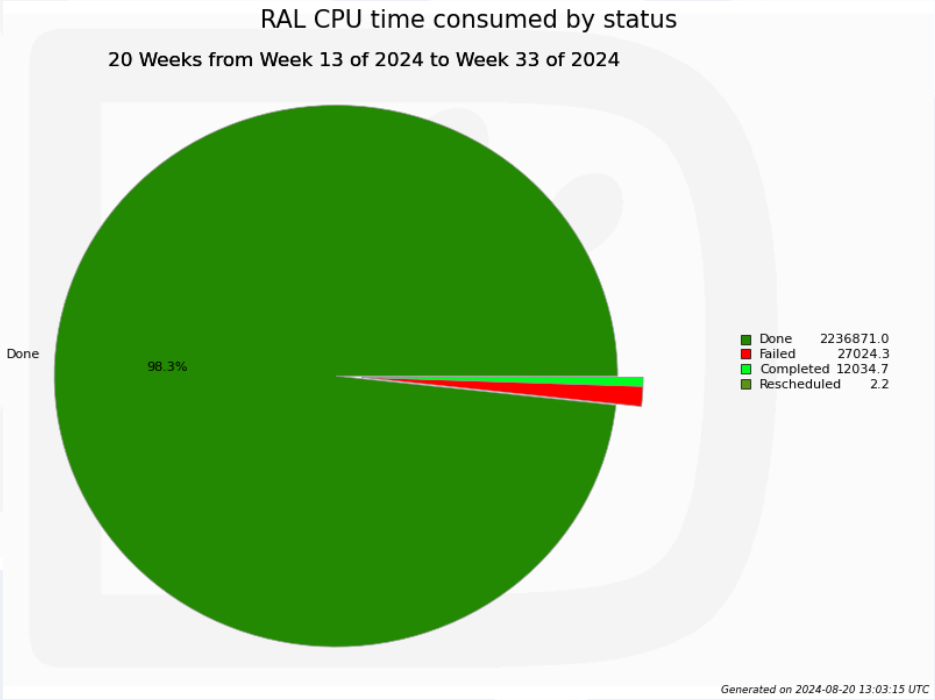

# RAL T1 CPU efficiency (backup)

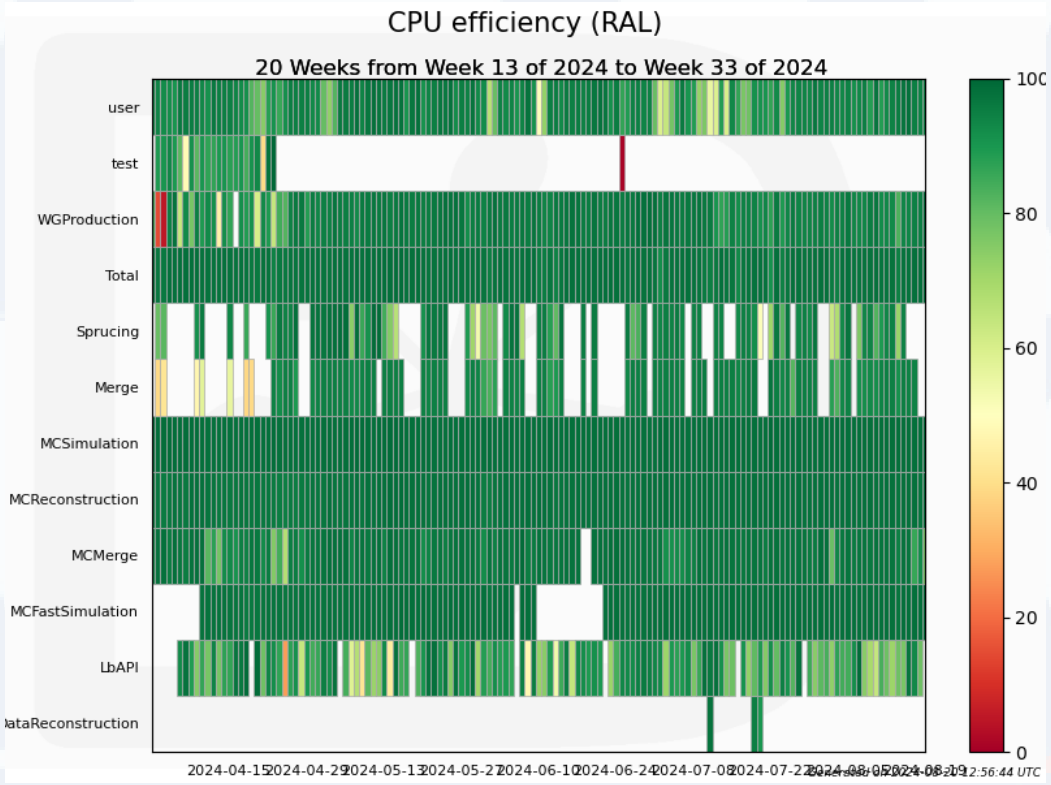


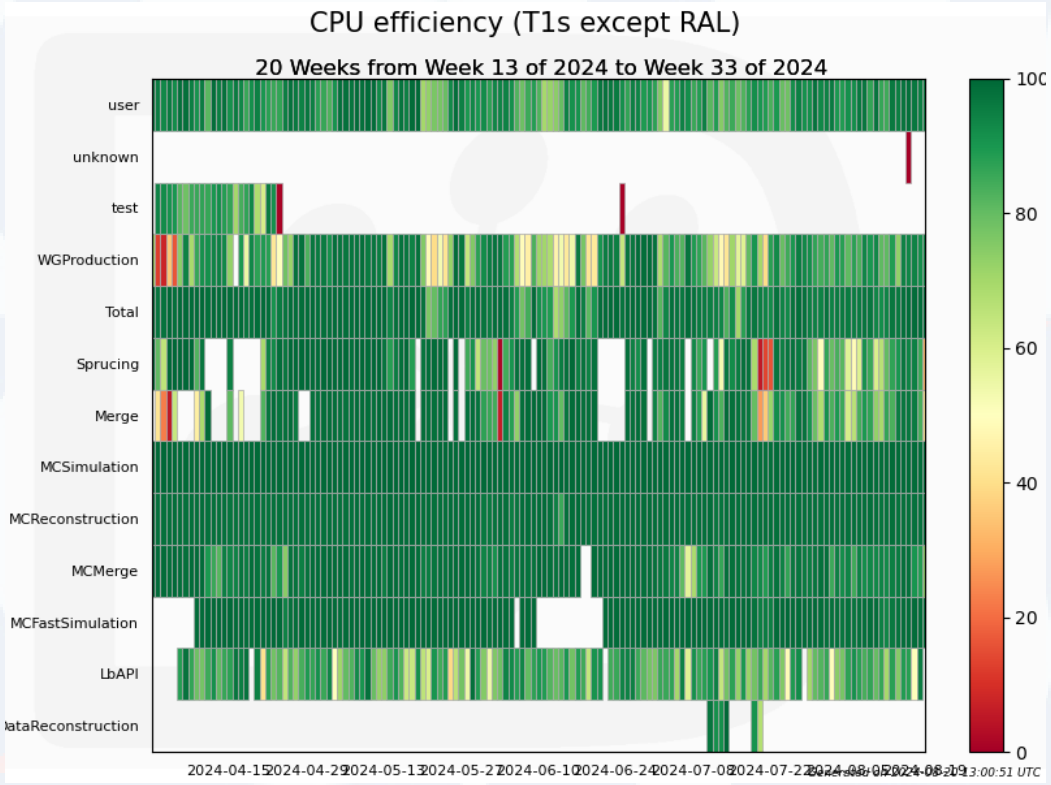

# Transfers to ECHO (backup)

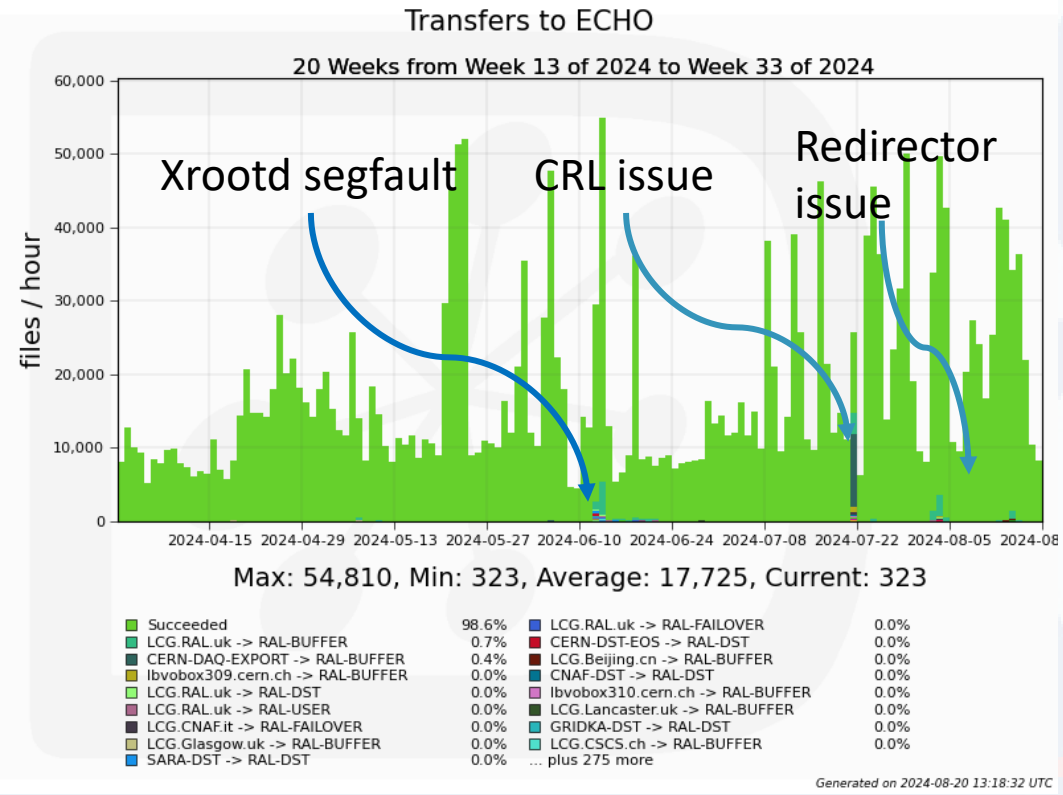


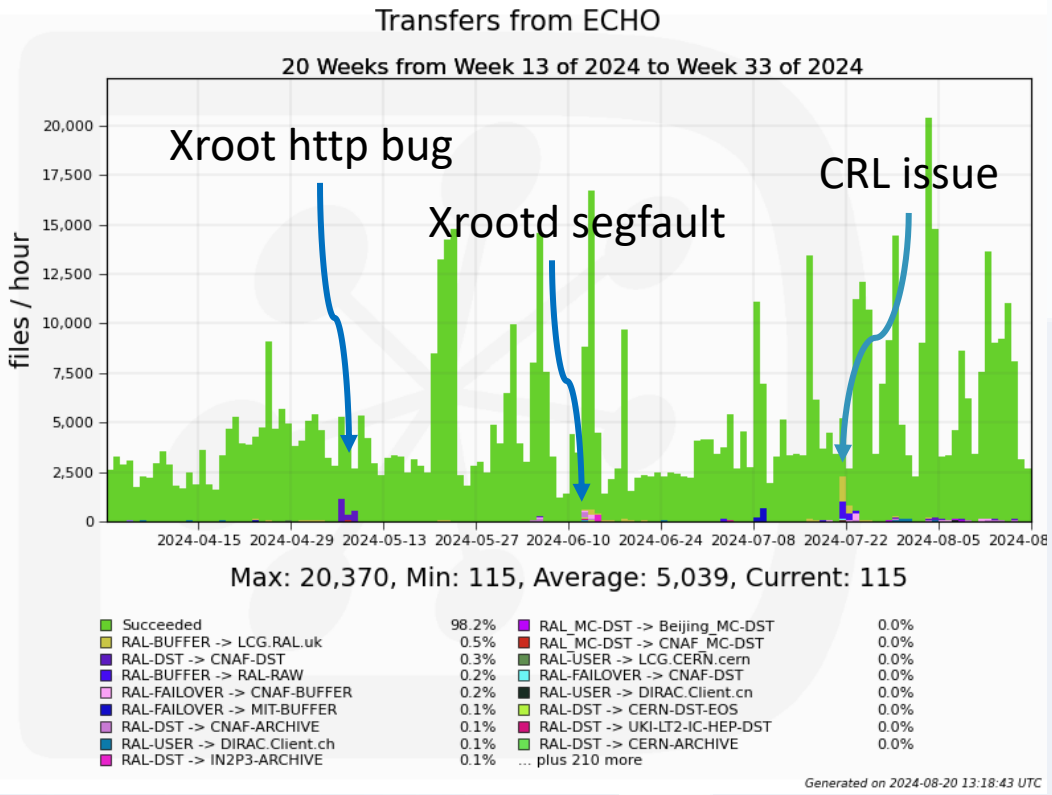

# Transfers to Antares (backup)