U.S. DEPARTMENT OF ENERGY
Office of Science

BERKELEY LAB

ALICE-USA Computing Project Status Report

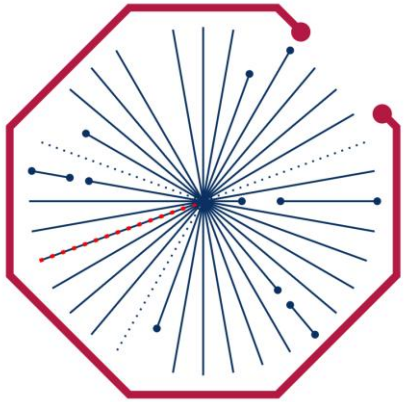

ALICE-USA Computing Annual Meeting @ LBNL

April 20, 2026

Irakli Chakaberia

ALICE

Brief History

- Original Project Proposal – 2009
- In operational since 2010
- 3-Year Project Review Cycle
- Initial plan requested in 2014
- Latest review in July 2021
- Potential review in 2026
- BUT, DOE is dealing with quite some changes
- I do not expect the review this year
- Project Execution & Acquisition Plan (PEAP)
  - Multi-year plan containing resource delivery, milestones, budget estimates, operations guide
  - updated annually & submitted to DOE
  - Latest one submitted in October 2025
- ALICE-USA Computing Website: https://sites.google.com/lbl.gov/alice-usa-computing/meetings?authuser=1

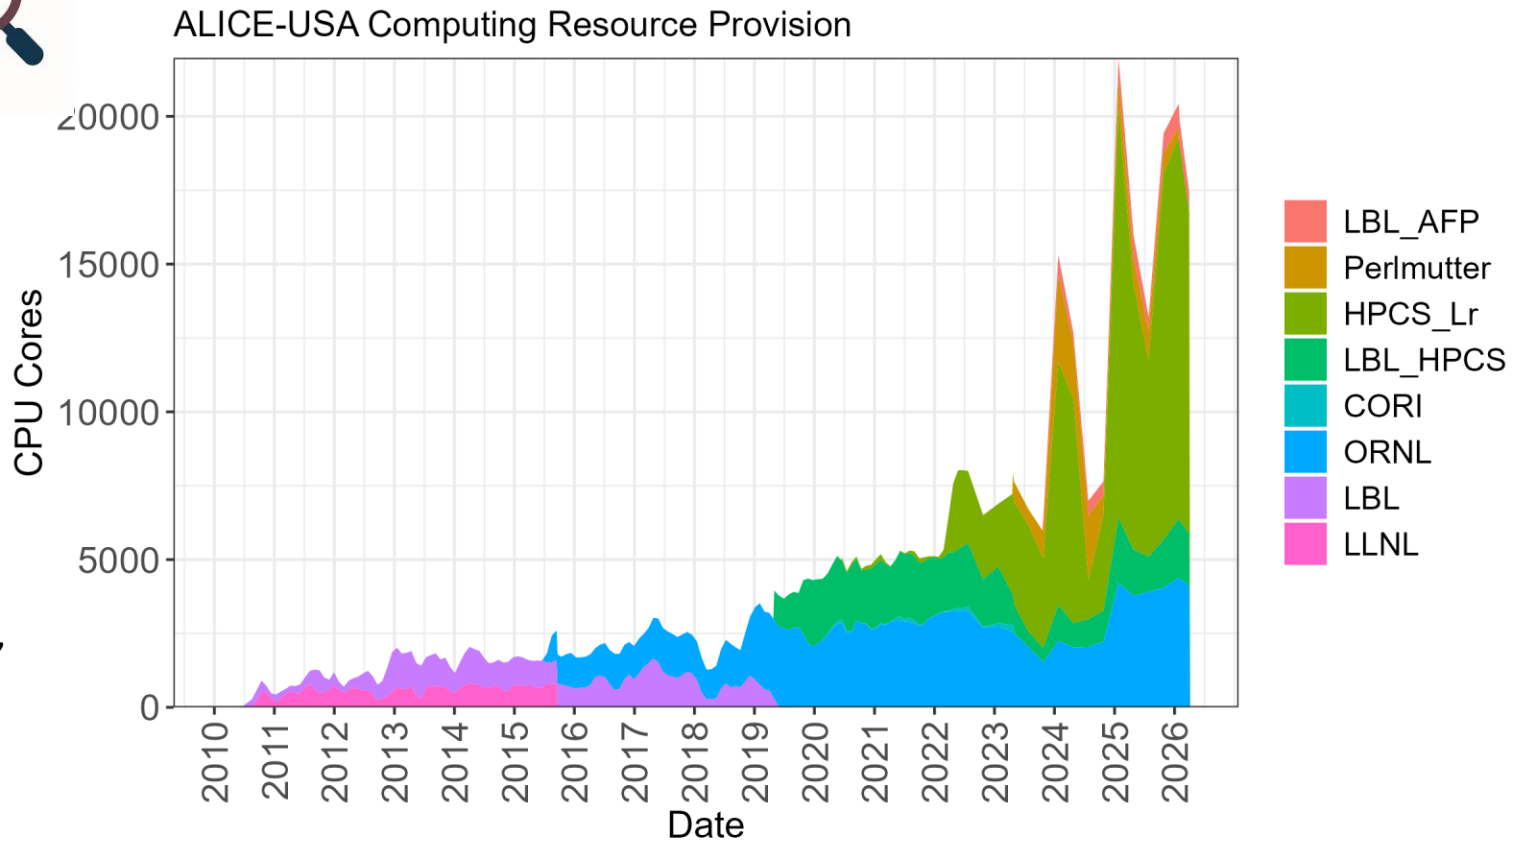

Personnel Change

Since our last meeting in September, the project has significant personnel change

- ORNL
  - Pete and Brian had left the project
  - While Steve has been an extremely reliable pillar, his time contribution is limited
  - The new position was open at ORNL Physics division to hire a replacement for Pete
  - The search is still underway…
- LBNL
  - John has left the lab
  - Big thanks:
    - To John for finalizing the AF EOS work before he left, and for training Karen and Torben and leaving the documentation
    - To ITD HPCS group and Wei in particular for proactively working with us when John left
    - To Karen, for picking up quite a bit of work that needed to be done when John left
    - To Torben, for being interested in this project and for sharing the work with Karen to have the transition as smooth as possible

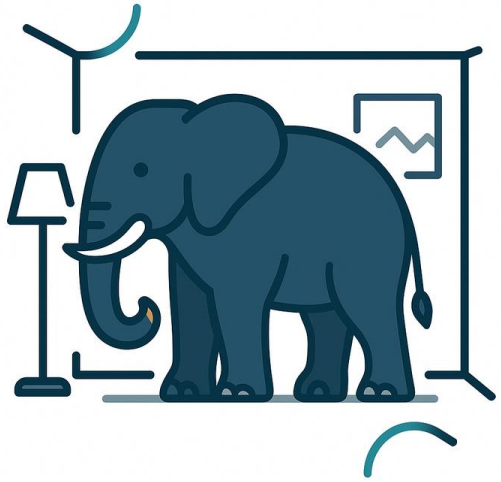

Other changes – Resource Delivery

- ALICE-USA remains strongly involved in ALICF
- However, we see some reduction of the M&O_A in 2027 which I expect to remain on that level
- ALICE-USA Computing share will be reduced from current 8.2% to 7.2%
- This reduction still results in a resource growth for in the future

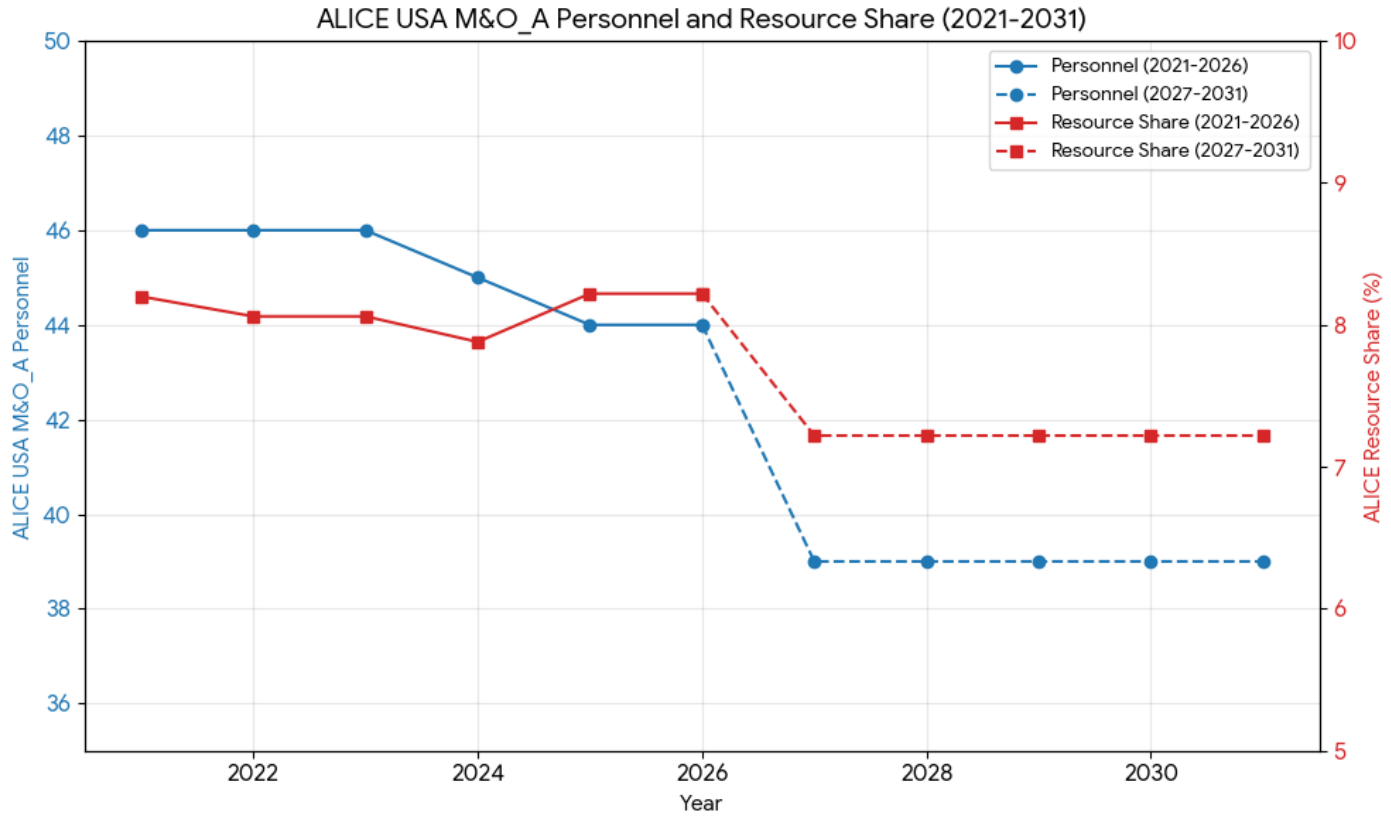

US Department of Energy Supported ALICE-USA Group

- The project's main goal is to fulfills DOE funded MoU-based ALICE USA obligations for
  - 11 Institutions
  - 44 M&OA, 39 in 2027
- Operates ALICE grid facilities at 2 DOE labs
- ALICE-USA Computing core group meets monthly, every third Tuesday of the Month
- We use slack for issues during the operations and other coordination
- **Nominally we try to hold two annual grid review meetings**
  - **Spring meeting at LBL**
  - **Fall meeting at ORNL – we might revisit the format of this meeting**

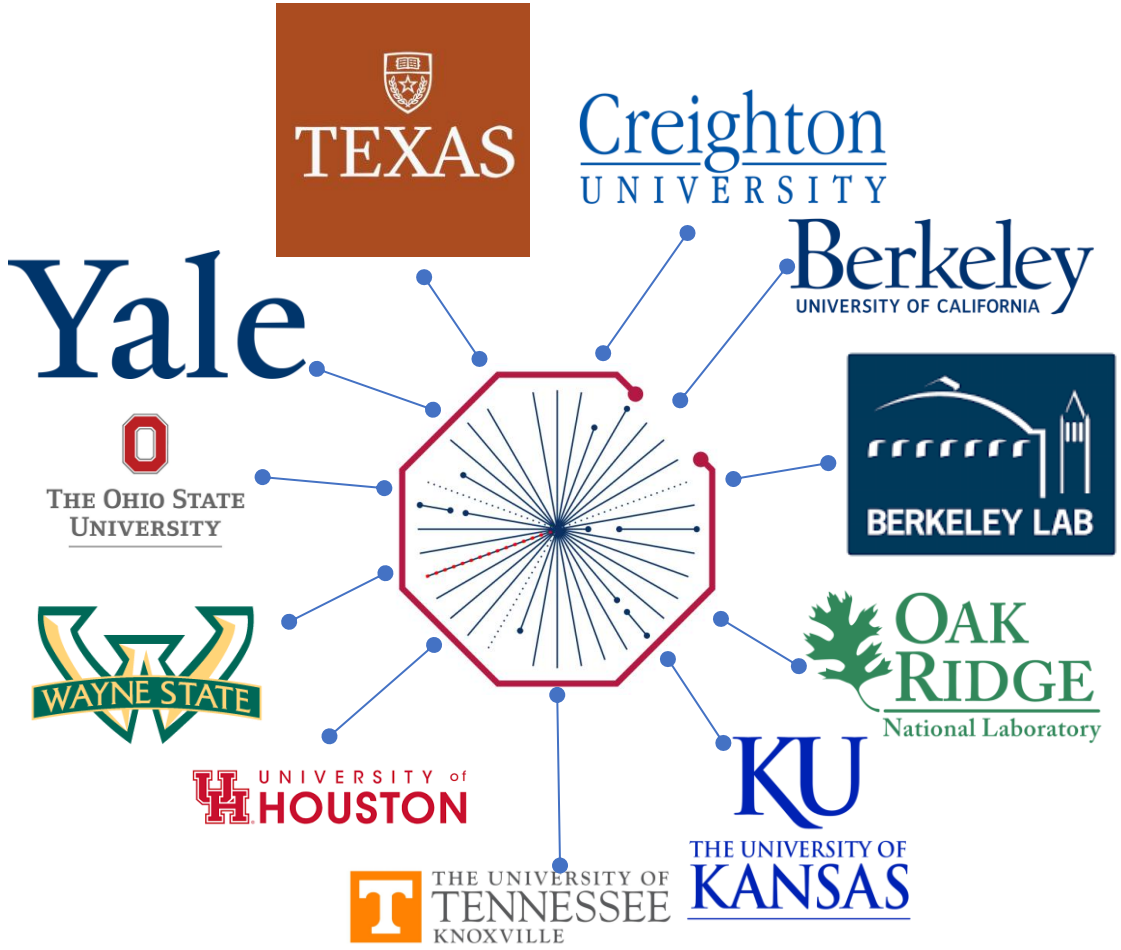

ALICE-USA T2 Sites & Analysis Facility (Prototype)

- Project continues to operate two T2 sites at ORNL and LBNL
- The analysis facility prototype was deployed two years ago and continues stable operations
- In addition, we provide resources on Lawrencium (opportunistic) and Perlmutter (semi-opportunistic) HPCs

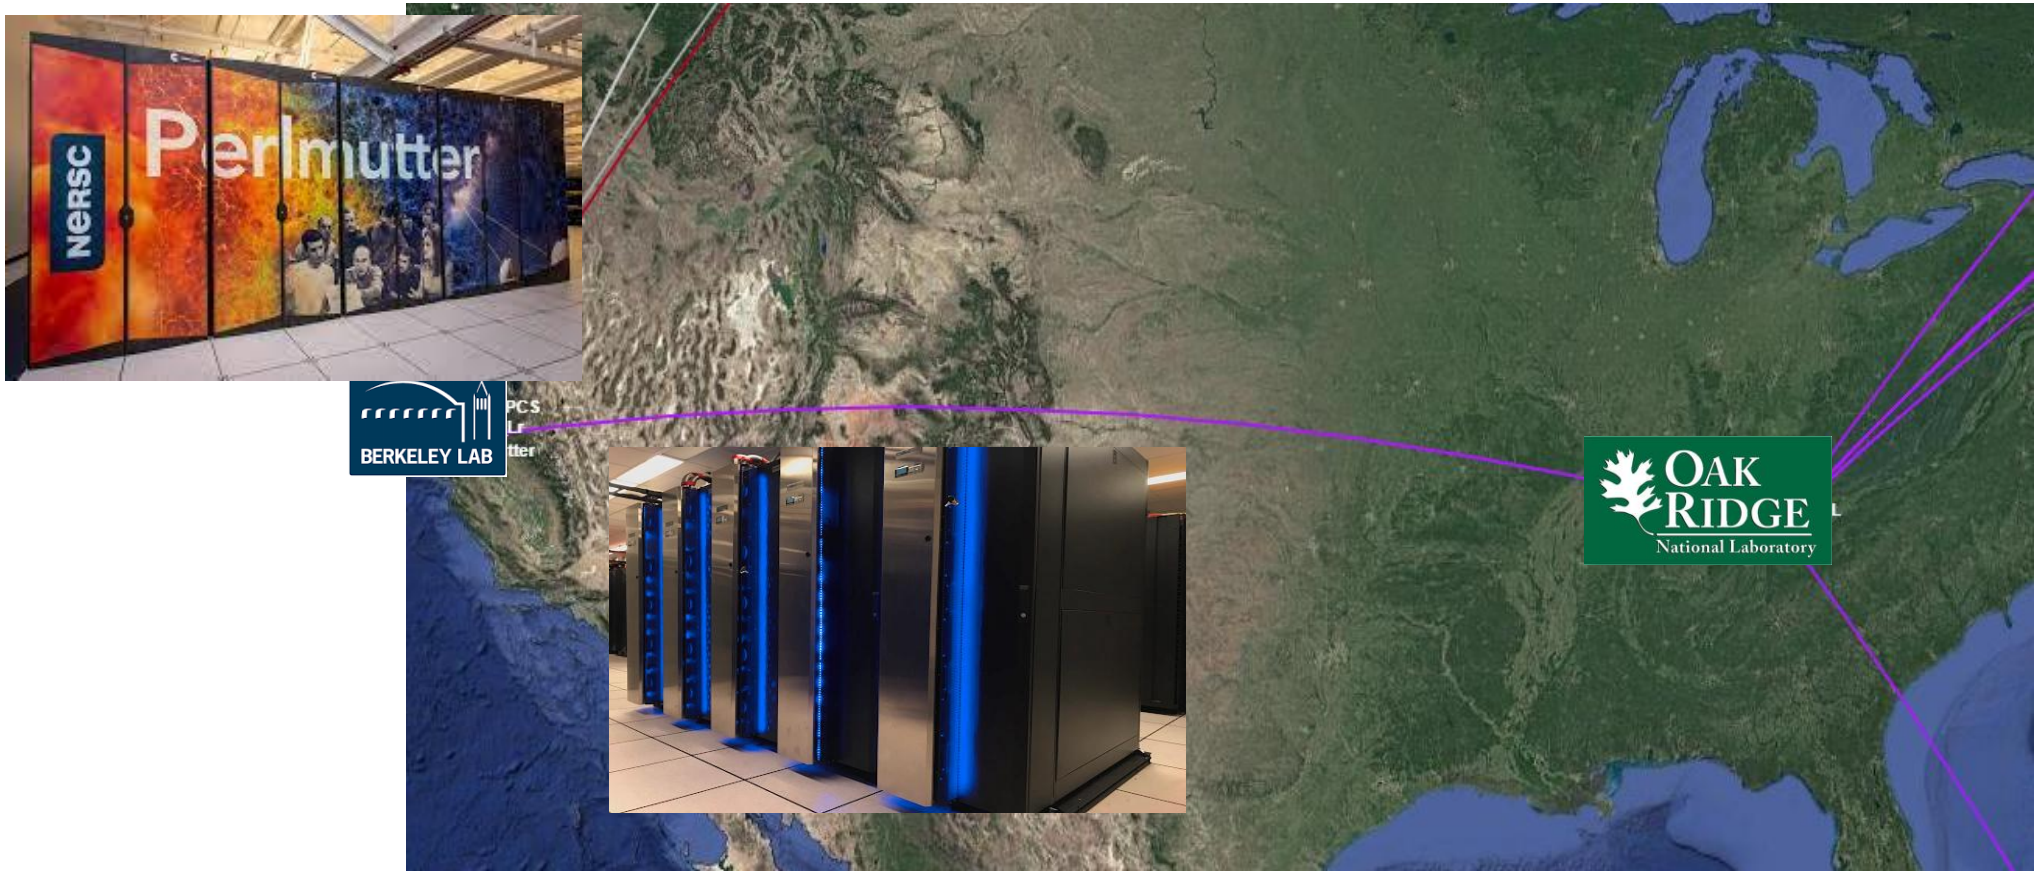

US-ALICE Computing Annual Meetings

- ALICE-USA Computing Meeting at LBNL in September 2024

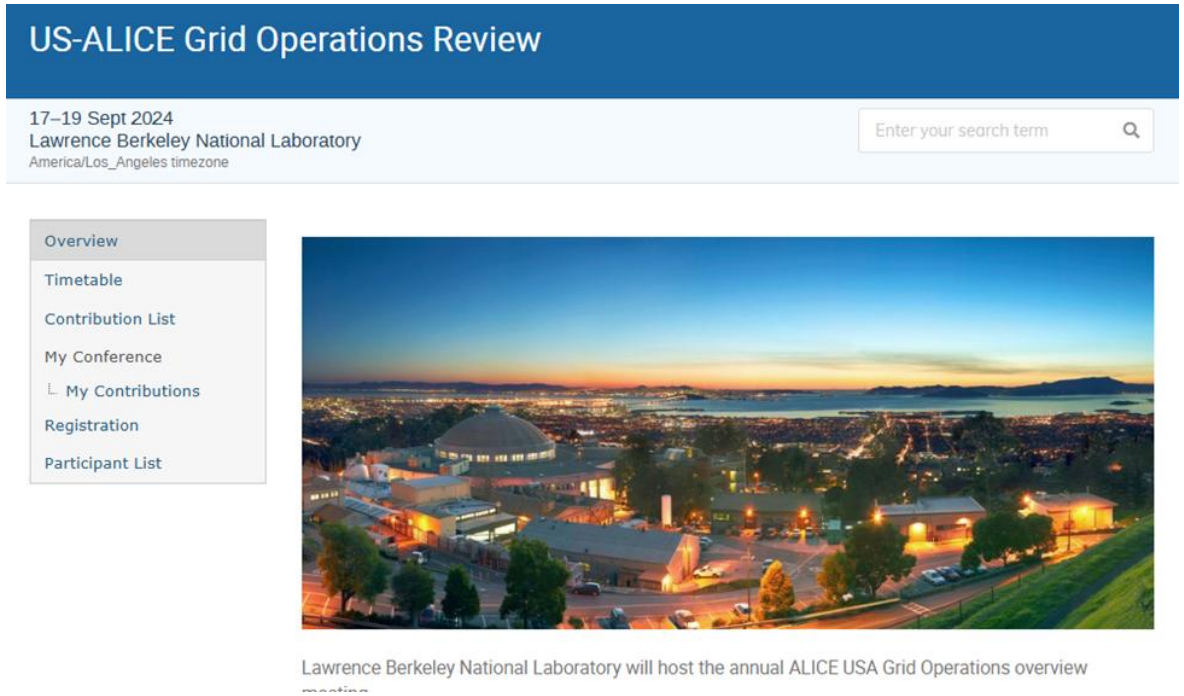

- ALICE-USA Computing Meeting at LBNL in September 2025

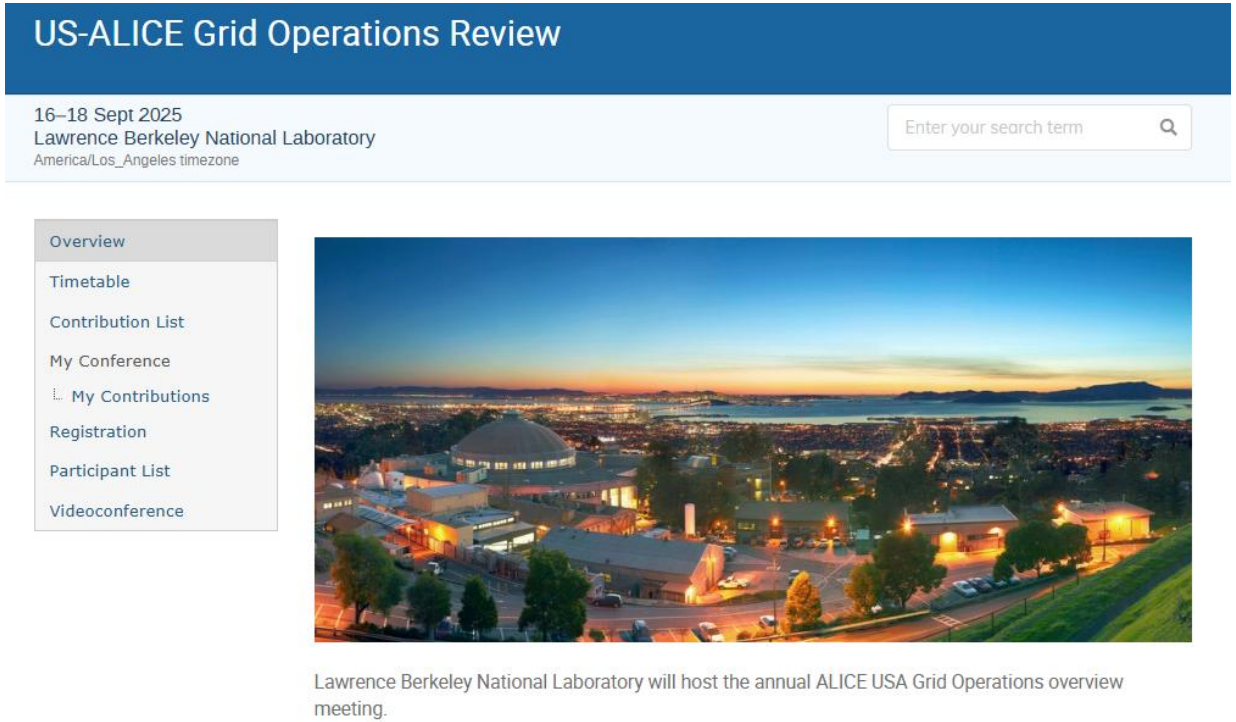

Task Tracking (in now deprecated Area120 tables)

- These tasks were identified 2.5 years ago
- The remaining tasks have not changed
  - CCDB at LBNL has not been pressing as we do not see the bottleneck with CCDB access after putting ORNL replica in place
  - Perfsonar has not been really required anywhere – this task has been scrapped
  - IPv6 setup at ORNL has progressed but not finalized
- Done but waiting for a favorable time
  - New AF proposal to DOE

1	LBNL	☐	Perfsonar	3/1/23	ASAP	1 year ago
2	LBNL	☑	LHCOne	4/1/23	ASAP	1 year ago
3	LBNL	☐	CCDB repository setup	1/20/23	ASAP	1 year ago
4	LBNL	☑	Migrate from Cori to Perlmutter	6/15/23	ASAP	1 year ago
5	LBNL	☑	Migrate to whole-node scheduling	5/31/23	ASAP	1 year ago
6	LBNL	☑	RHEL 8 switch on HPCS	2/1/23	ASAP	just now
7	LBNL	☑	aliprod shared user on sfapi-vobox		ASAP	1 year ago
+ Add row Shift-Enter						
ry ORNL Count 4						
					ASAP: 3	Later: 1
8	ORNL	☐	Perfsonar	3/1/23	ASAP	1 year ago
9	ORNL	☑	CTF Production Resource (PRF)	1/10/23	ASAP	1 year ago
10	ORNL	☑	CCDB repository setup	1/20/23	ASAP	1 year ago
11	ORNL	☐	IPv6		Later	1 year ago
+ Add row Shift-Enter						
ry General Count 1						
					ASAP: 1	
12	General	☐	New AF Proposal	4/1/23	ASAP	1 year ago

Status of operations

Monthly - March (since today we nominally would have had a monthly meeting)
+ semi-annual since the last meeting @ LBNL

WLCG: August 2025

Tier-2 Availability and Reliability Report
ALICE

March 2026

Federation Details

Availability Algorithm: @ALICE_CE * @ALICE_VOBOX * all AliEn-SE

Color coding: N/A <30% <60% <90% >=90%

Federation & Sites	Pledge CPU	Pledge Disk	Availability	Reliability	Unknown	Availability History Dec-2025	Jan-2026	Feb-2026
US-LBNL-ALICE	**119700**	**13100**						
LBL_HPCS			97%	97%	0%	95%	60%	62%
ORNL			100%	100%	0%	94%	97%	96%

Our availability/reliability has been fully driven by the SE

Significant problems around (perhaps coincidentally) the sitewide IB upgrade at LBL… details later

…but the problem is solved

WLCG: October 2024 to Now

- Availability Algorithm:

  @ALICE_CE * @ALICE_VOBOX * all **AliEn-SE**

- Since October 2024 we have reported
  - **31 days and 22 hours**. This adds up to a total of **766 hours.** (Most of it comes from the last two weeks of power work).
  - Only 9 hours downtime at ORNL.
- Low availability is solely due to the challenges related our SEs @ both sites, which have been “all addressed”

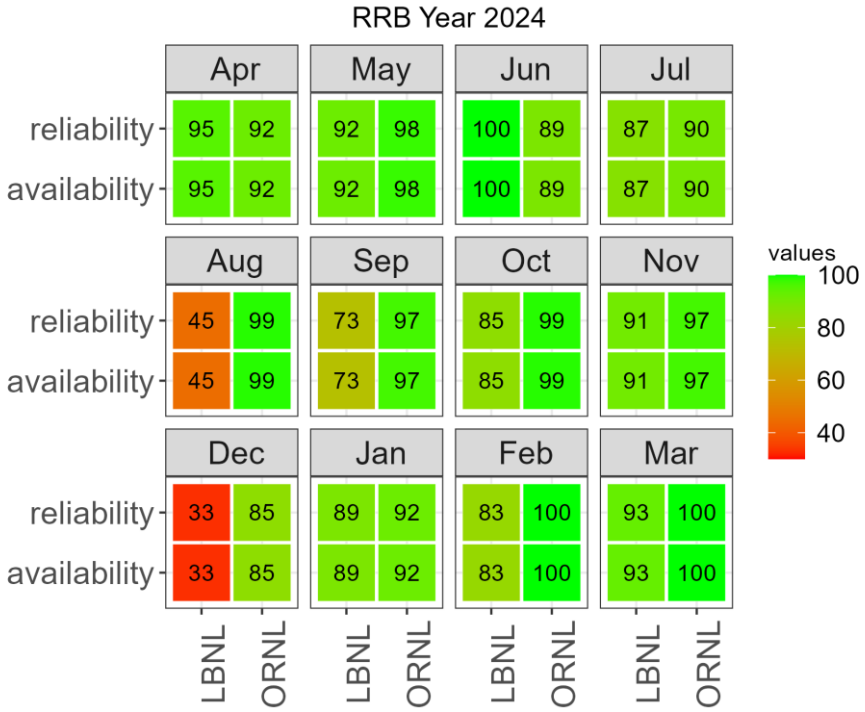

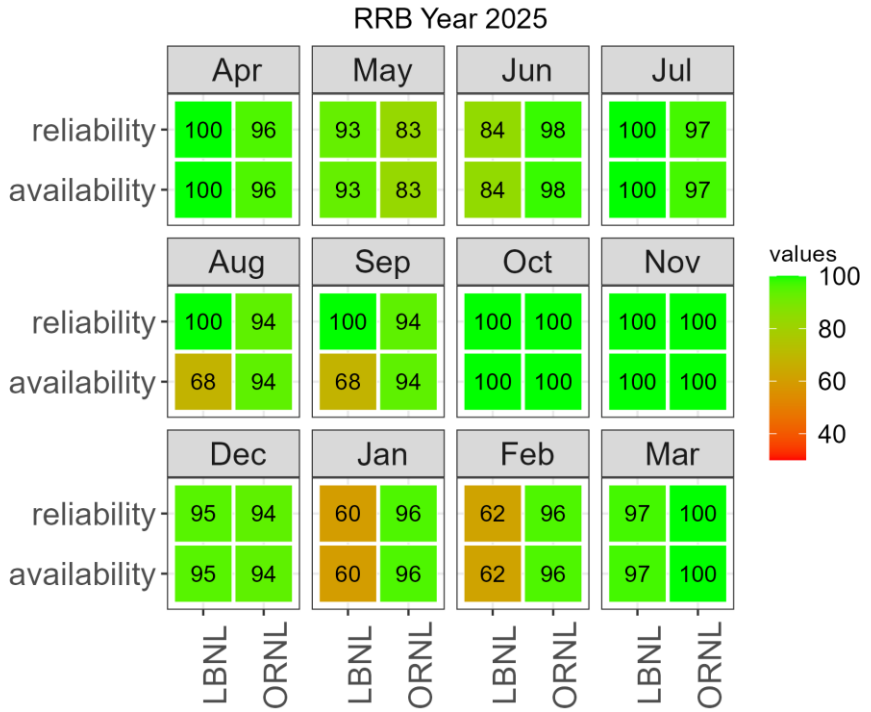

March 2025

- All sites had a very stable load and handled it well
- Still no jobs @ Perlmutter, but I see the slowdown in the user consumption, which means we can crank up the grid usage…
- … however, we might need to implement a split between GPU and CPU in JAliEn
- Dips in ORNL load are from the job cancellation issue that started in August 2025… (fingers crossed after the HTCondor update)

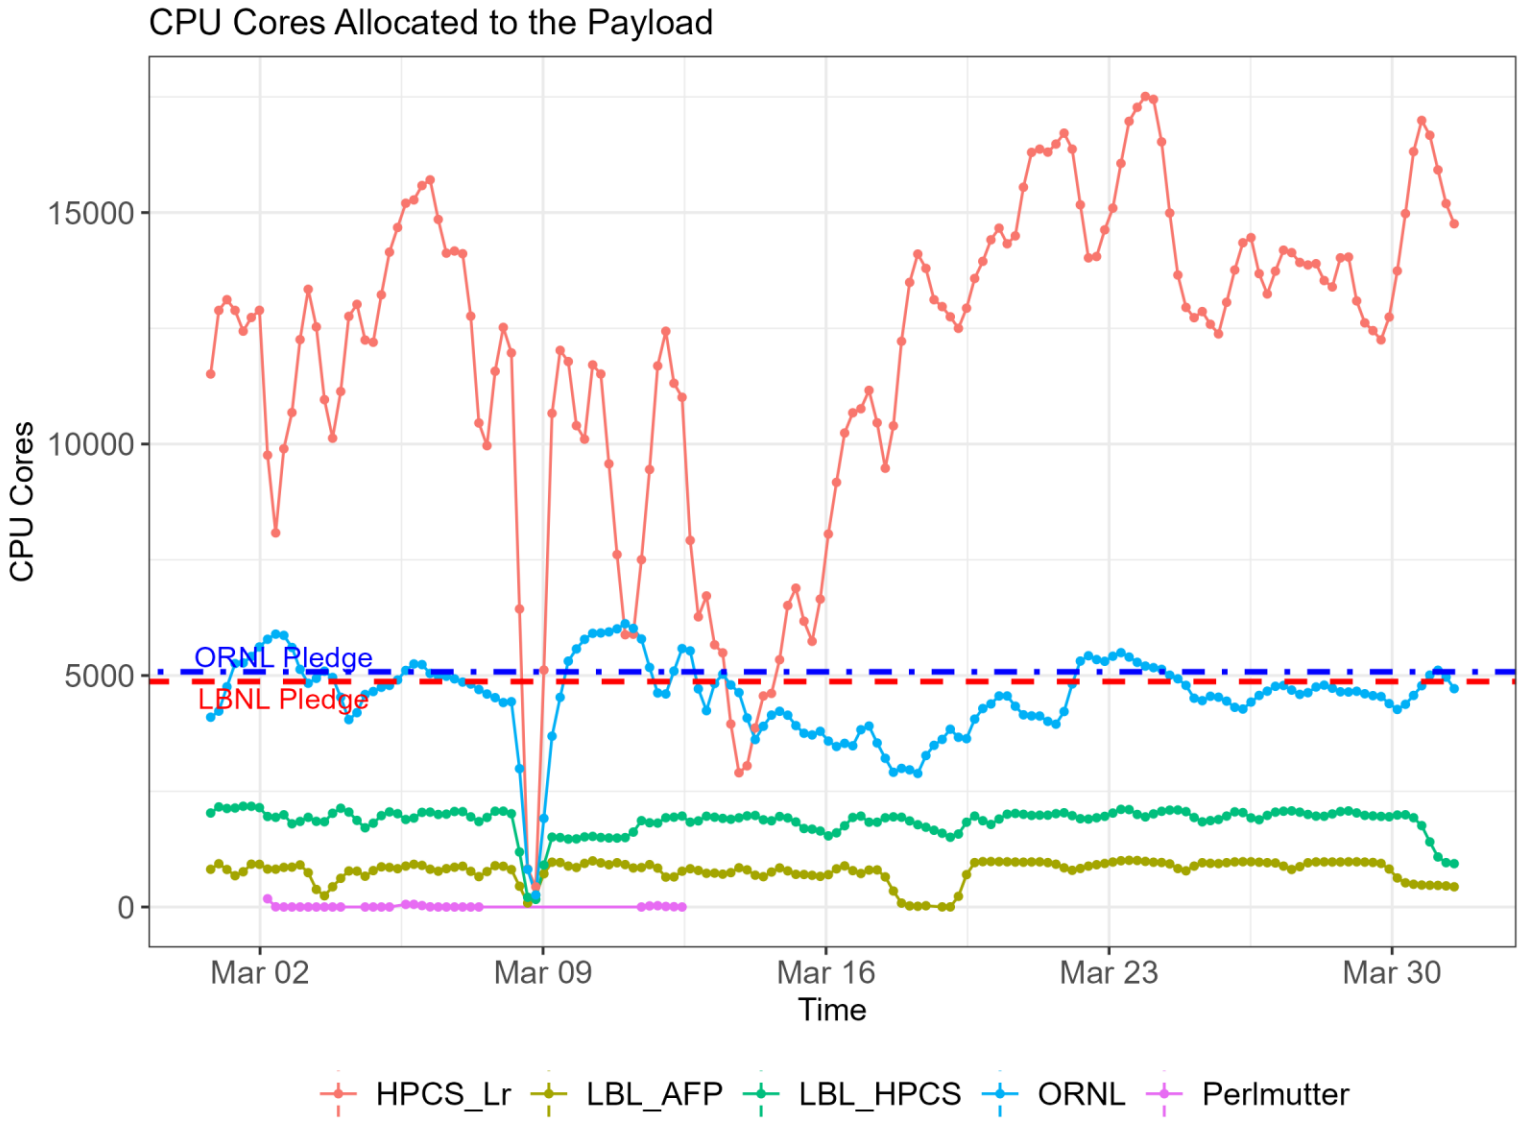

October 2025 to Now

Compute Elements [eventful half a year]

- **LBL-HPCS**
  - Steady performance
  - IB network upgrade
- **ORNL**
  - Steady performance…
  - …with a job cancellation problem
- **Lawrencium**
  - Very reliable and valuable resource
  - Some downtime after scratch system issues
  - Stabilized performance by moving away from the scratch
- **Perlmutter**
  - Up when we had enough resources
  - 2026 ERCAP allocation only 50k node-hours, significantly smaller than in 2025
- **LBL_AFP**
  - Very steady performance
  - Problematic nodes are being switched off
  - Two nodes lost so far

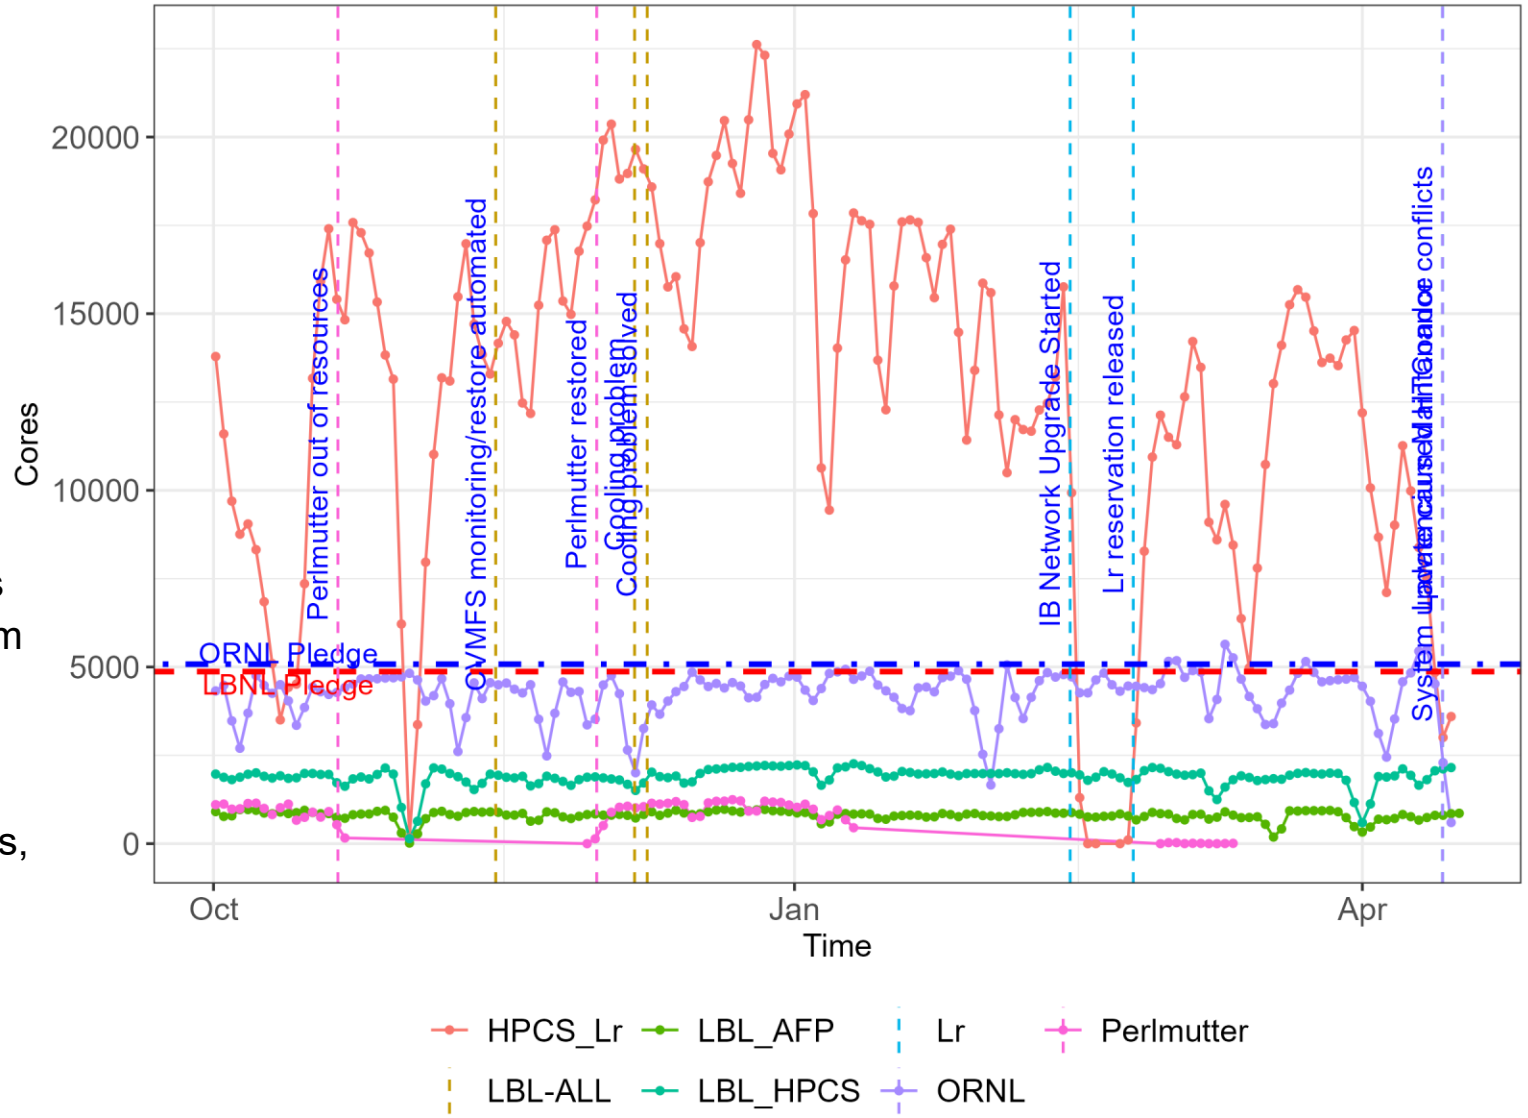

Efficiency: March 2025

Jobs efficiency (cpu time / wall time)

	Series	Last value	Min	Avg	Max
1.	Hiroshima	84.23	0.138	74.53	125.3
2.	HPCS_Lr	86.33	6.153	84.5	134.1
3.	LBL_AFP	100.9	0.285	98.12	247.8
4.	LBL_HPCS	101.6	0.326	99	164.5
5.	NIHAM	72.09	24.39	80.81	115.2
6.	ORNL	94.76	3.101	101.5	144.5
7.	UPB	60.01	5.335	78.47	150.9
	Total	**85.7**		**88.14**	

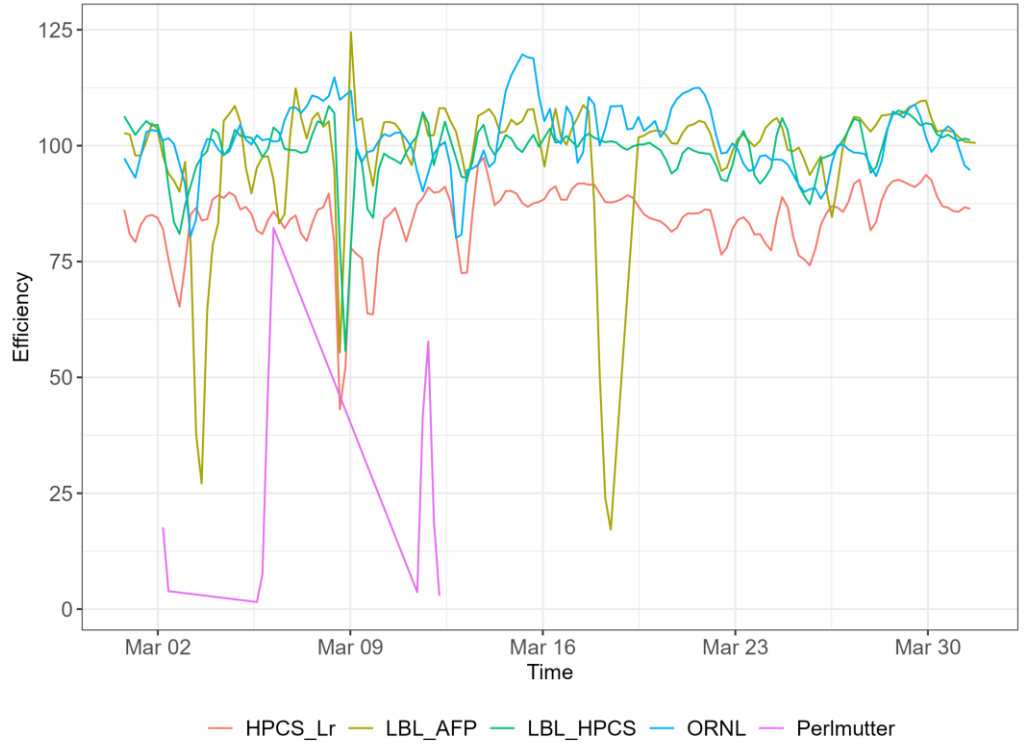

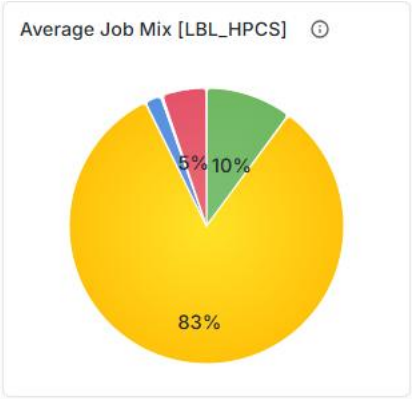

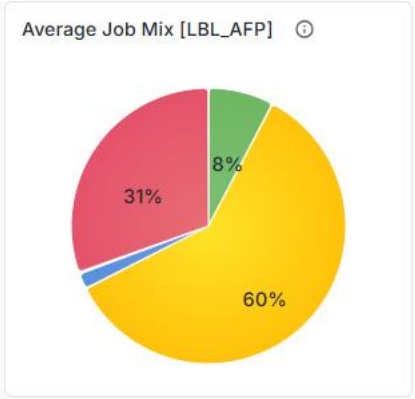

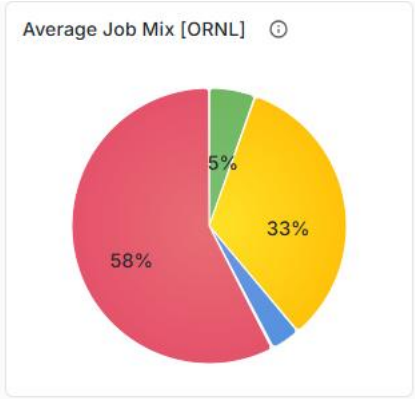

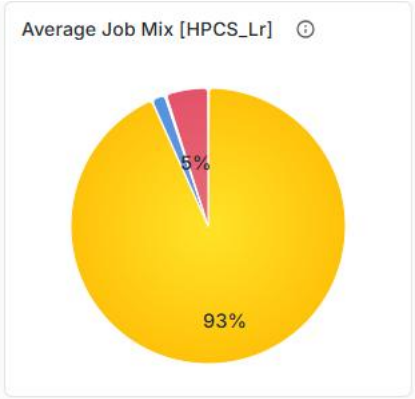

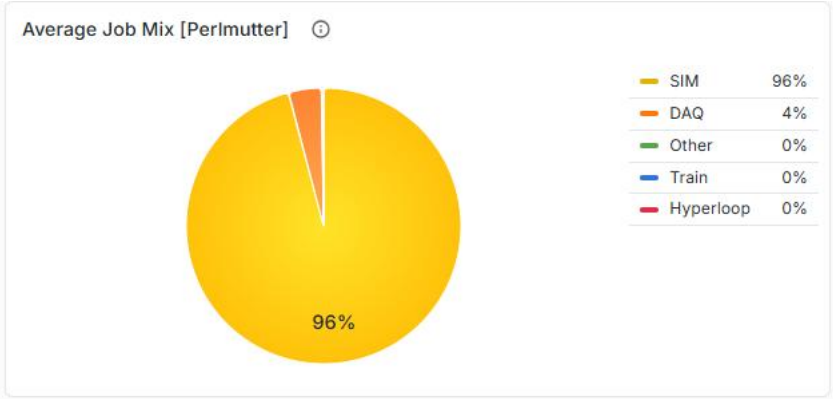

Accounting: March 2025

CPU Obligations	kHS23
ALICE-USA	116.8
LBNL HPCS T2	58.4
ORNL T2	58.4

T2 Site	Monalisa (khr)	WLCG (khr)	%diff
HPCS_Lr	8752.60	33526.00	68.11
LBNL	1366.00		
AFP	574.30		
ORNL	2882.50	3017.00	4.46
Total	**13,575.40**	**36,543.00**	**62.85**

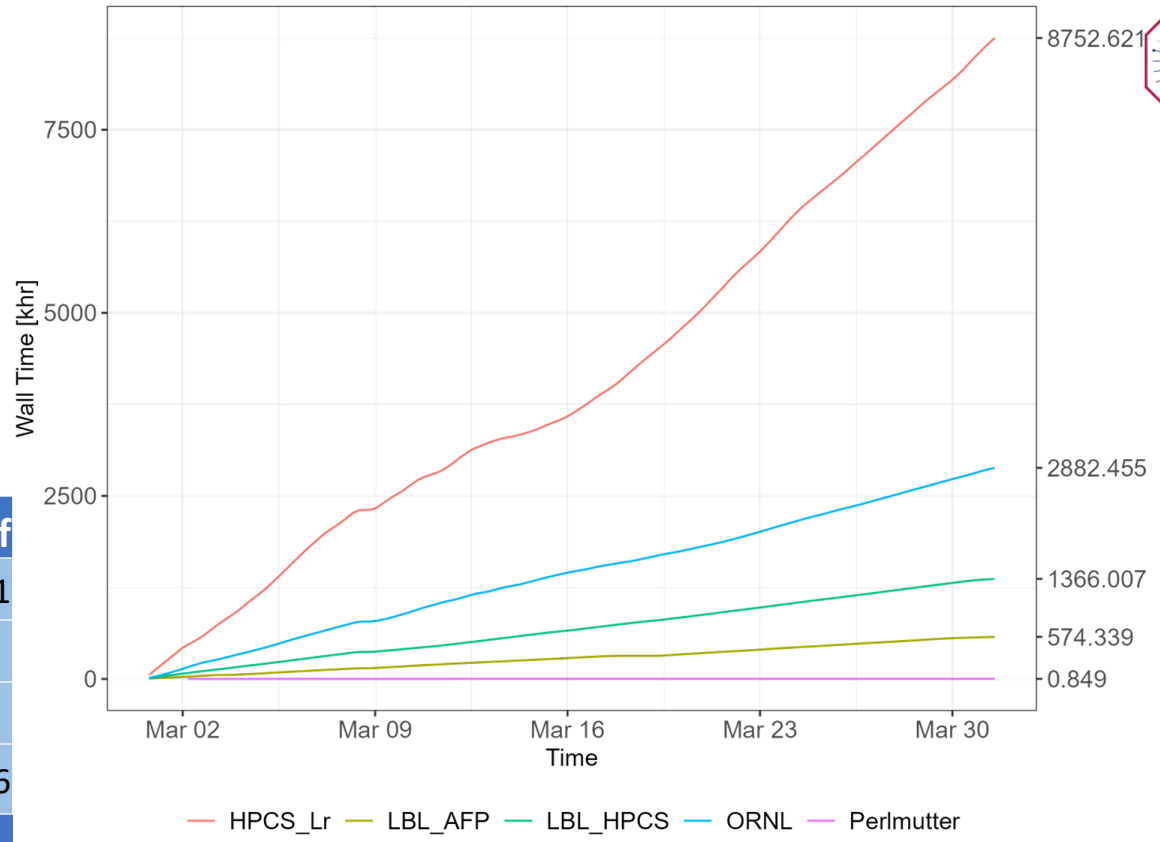

N.B. We do not run accounting on Perlmutter

T2 Site and HPCS_Lr	CPU/Core [HS23/Core]	ALICE-USA Obligation [MHS23 x hr]	CPU Delivered [MHS23 x hr] (ML reporter)	CPU Delivered [MHS23 x hr] (WLCG reporter)	Delivered per Obligation [%] (ML reported)	Delivered per Obligation [%] (WLCG reported)
HPCS_Lr	10.0	43.45	87.53	335.26	201.4	771.6
LBNL	12.0		16.39		37.7	
AFP	12.0		6.89		15.9	
ORNL	11.5	43.45	33.15	34.70	76.3	79.9
Total		**86.90**	**143.96**	**369.96**	**165.66%**	**425.73%**

Accounting: October 2025 to Now

- This period covers second half of RRB2025
- Over 91,000 kHrs of CPU delivered, corresponding to about 190 kHS23 – about 60% over the obligation
- Lr continues to significantly augment our resources to help deliver the obligations
- In 2026 PEAP we put 20 kHS23 from Lr, which was very conservative estimate [*it was over 100 kHS23 of pure payload*]
- In 2027 we will probably continue to pad our resources with 20 kHS23 from Lr

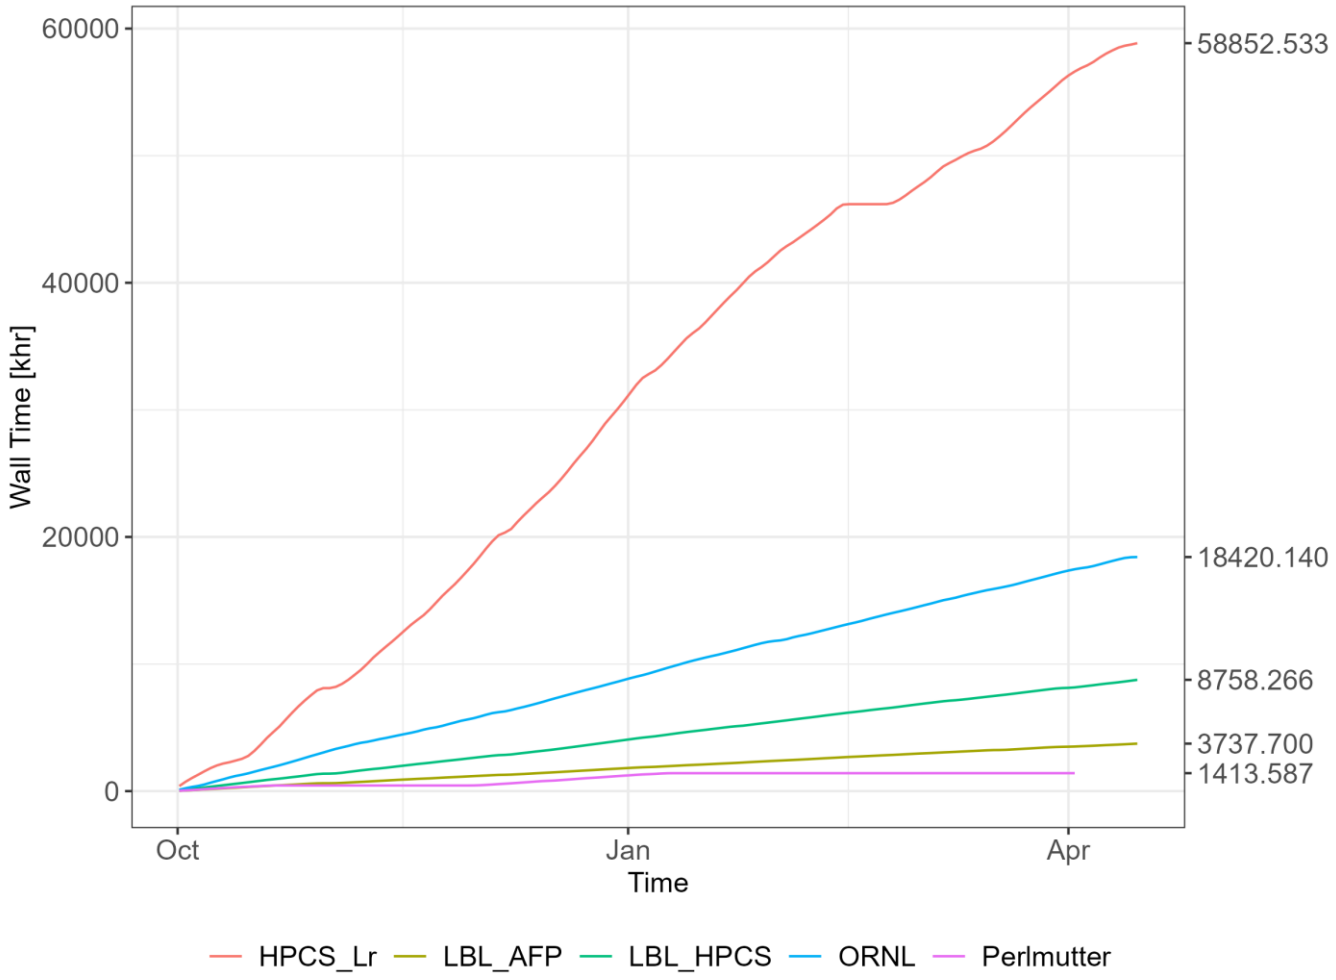

EOS: March 2025

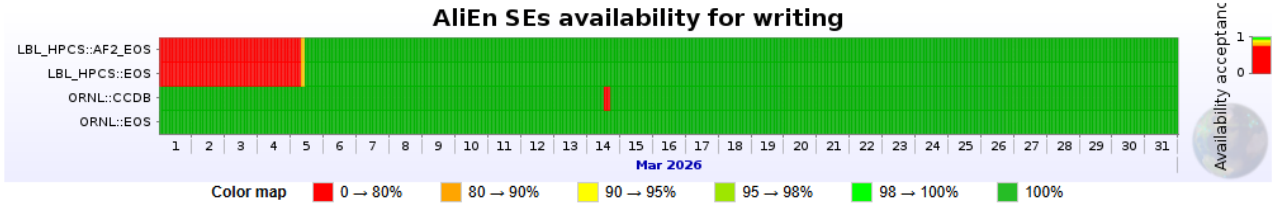

Statistics

Link name	Data		Individual results of writing tests			Overall
	Starts	Ends	Successful	Failed	Success ratio	Availability
LBL_HPCS::AF2_EOS	28 Feb 2026 23:33	31 Mar 2026 23:57	648	104	86.17%	86.13%
LBL_HPCS::EOS	28 Feb 2026 23:39	01 Apr 2026 00:04	647	103	86.27%	86.13%
ORNL::CCDB	28 Feb 2026 23:28	31 Mar 2026 23:53	753	2	99.74%	99.74%
ORNL::EOS	28 Feb 2026 23:40	01 Apr 2026 00:05	748	0	100%	100%

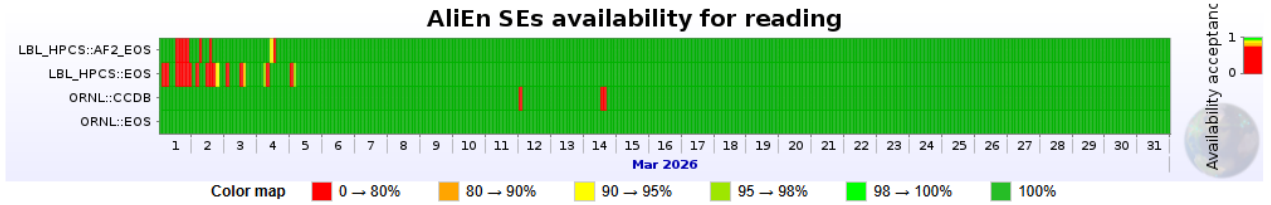

Statistics

Link name	Data		Individual results of reading tests			Overall
	Starts	Ends	Successful	Failed	Success ratio	Availability
LBL_HPCS::AF2_EOS	28 Feb 2026 23:33	31 Mar 2026 23:57	741	11	98.54%	98.52%
LBL_HPCS::EOS	28 Feb 2026 23:39	01 Apr 2026 00:04	729	21	97.20%	97.18%
ORNL::CCDB	28 Feb 2026 23:28	31 Mar 2026 23:53	752	3	99.60%	99.60%
ORNL::EOS	28 Feb 2026 23:40	01 Apr 2026 00:05	748	0	100%	100%

	AliEn SE		Catalogue statistics (1024-base units)						Storage-provided information (1024-base units)					
SE Name	AliEn name	Tier	Size	Used	Free	Usage	No. of files	Type	Size	Used	Free	Usage	Version	EOS Version
1. LBL_HPCS - EOS	ALICE::LBL_HPCS::EOS	2	6.313 PB	5.733 PB	593.6 TB	90.82%	104,685,030	FILE	7.369 PB	6.744 PB	640.5 TB	91.51%	Xrootd 5.8.4	5.3.27
2. ORNL - EOS	ALICE::ORNL::EOS	2	6.586 PB	5.809 PB	796.1 TB	88.2%	73,108,886	FILE	6.586 PB	5.877 PB	726.5 TB	89.23%	Xrootd 5.8.3	5.3.15
Total			**12.9 PB**	**11.54 PB**	**1.357 PB**	**89.48%**	**177,793,916**		**13.96 PB**	**12.62 PB**	**1.335 PB**	**90.43%**		

EOS: October 2025 to Now

- SE availability
  - In general, storage is very stable at both sites
  - LBNL EOS issues during Jan-Feb are solved and the site is back to 100% stable mode
  - Some fsts needed to be worked on at ORNL
- IPv6 has been up at LBNL since early RRB2021
- IPv6 at ORNL is to be finalized

	READ	WRITE
LBNL	86.68	84.35
ORNL	96.13	97.17

IPv4/IPv6

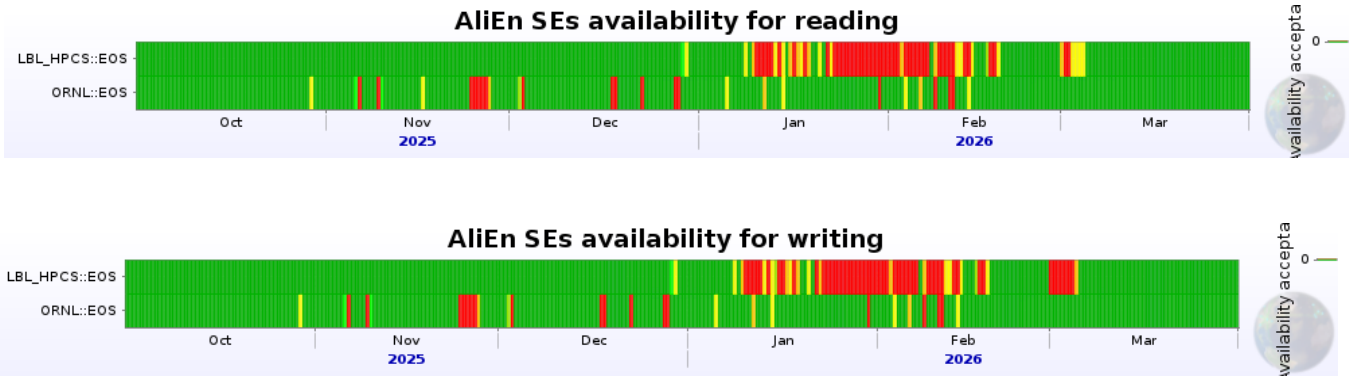

IPv6 Only

EOS: Success Ratio: October 2025 to now

- Since the file corruption mysteries of 2024 early 2025 this figure of merit has been under control and fully understood
- The deficiencies seen on the figure are
  - LBNL – overall problems with the EOS in Jan-Feb and then network issues at a couple of fsts
  - ORNL – two fsts were out to be repaired by the vendor (your might remember the intricate gymnastics Pete had to do every time he needed to reboot the server)

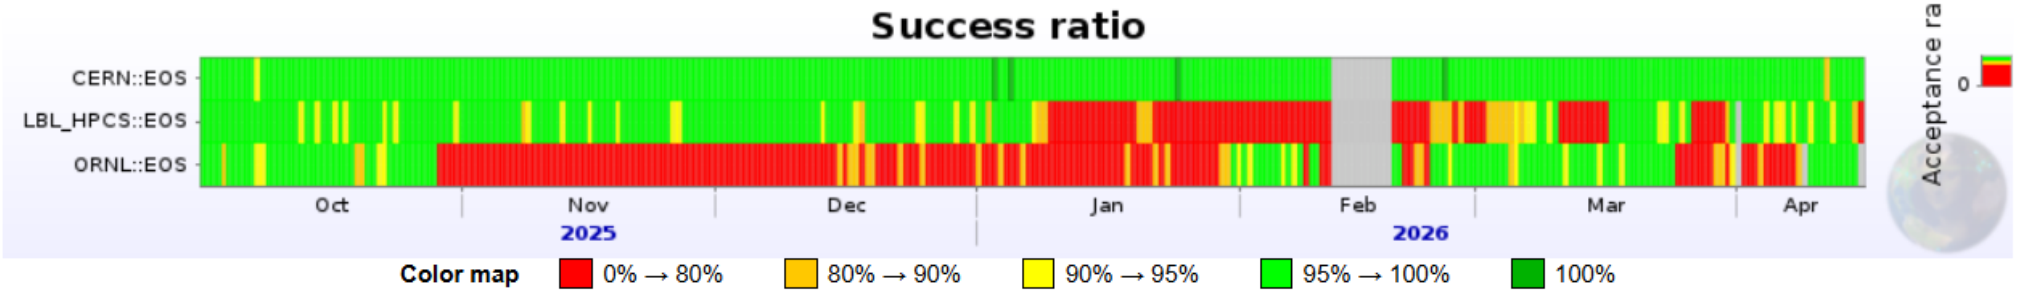

Averaged metrics for the selected interval

SE Name	Start	End	Success ratio	Corrupt ratio	Inaccessible ratio	Internal error ratio
CERN::EOS	01 Oct 2025 04:01	15 Apr 2026 23:50	99.75 %	0.01 %	0.21 %	0.03 %
LBL_HPCS::EOS	01 Oct 2025 04:03	15 Apr 2026 22:59	80.24 %	0.53 %	19.21 %	0.01 %
ORNL::EOS	01 Oct 2025 04:03	15 Apr 2026 00:19	82.07 %	0.16 %	17.75 %	0.02 %

Job Core Mix: October 2025 to Now

- A healthy mix of job types across all sites
- ORNL in general runs more hyperloop jobs than others (even the AF)
- Lawrencium has been running hyperloop jobs since October

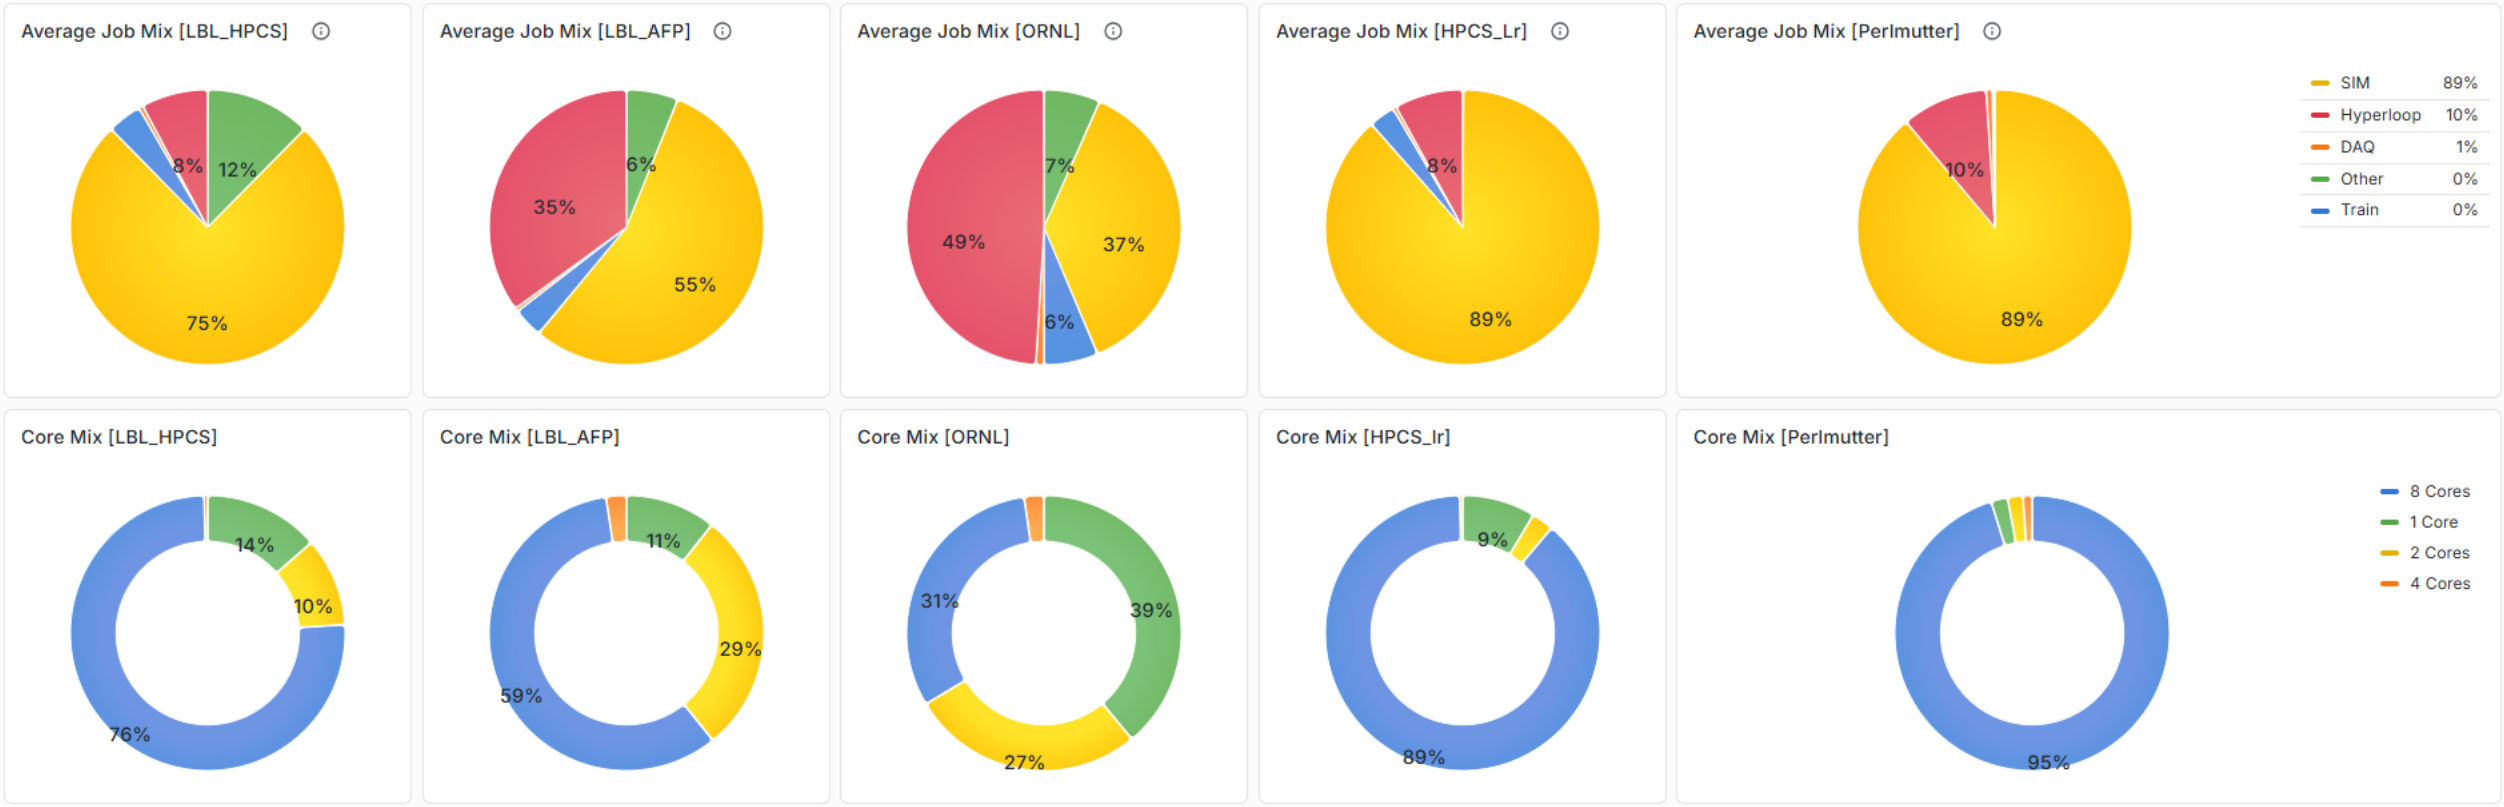

Current Issues

- We've been resolving all the issues and warning at our sites, like CVMFS update, etc
- ORNL does not even show up under issues anymore
- IPv6 at LBNL WNs is a long-term project, which we will come back to at the end of the year

Site name	VOBox and SE problems
HPCS_Lr lrc-vobox.lbl.gov	Networking: No IPv6 public address
LBL_HPCS alice-vobox.lbl.gov	Networking: No IPv6 public address
LBL_AFP alice-afp.lbl.gov	Networking: No IPv6 public address
Perlmutter alice-sfapi.lbl.gov	Networking: low buffer size kernel parameters: tcp_wmem_max=4 MB only Networking: No IPv6 public address
4 sites	**5 issues**

Some other issues/TODOs we are tackling

- Job cancellations at ORNL
  - Many ends were tied
    - Explicit resubmission removed
    - The site runs long TTL, so HTCondor dir cleanup was extended
    - TTL between CE, HTCondor, HTCondorCE, and SLURM syncronised
    - BUT… nothing helped…
    - Until… perhaps an HTCondor upgrade from ver. 23 to ver. 25 perhaps, hopefully, did
- LBNL
  - Perlmutter job log cleanup is still manual, needs to be automated
  - This is a very low priority at we are not even running at Perlmutter currently

EOS merger

- Since the deployment of the AF we operated two EOS instances
  - About 1.1 PB for the AF and the rest for the T2
- Some shortcomings in performance of AF were observed and tracked to the AF EOS performance
- The solution was to merge two EOS instances thus serving data using more fsts
- This merger was done last November and seemingly the stability has increased
- TODO: put a logical separation of quotas in EOS so we report them properly to alimonitor

	AliEn SE		Catalogue statistics (1024-base units)						Storage-provided information (1			
SE Name	AliEn name	Tier	Size	Used	Free	Usage	No. of files	Type	Size	Used	Free	Usage
1. LBL_HPCS - EOS	ALICE::LBL_HPCS::EOS	2	6.313 PB	5.734 PB	592.8 TB	90.83%	104,687,831	FILE	7.369 PB	6.745 PB	639.5 TB	91.52%
Total			**6.313 PB**	**5.734 PB**	**592.8 TB**	**90.83%**	**104,687,831**		**7.369 PB**	**6.745 PB**	**639.5 TB**	**91.52%**

	AliEn SE		Catalogue statistics (1024-base units)						Storage-provided information (			
SE Name	AliEn name	Tier	Size	Used	Free	Usage	No. of files	Type	Size	Used	Free	Usage
1. LBL_HPCS - AF2_EOS	ALICE::LBL_HPCS::AF2_EOS	2	1.083 PB	1019 TB	90.09 TB	91.88%	1,730,503	FILE	1.2 PB	1020 TB	208.8 TB	83.01%
Total			**1.083 PB**	**1019 TB**	**90.09 TB**	**91.88%**	**1,730,503**		**1.2 PB**	**1020 TB**	**208.8 TB**	**83.01%**

HTCondor

- We've been successfully running our setup for quite some time now
- Until a week ago it started glitching at ORNL
- Condor jobs on the VOBox were Idle regardless of their actual status ir SLURM
- The problem was initially coming from the mismatch between the hostnames of our VOBox `alice-vobox.ornl.gov` and `alice.vobox.ornl.gov`
- Fixing this in DNS (for the reverse-proxy lookup that HTCondor does) made the VOBox happy, however that backwars handshake between HTCondorCE and HTCondor schedule was not happening
- After some digging, I opted for the version upgrade
- This made system a little happier, but it took some config changes to make things work again

```
JOB_ROUTER_HISTORY_DIR = /var/lib/condor-ce/spool
JOB_ROUTER_SCHEDD2_SPOOL = /var/lib/condor-ce/spool
NETWORK_HOSTNAME = vobox-alice.ornl.gov
UID_DOMAIN = vobox-alice.ornl.gov
TRUST_UID_DOMAIN = True
GRIDMANAGER_SELECTION_EXPR = (Owner == "aliprod")
```

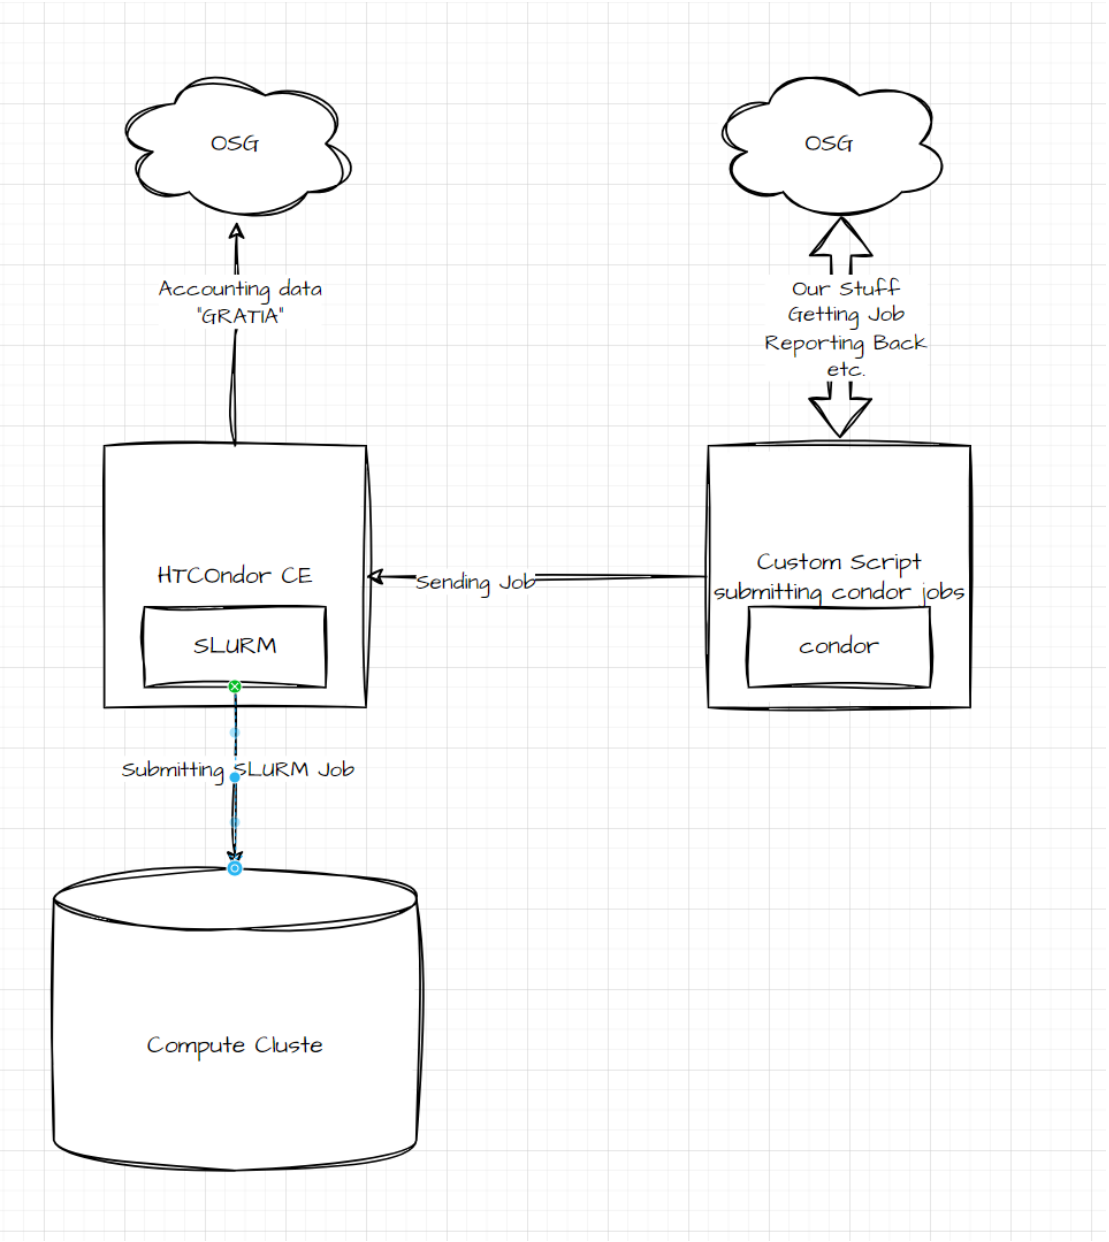

Perlmutter

- Current average usage – only users so far
- In 2025:
  - Initial allocation: 110,000 node-hours
  - At the end nearly 140,000 node-hours used
  - 80% was used as a contribution to the ALICE Grid
  - Additional allocation throughout the year: 30,000 in October

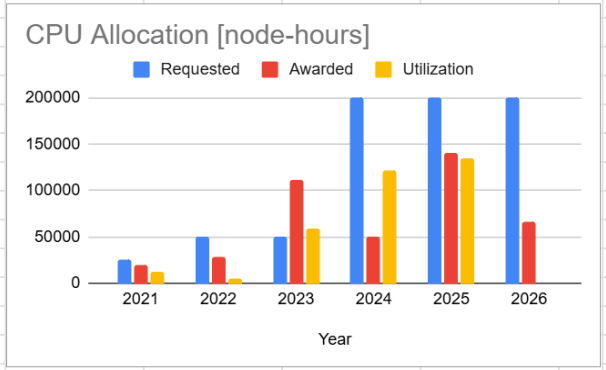

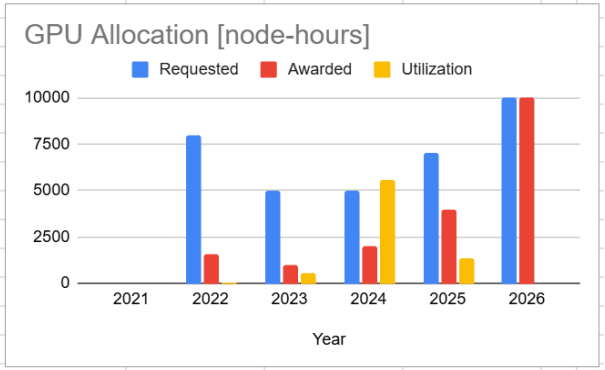

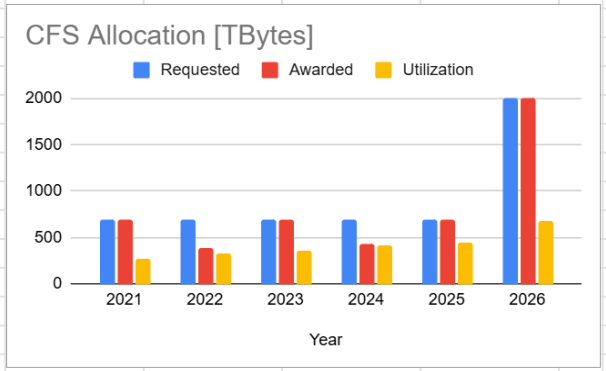

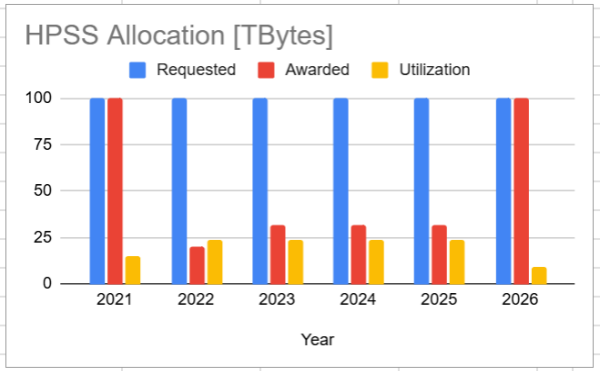

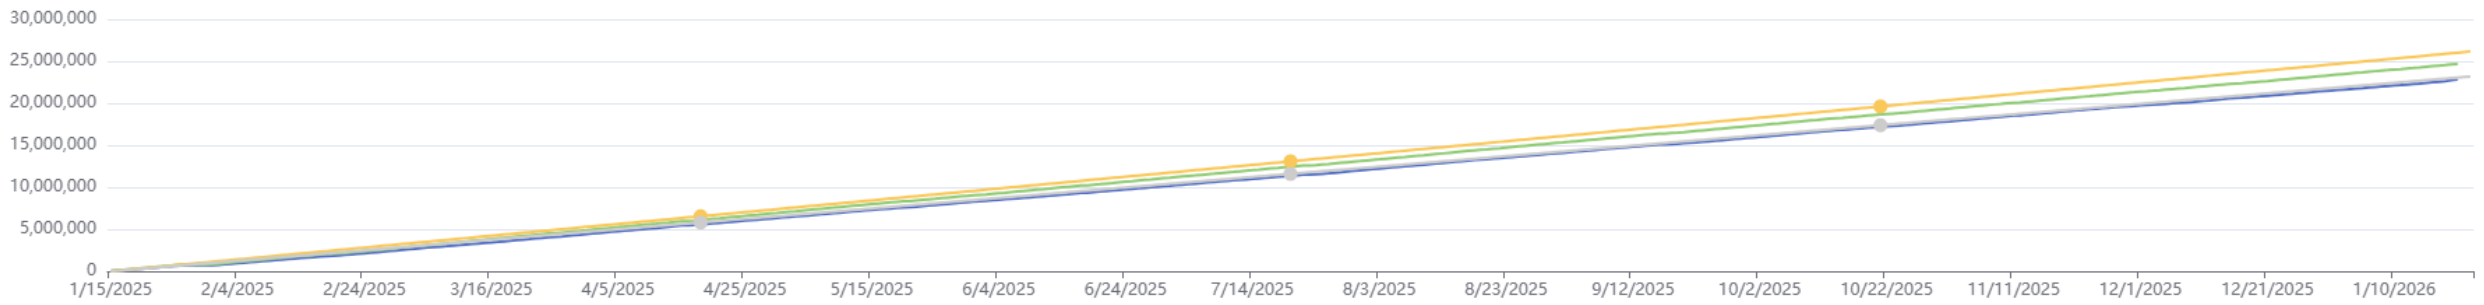

GPU @ Perlmutter

- ALICE Pb-Pb data reconstruction benefits from about 2.5 acceleration factor due to 60% of event reconstruction offloading to the GPU
- This reconstruction is constrained entirely to the ALICE EPN farm
- The software was compiled and tested for the NERSC site architecture
- Over the weekend we succeeded in submitting the jobs to the Perlmutter GPU using our SFAPI based batchqueue from the VOBox
- The details need to be sorted out

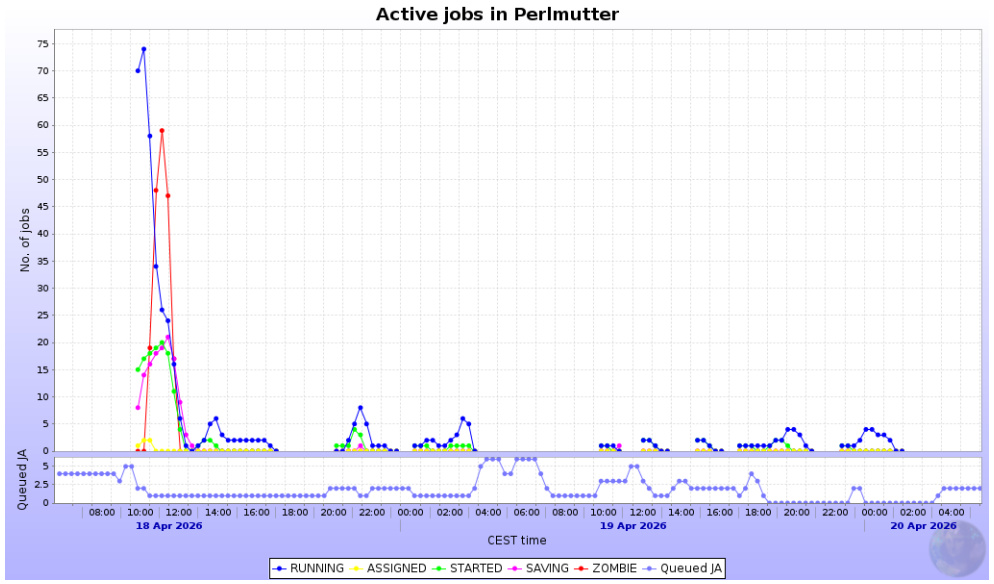

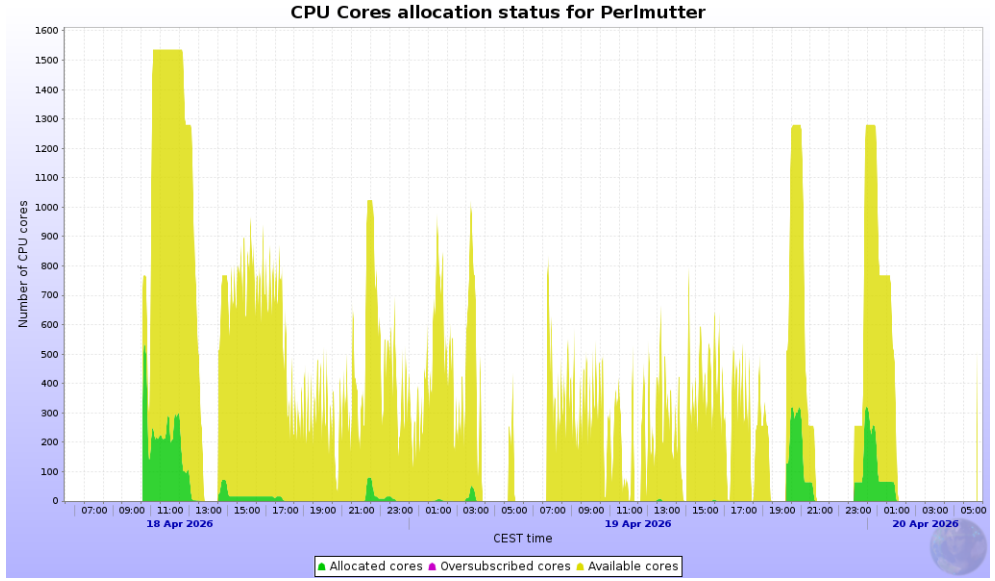

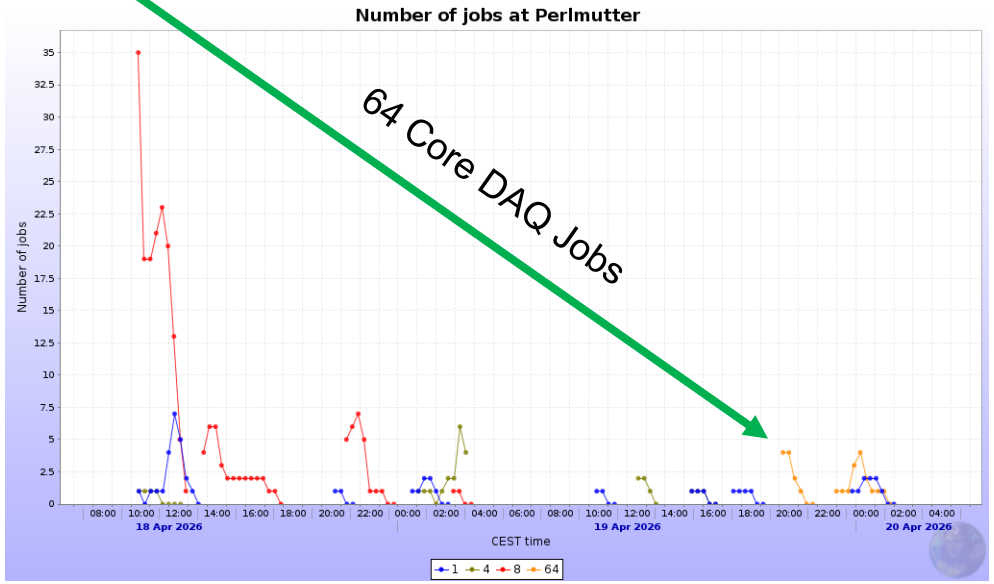

64 Core DAQ Jobs

Summary

- We continued stable operations at both sites regardless of the significant personnel change
- We managed to run sites at a high level with the aging hardware and budget shortage
- EOS stability was brought back to near 100%
- Analysis Facility stability increased with AF EOS integration into main HPCS EOS
- We now have jobs submitted to the Perlmutter GPU partition using the SFAPI from the VOBox
- Detailed site reports will be given tomorrow