

**RAL Tier1 Resource Review Meeting**  
**29 Jan 2025**

**Tier-1 Experiment delivery sign off: ATLAS**  
**01 Oct to 31 Nov 2024 (2024Q4)**

Brij Kishor Jashal, Jyoti Prakash Biswal  
Rutherford Appleton Laboratory





- Previous quarter reports [2024Q1](#), [2024Q2](#) and [2024Q3](#)
- Key numbers for 2024Q4
  - WLCG Availability / Reliability – **99.27%**
  - Delivered compute HS23 vs Pledge: **80%**
    - Lower as compared to the last quarter of 96%
  - Walltime efficiency of jobs success/all - **93%**
  - Transfer efficiency [96.6%](#)
  - ATLAS storage occupancy: **96 PB** [Link](#)

VO atlas Profile ATLAS\_CRITICAL Tier All Country UK Federation All Site RAL-LCG2 Flavour All Hostname All Recomputations status More Dashboards

Website

This dashboard shows the availability, reliability, and status computation from ETF tests. To use it:  
1. Apply the desired selection using the filters on top (dashboard will also load faster).  
2. Click on "Availability & Reliability", "Site Status", "Endpoint Status", or "Test Status" to see results. Scroll up/down for more results.


### Availability & Reliability



# Storage: Occupancy

[Link]

Binning 1h Group by jobs dst\_experiment\_site Group By DDM dst\_experiment\_site Cloud UK Country All Federation All Site RAL-LCG2 Panda Queue All Job type All






Name	Mean	Last *	Max	Min
RAL-LCG2-ECHO_DATADISK	25.0 PB	25.2 PB	25.2 PB	24.6 PB
RAL-LCG2-ECHO_SCRATCHDISK	195 TB	199 TB	209 TB	165 TB
RAL-LCG2-ECHO_TEST	3.54 TB	3.58 TB	3.58 TB	3.52 TB
RAL-LCG2_DATATAPE	29.3 PB	30.0 PB	30.0 PB	27.9 PB
RAL-LCG2_MCTAPE	39.6 PB	40.5 PB	40.5 PB	38.9 PB

# Storage: pledge vs occupancy

~96 PB total storage

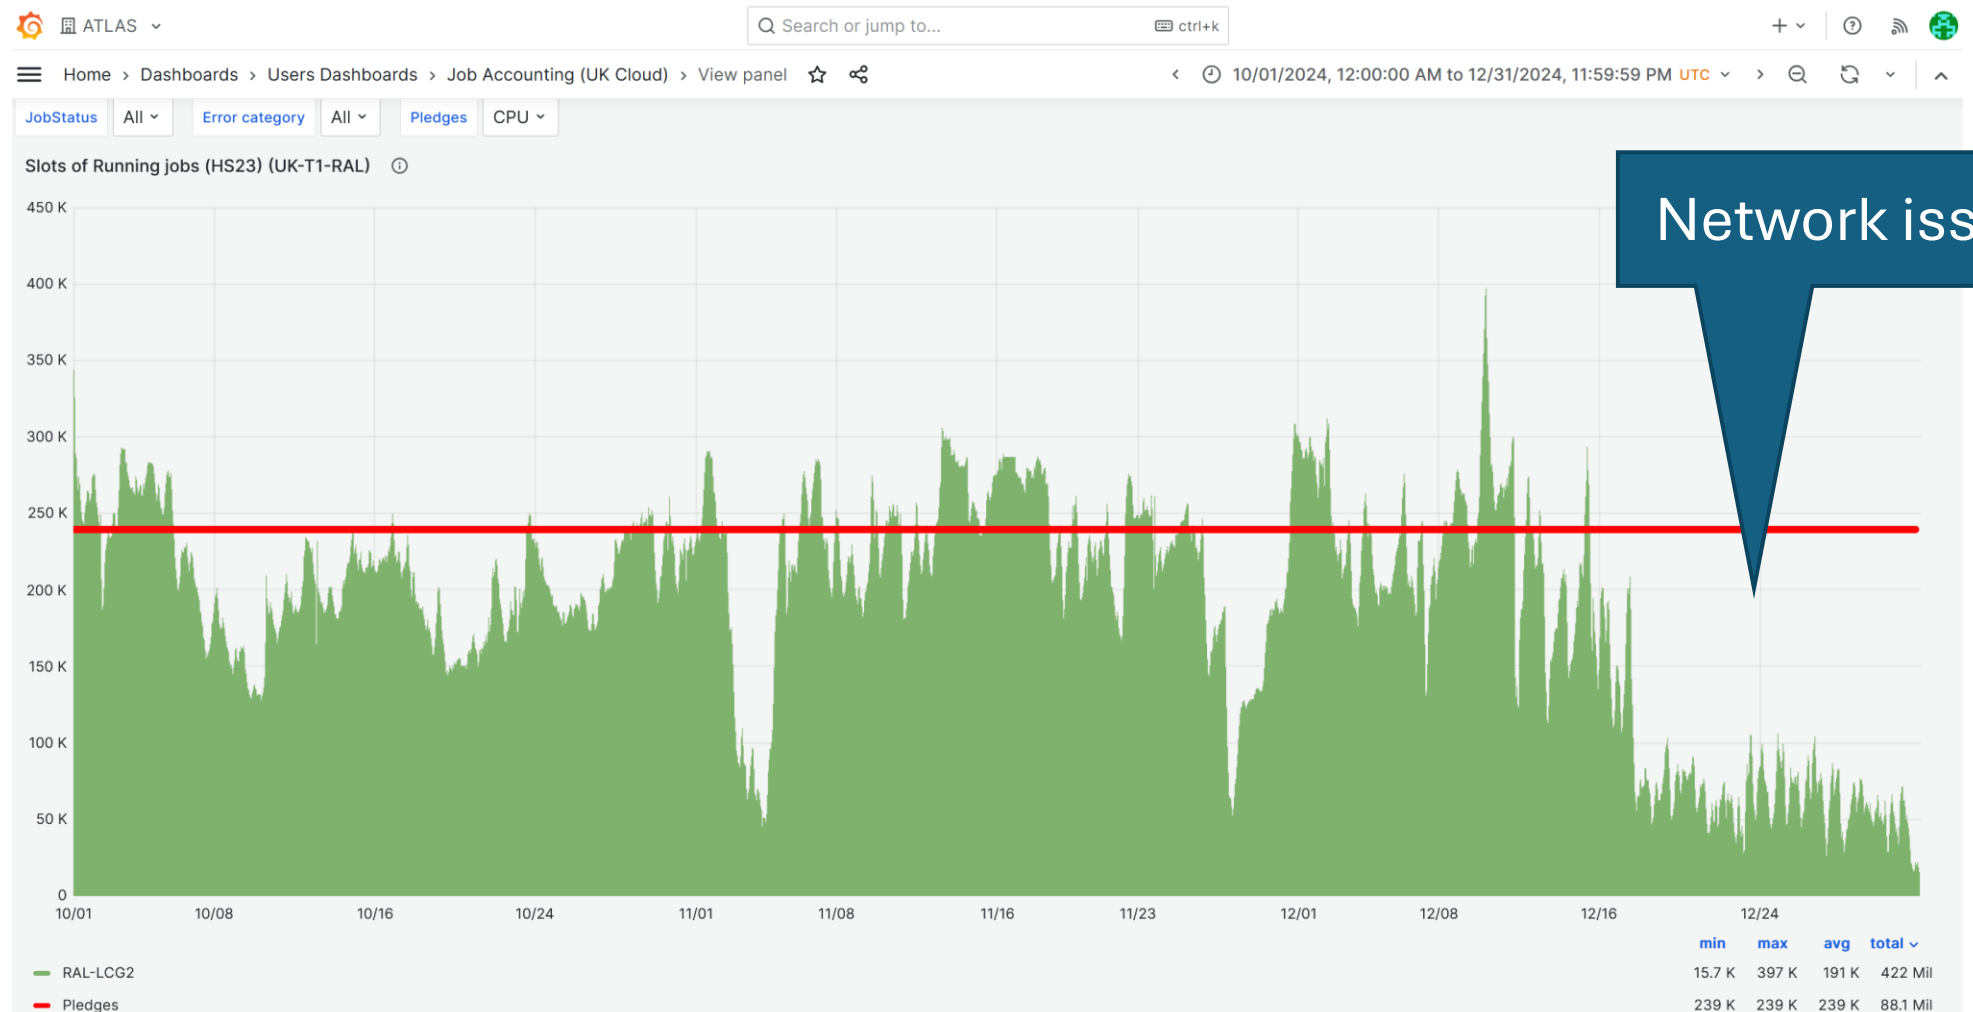
- 25.5 PB DISK
- 70.5 PB Tape
  - 30 DataTape
  - 40.5 MCTape



# Compute: RAL share among ATLAS T1s

[[Dashboard link](#)]

- Delivered compute HS23 vs Pledge: **80%**
- Main reason for under-delivery was network issue, documented in this ticket [169483](#) and weekly reports.
  - This was traced down to a fault with a transceiver on one of RAL's internal routers, specifically ral-int-rtr-02



# Compute: RAL share among ATLAS T1s

[\[Dashboard link\]](#)

Home > Dashboards > Jobs Monitoring > Job Accounting (Historical data) > View panel ☆ 🔗


10/01/2024, 12:00:00 AM to 12/31/2024, 11:59:59 PM UTC 🔍 ↺ ↻

Binning 1h Group by atlas\_site Country All Fed All Resource Type All Tier 1 Cloud All Site All Panda Queue All Nucleus All Cores All

Event Serv All Groups All Input Data All Input Proj All Output Proj All GlobalShare All Resources Reporting All ProcessingType All JobType All

Prodsourcelabel All JobStatus All Error category All Container name All Job resource type All Generator All CPU architecture All CPU type All

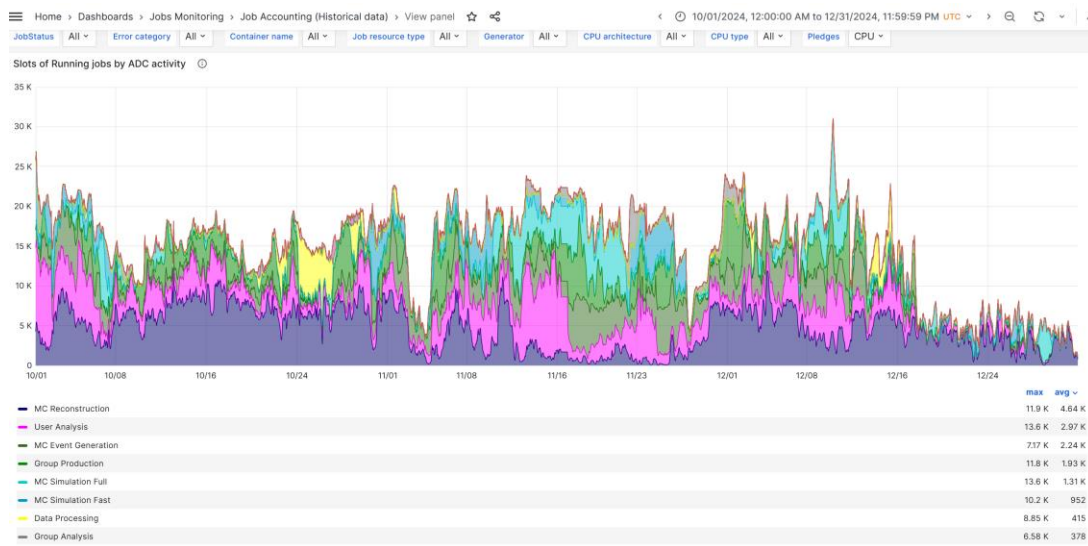
Pledges CPU



12% of ATLAS Tier1 compute  
2<sup>nd</sup> Rank

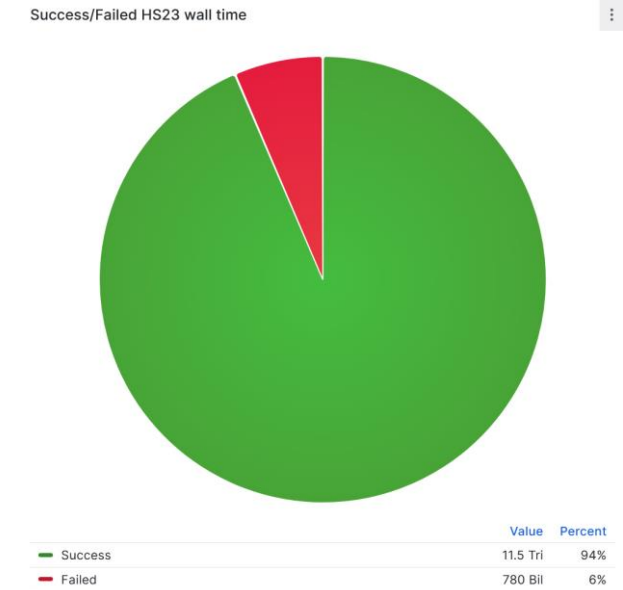
	Value	Percent
BNL-ATLAS	8.20 Mil	30%
RAL-LCG2	3.30 Mil	12%
IN2P3-CC	3.25 Mil	12%
FZK-LCG2	2.49 Mil	9%
TRIUMF-LCG2	2.23 Mil	8%
NDGF-T1	1.91 Mil	7%
INFN-T1	1.73 Mil	6%
pic	1.68 Mil	6%
SARA-MATRIX	1.59 Mil	6%

## Slot occupancy and different types of workflows



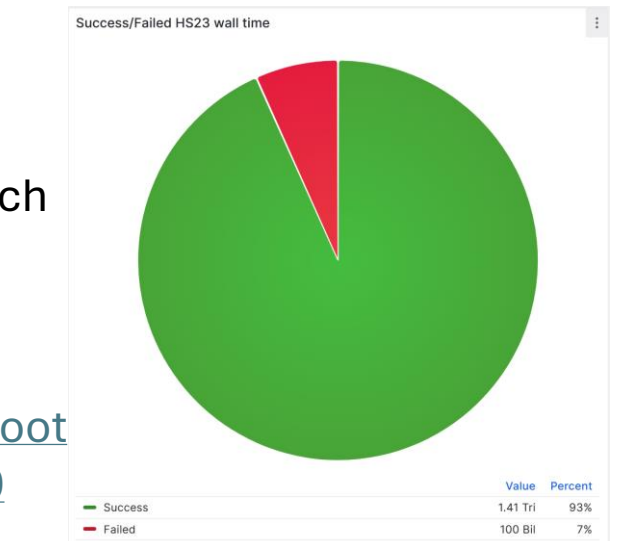
Average ATLAS Tier-1 job failure HS23 walltime at 6%

Average RAL Tier-1 job failure HS23 walltime at 7%



Two main reasons for this



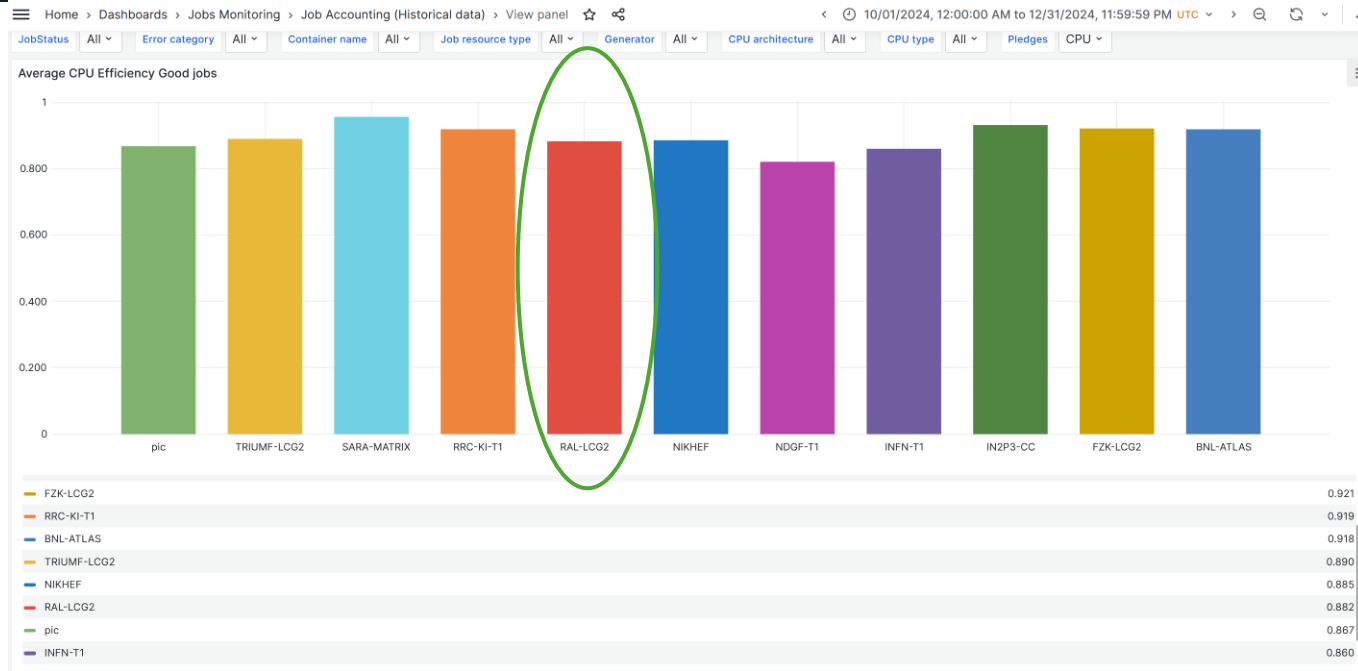
- Network issue [169483](#)
- Xrootd bug due to mismatch of client/server versions affecting Event Index Jobs
  - Ticket – [169427](#)
  - <https://github.com/xrootd/xrootd/issues/2410>





# Compute: Walltime efficiency

[Dashboard link]



# Data transfers: Efficiency


[\[Dashboard link\]](#)

Home > Dashboards > Site monitoring > Site-oriented dashboard > View panel ☆ 🔗

< 🕒 10/01/2024, 12:00:00 AM to 12/31/2024, 11:59:59 PM UTC > 🔍 ↻ ⌵ ⌶

Binning 1h Group by jobs dst\_experiment\_site Group By DDM dst\_experiment\_site Cloud UK Country All Federation All Site RAL-LCG2 Panda Queue All Job type All

## Transfer Efficiency




**Good data transfers efficiency overall at 96.6%**

Name	Mean	Last *	Max	Min
RAL-LCG2	96.6%	85.6%	99.9%	11.7%


Transfer throughput with Destination as RAL\_LCG2




Cumulative transfer volume with RAL\_LCG2 as destination ~70PB



Transfer throughput with Source as RAL\_LCG2



Cumulative transfer volume with RAL\_LCG2 as Source ~98 PB



# Summary

- Compute resource delivery vs pledge lower at 80% affected due to network issue.
  - Significantly longer duration to identify and rectify the issue due to multiple reasons
    - Complicated issue which was not consistent.
    - Initially thought to be a problem with Batch system due to token transition overlapping timeline and symptoms
    - Holiday season.
- Token transition was rolled back. ATLAS is not submitting any production jobs anyways. This will be rolled back after few weeks of stability or with new ArcCE7s
- Apart from the above, remaining all key critical parameters of performance are good and satisfactory
- ATLAS has been requested to prioritize RAL\_LCG2 for consistent load of workflows and we are seeing good occupancy during past weeks.
- More details of operations are available here in my weekly reports during this quarter
  - Until Dec2024 <https://codimd.web.cern.ch/s/sUxhgPxC8#>
  - From Jan 2025 <https://codimd.web.cern.ch/s/QYOGneO7E#>

**Thank you**