

RAL Tier1 Resource Review Meeting
16th Oct 2024

Tier-1 Experiment delivery sign off: ATLAS
01 July to 30 Sep 2024 (2024Q3)

Brij Kishor Jashal, Jyoti Prakash Biswal
Rutherford Appleton Laboratory




- Previous quarter reports [2024Q1](#) and [2024Q2](#)
- Key numbers for 2024Q2
 - WLCG Availability / Reliability – **98.23%**
 - Delivered compute HS23 vs Pledge: [96%](#)
 - Much higher as compared to the last quarter of 86%
 - Walltime efficiency of jobs success/all - [95%](#)
 - Transfer efficiency [96%](#)
 - ATLAS storage occupancy: **95.5 PB** [Link](#)
 - DATATAPE Occupancy at **82%** [Link](#)
 - MCTAPE(TB) Occupancy at **82%** [Link](#)
 - ECHO_DATADISK(TB) Occupancy at **98%** [Link](#)

WLCG Availability / Reliability


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

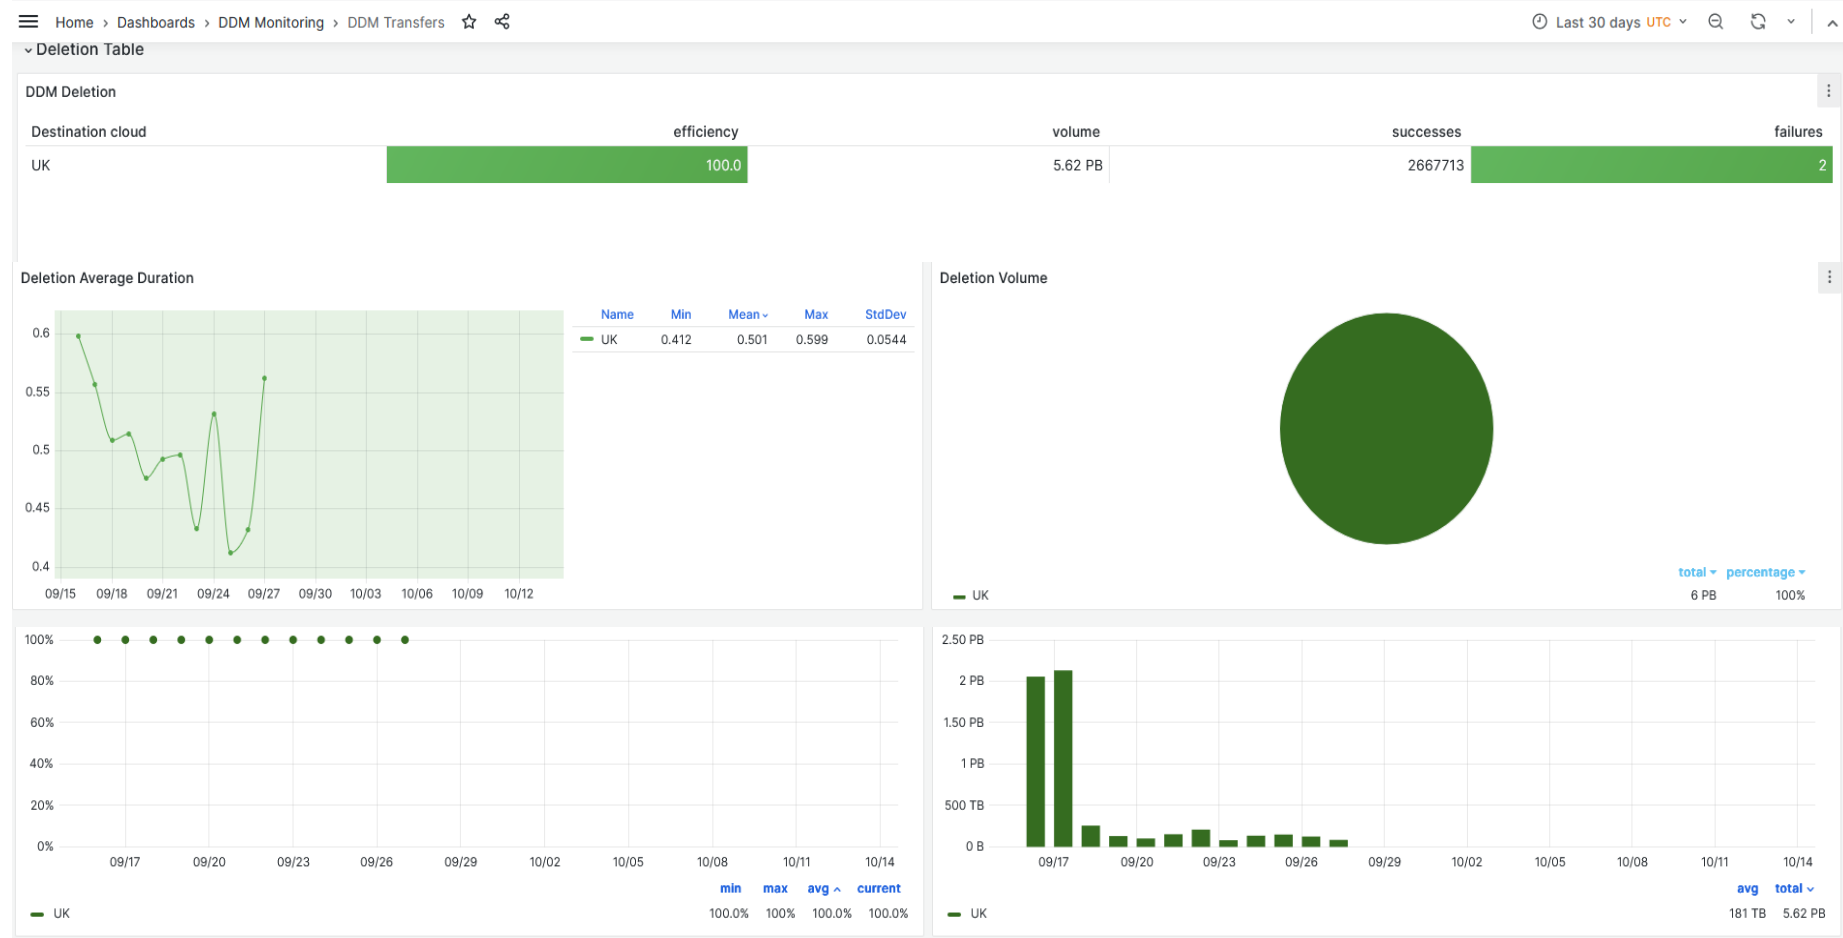



Site Status



News items from the reporting quarter (selected few)


- Deletion campaign, 6 PB deleted from RAL-Tape [Link](#)




Improved deletion performance overall

- After thorough deletion performance analysis, identified CRIC misconfiguration with old gateway endpoint.
- After implementing the correction, the deletion performance metrics of ATLAS look much better.


Deletion Average Duration




Deletion Failures



Deletion Efficiency




Deletion Volume

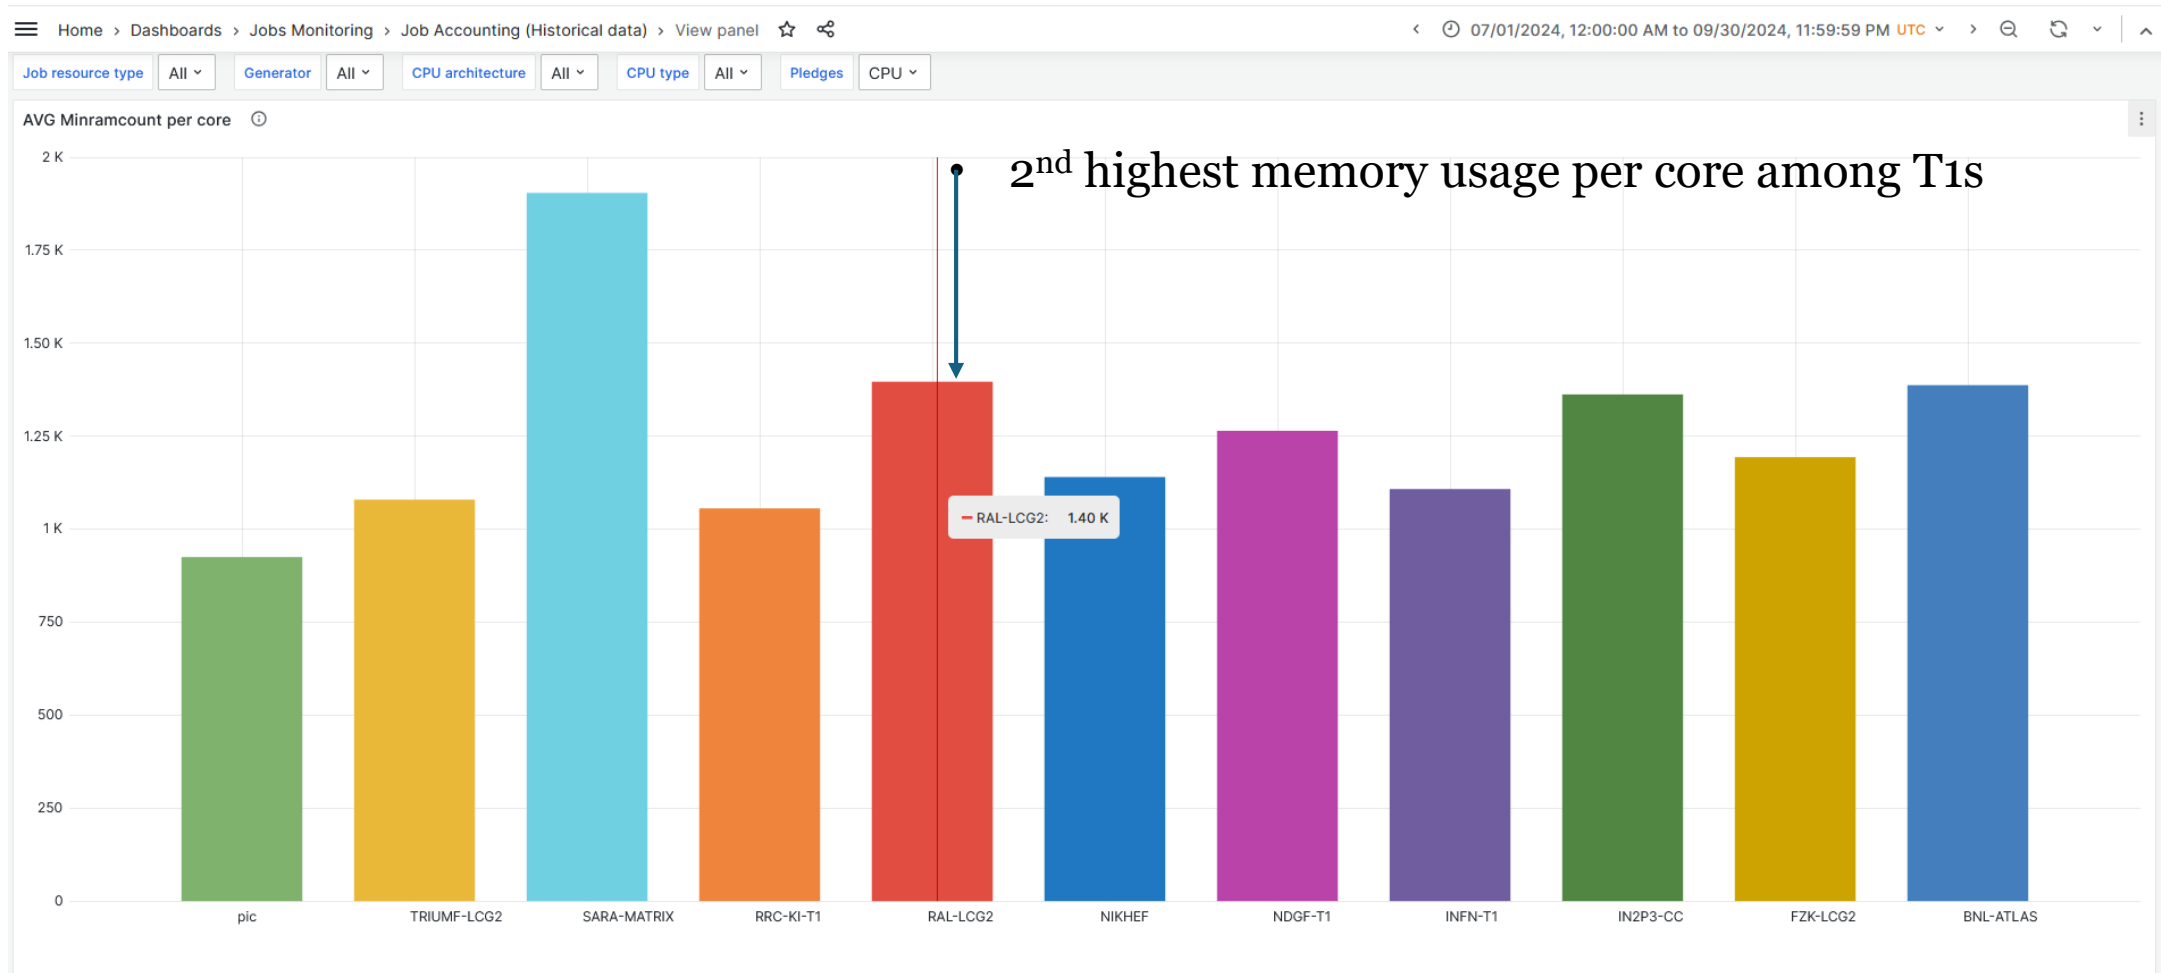


News items from the reporting quarter (selected few)


- For increasing the overall memory utilization from high-memory jobs, updated the MaxRSS (from 3000 to 6000) and MeanRSS (2000 to 4000) values at RAL Panda CRIC endpoints)
- Resulting into higher average memory utilization



- 2nd highest memory usage per core among T1s



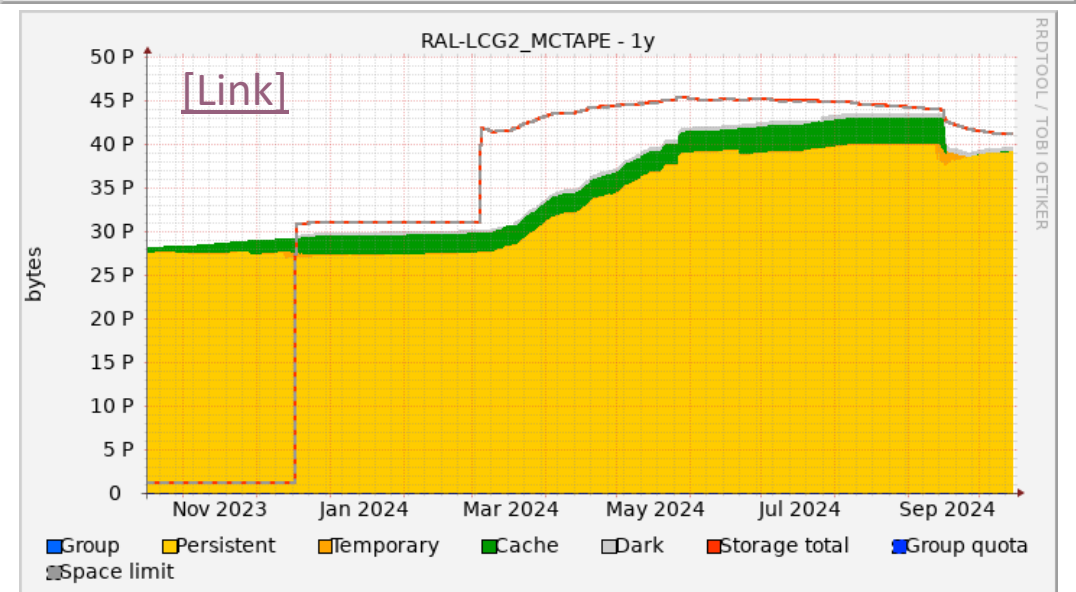


Storage: Occupancy




Storage: pledge vs occupancy

~95.5 PB total storage

- 28.9 PB DISK
 - 25.2PB – Datadisk
 - 18PB persistent
 - 7.2 PB Cache and temp
 - 3.5PB - Scratchdisk
- 66.6 PB Tape
 - 27.8 DataTape
 - 38.8 MCTape




Wall clock time. All jobs (HS23 seconds)





	Value	Percent
BNL-ATLAS	3.20 Tri	22%
RAL-LCG2	1.81 Tri	13%
IN2P3-CC	1.74 Tri	12%
FZK-LCG2	1.53 Tri	11%
pic	1.34 Tri	9%
TRIUMF-LCG2	1.34 Tri	9%
NDGF-T1	1.22 Tri	9%
INFN-T1	1.00 Tri	7%
SARA-MATRIX	40.3 Bil	3%

- Delivered compute HS23 vs Pledge: **96%**
- RAL T1 delivered above pledge, the reason for the delivered resource number being not above 100% is the underutilization due to non availability of workflows from ATLAS





Slots of Running jobs by ADC activity

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



Slots of Running jobs by ADC activity


Average RAL Tier-1 job failure HS23 walltime at 4%




Slot occupancy and different types of workflows

Wallclocktime successful/all

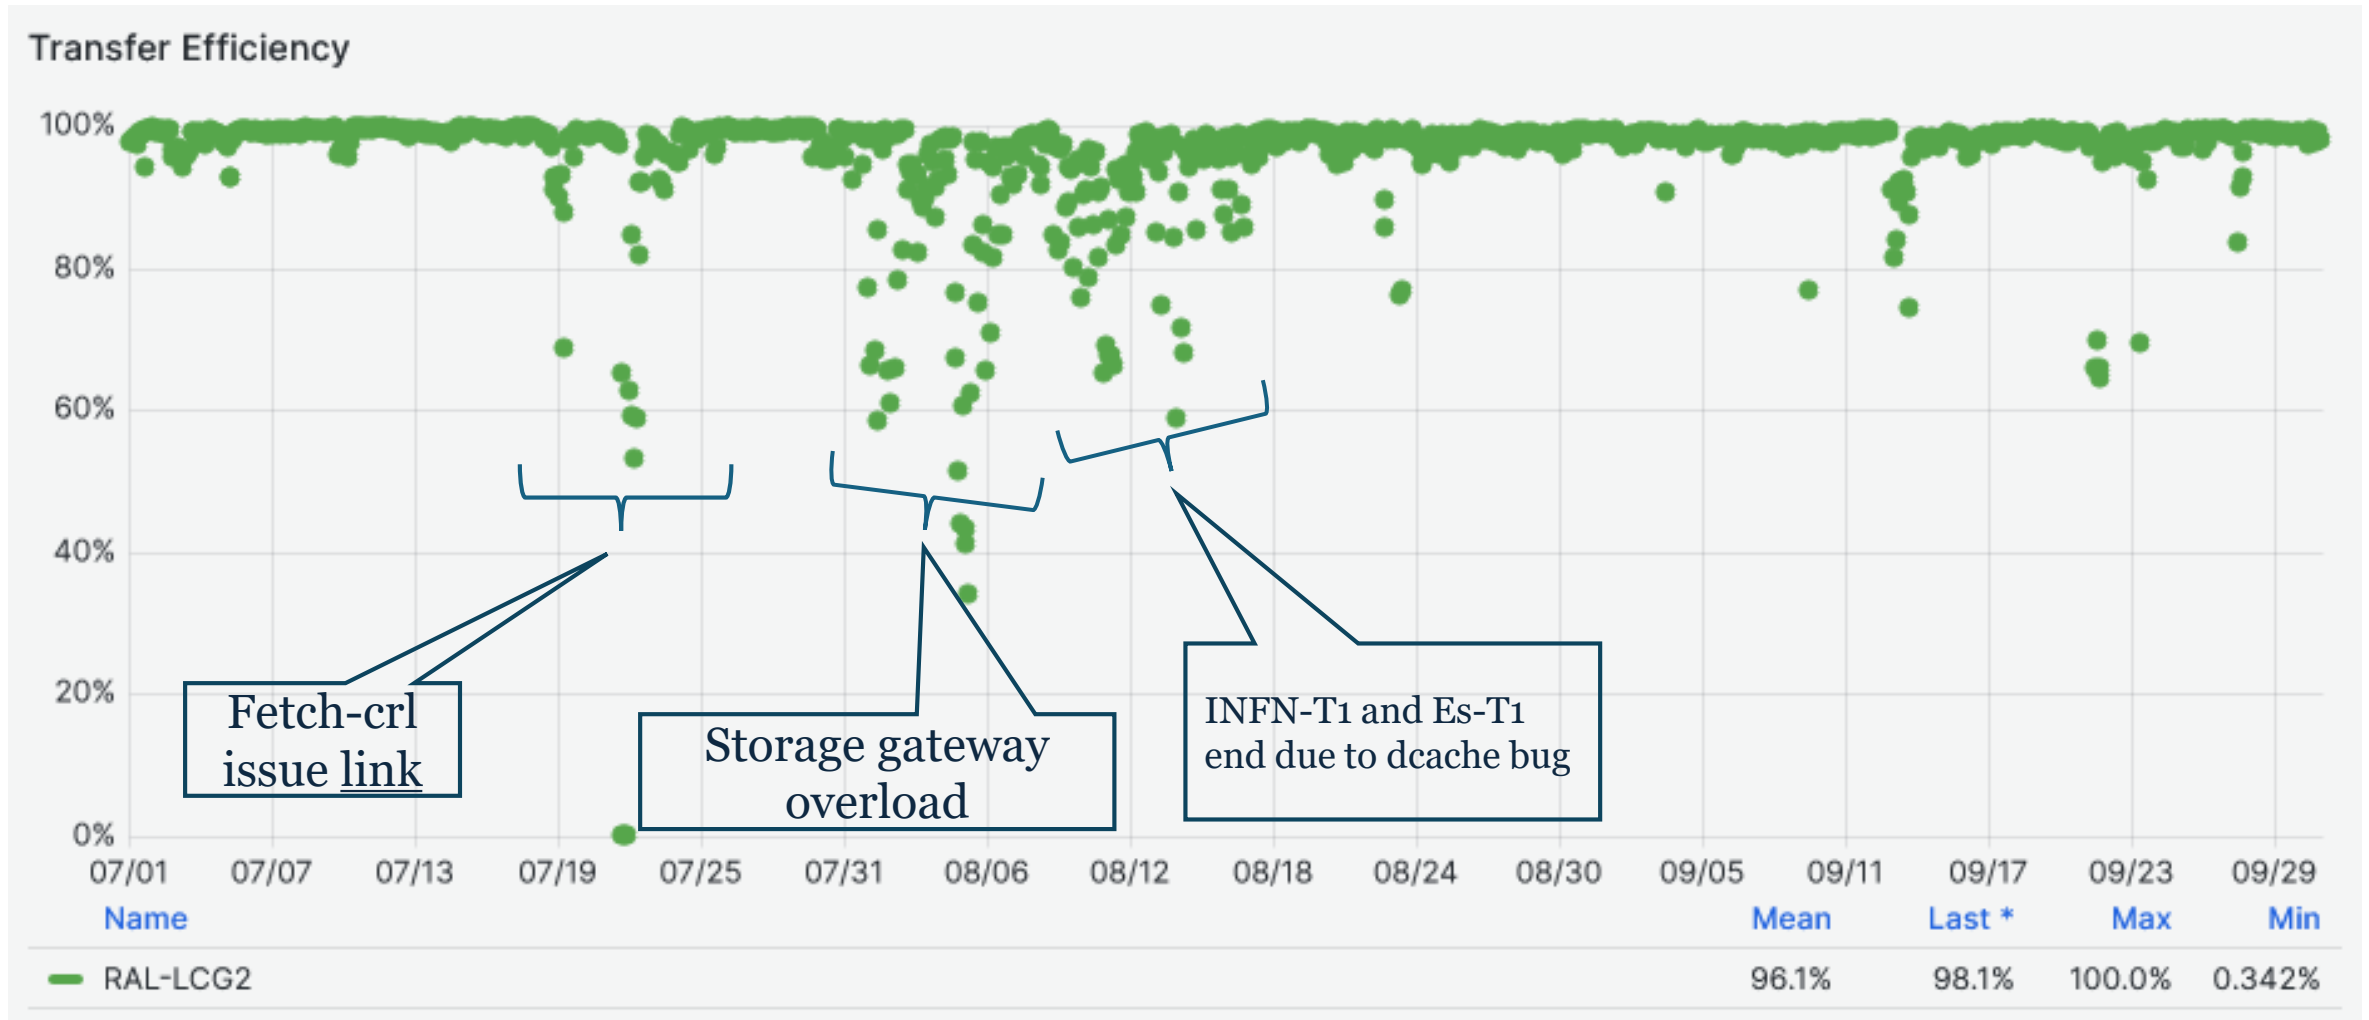
- CPU wall time Successful vs all at ~95%





Data transfers: Efficiency


Good data transfers efficiency overall at 96%



Data transfers: Transfers from RAL-LCG2 (Uploads)



[\[Dashboard link\]](#)

Cumulative 70 PB with source as RAL-LCG2



Data transfers: Transfers to RAL-LCG2 (Downloads)

Cumulative 85 PB with destination as RAL-LCG2



- Compute resource delivery vs pledge significantly improved as compared to the previous quarter (96% current vs 86%), thanks to tweaks in scheduling policy favoring multicore jobs
- Deletion performance very satisfactory after changes to the CRIC gateway.
- Among first T1s to implement HIMEM configuration, before getting any ggus tickets
- Delivered resources with high efficiency and throughput.
- Planned Token transition during next quarter.

Some noteworthy points from ATLAS

- PHYS vs PHYSLITE - Good efforts and progress in trying to reduce production and analysis size and compute cost. Will result in reduction of resource projection for Run4 - [Link](#)

Date	Production	Size/event (kB)	
		PHYS	PHYSLITE
	TARGET	30-50	10-12
09/2023	24.0.12	28.2	11.2
01/2024	24.0.21	25.6	9.14
06/2024	25.0.8	27.9	11.7


- SAM / Experiments Test Framework (ETF) migration to Arc-CE rest based interface under testing with the current production HTCondor framework

- TAPE REST deployment(ATLAS SRM retirement) SRM dead starting 2026,

TAPE recalls in a year

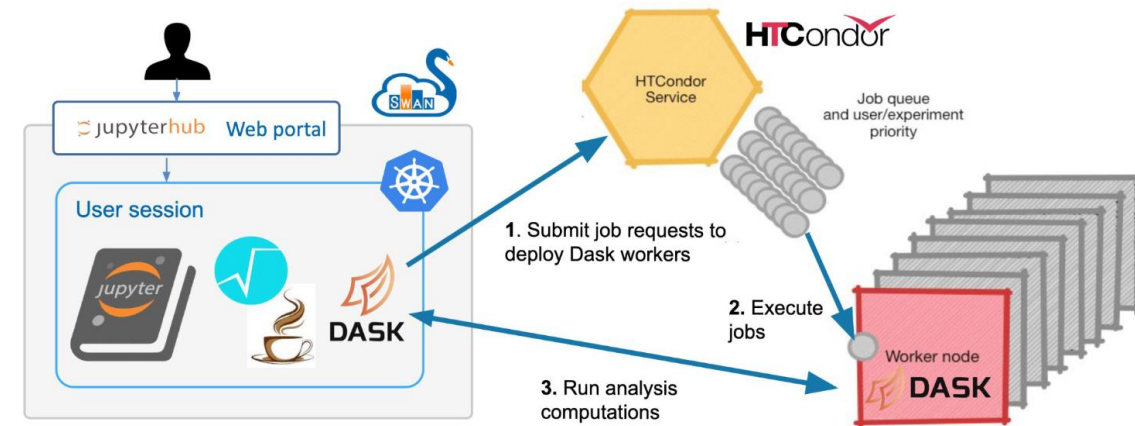
- o ATLAS – 187 PB
- o CMS – 51PB
- o LHCb – 19PB
- o BelleII – 1PB

Staging volume in last two years



Some noteworthy points from ATLAS

- New CERN Analysis Facility: integration with WLCG planned.
[Link](#)
- Grid facilities and services are evolving fast, building on Kubernetes based micro-architectures, several interesting talks at Oct A&S week
 - [ContainerPulling](#)
 - [RegistryHarbour](#)
 - [PanDAonKubernetes](#)
 - [Data carousel and centralized NTUP production](#)
- Hybrid architectures are the future, as a facility we should be ready with testbeds with ability to ramp-up the resources as per the demands from VOs and our users.



Thank you