U.S. DEPARTMENT OF ENERGY
Office of Science

BERKELEY LAB

ALICE-USA Computing Project Status Report

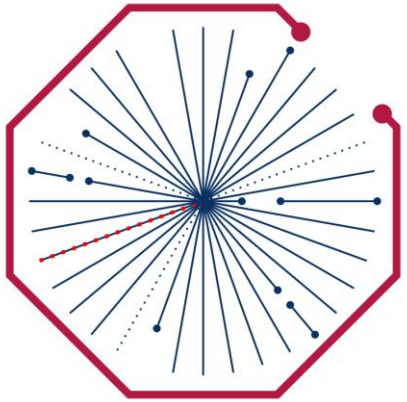

ALICE-USA Computing Annual Meeting @ LBNL

September 16, 2025

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

- 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 2023

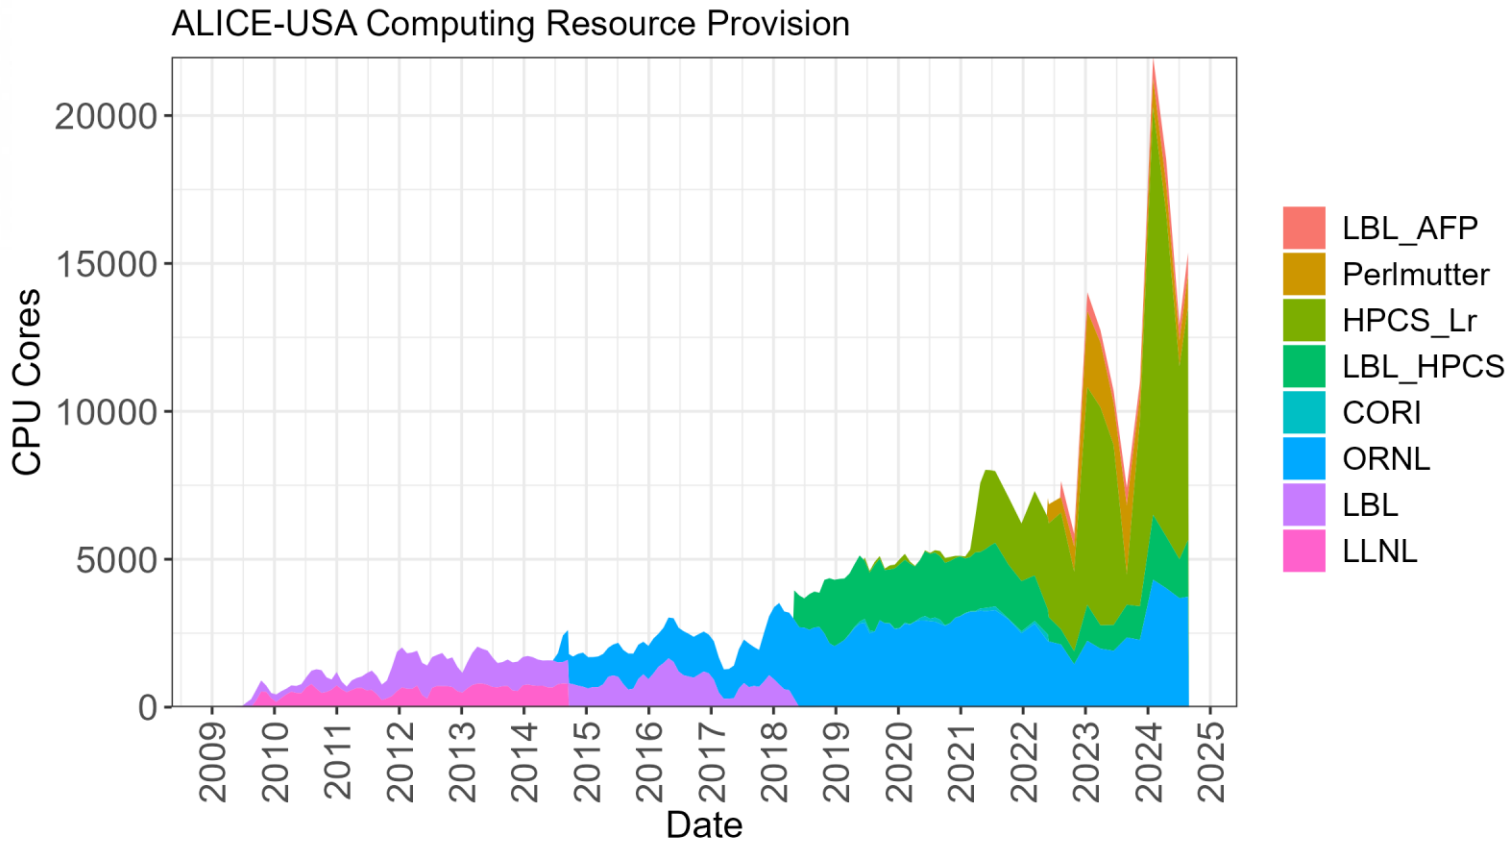

- ALICE-USA Computing Website: https://sites.google.com/lbl.gov/alice-usa-computing/meetings?authuser=1

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, 38 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

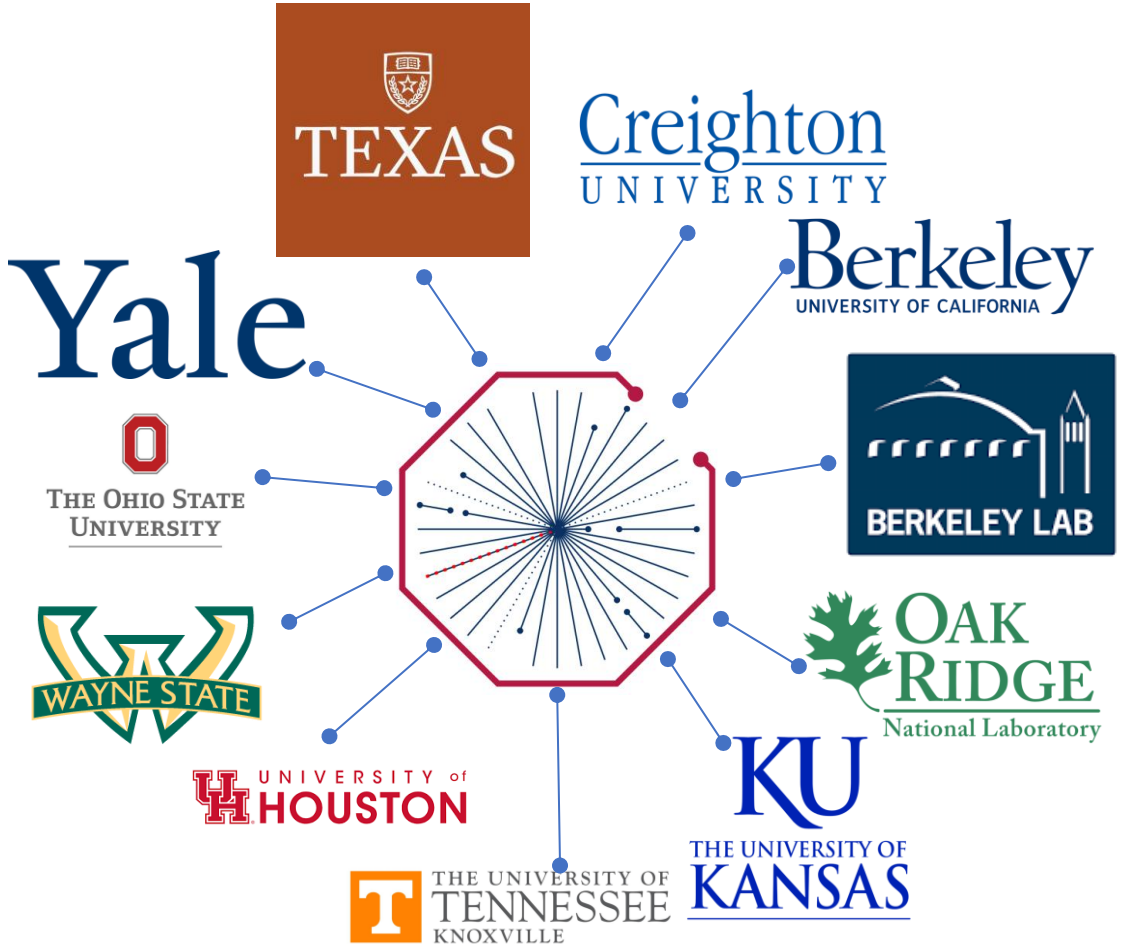

ALICE-USA T2 Sites & Analysis Facility (Prototype)

- 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 HPCs

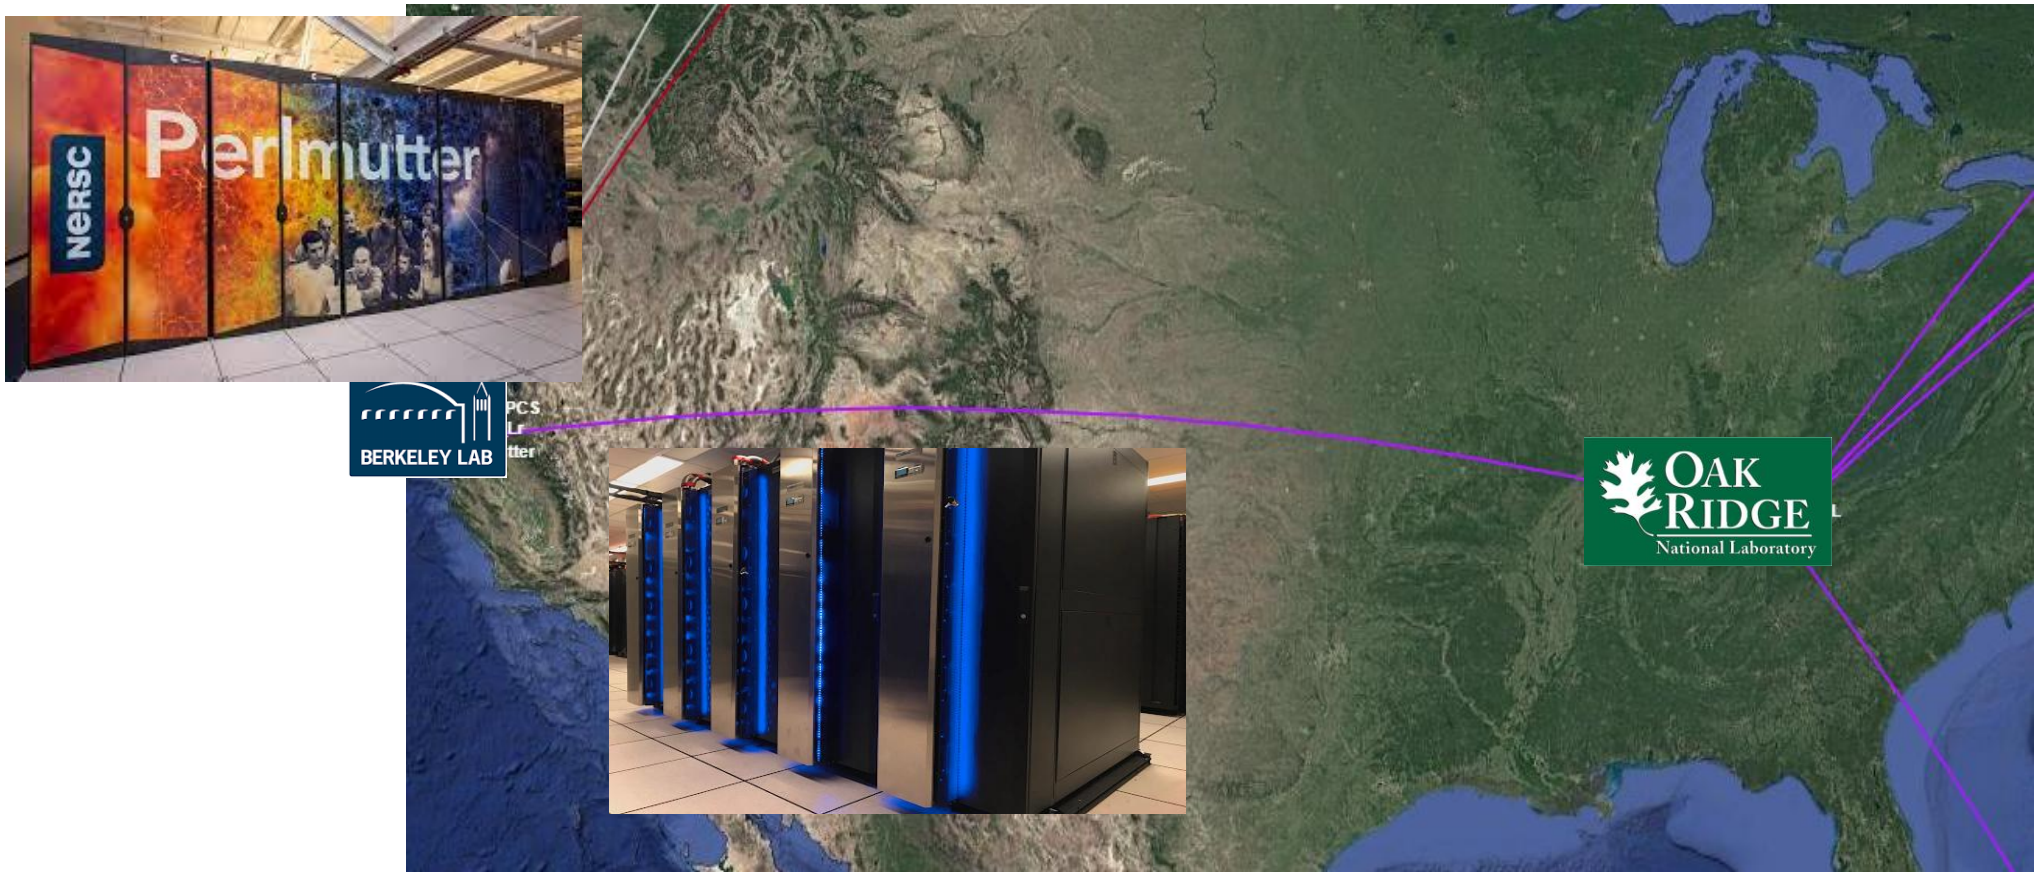

US-ALICE Computing Annual Meetings

- ALICE-USA Computing Meeting at ORNL in October 2023

3–5 Oct 2023
Oak Ridge National Laboratory
America/Los_Angeles timezone

Enter your search term

Overview
Timetable
Contribution List
My Conference
My Contributions
Registration
Participant List
Videoconference

Contact
iraklic@lbl.gov
moultonsa@ornl.gov
865 386 3733

Annual meeting hosted by the Oak Ridge National Laboratory to discuss progress and plan for the future of the ALICE-USA Computing Project.

- ALICE-USA Computing Meeting at LBNL in September 2024

US-ALICE Grid Operations Review

17–19 Sept 2024
Lawrence Berkeley National Laboratory
America/Los_Angeles timezone

Enter your search term

Overview
Timetable
Contribution List
My Conference
My Contributions
Registration
Participant List

Lawrence Berkeley National Laboratory will host the annual ALICE USA Grid Operations overview

Task Tracking (in now deprecated Area120 tables)

- These tasks were identified 2 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 but I think we need to reinstate this service
  - IPv6 setup at ORNL has progressed but not finalized
- Done but waiting for a favorable time
  - New AF proposal to DOE

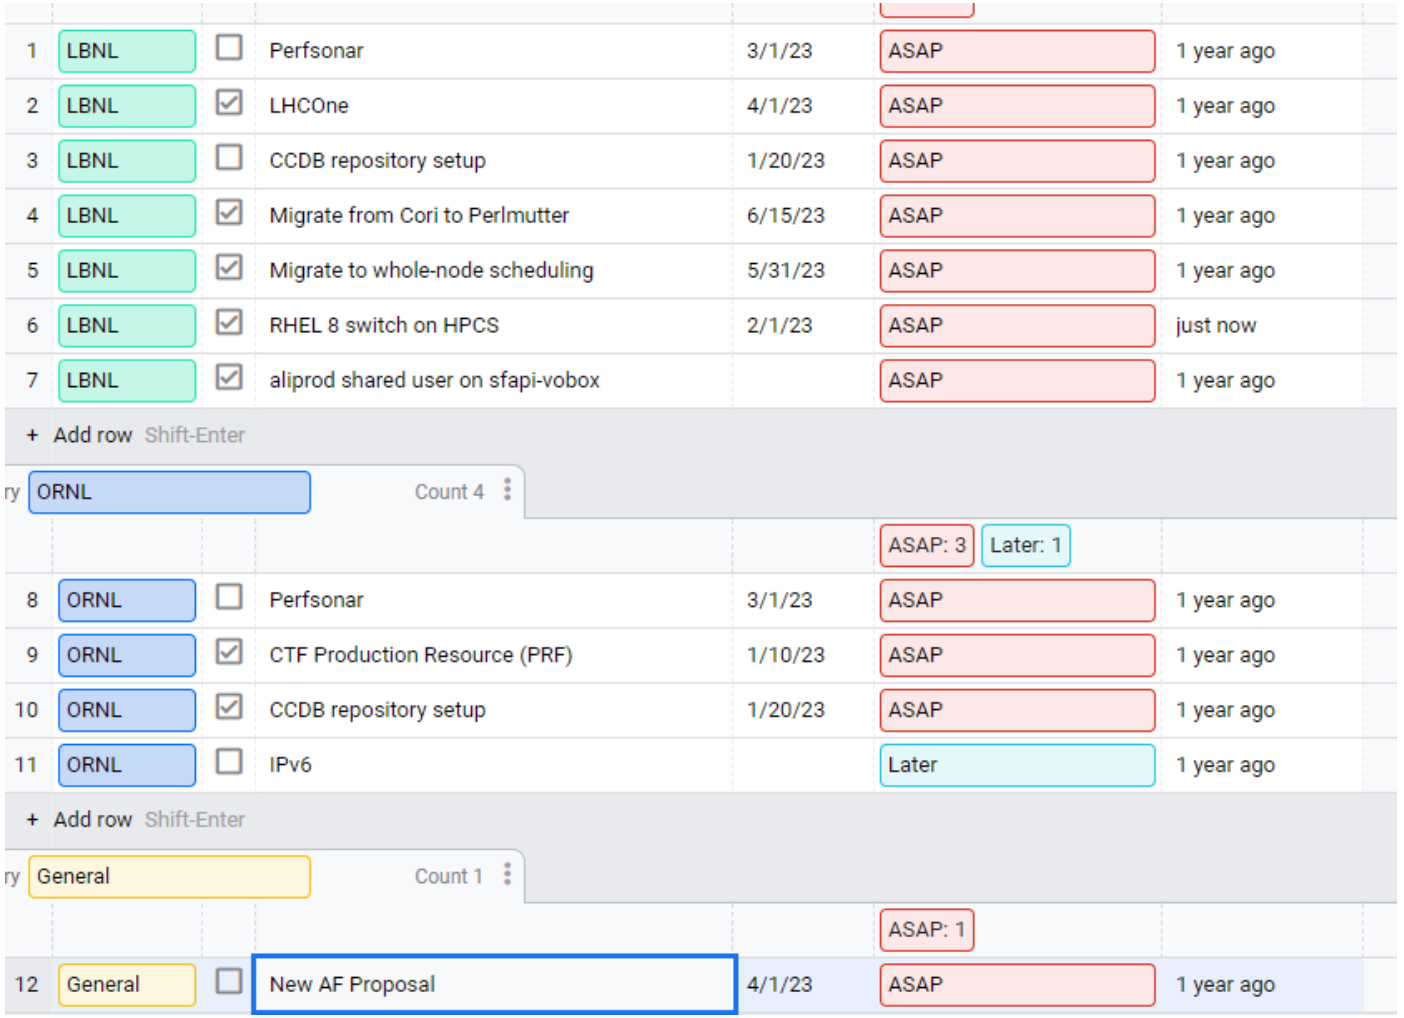

1	LBNL	☐	Perfsonar	3/1/23	ASAP	1 year ago
2	LBNL	☑	LHCOne	4/1/23	ASAP	1 year ago
3	LBNL	☐	CCDB repository setup	1/20/23	ASAP	1 year ago
4	LBNL	☑	Migrate from Cori to Perlmutter	6/15/23	ASAP	1 year ago
5	LBNL	☑	Migrate to whole-node scheduling	5/31/23	ASAP	1 year ago
6	LBNL	☑	RHEL 8 switch on HPCS	2/1/23	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 - August (since today we nominally would have a monthly meeting)
+ annual since the last meeting @ LBNL

WLCG: August 2025

Tier-2 Availability and Reliability Report

ALICE

August 2025

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 May-2025	Jun-2025	Jul-2025
US-LBNL-ALICE	**119700**	**13100**						
LBL_HPCS			68%	100%	0%	93%	84%	100%
ORNL			94%	94%	0%	83%	98%	97%

WLCG seems to have fixed the downtime reporting

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
- You can see that we've remedies these issues by automatic notifications and automatic restarts

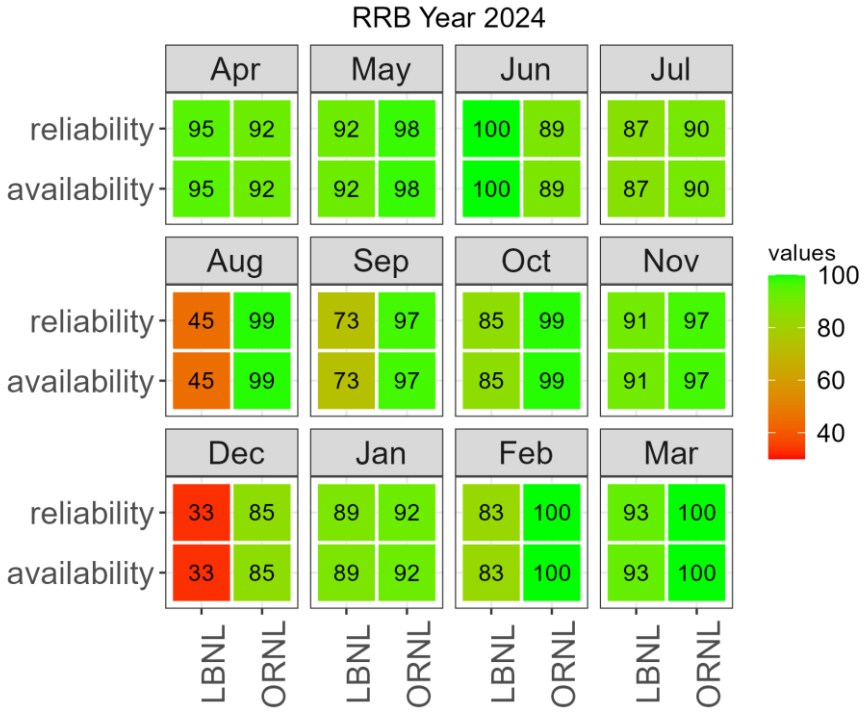

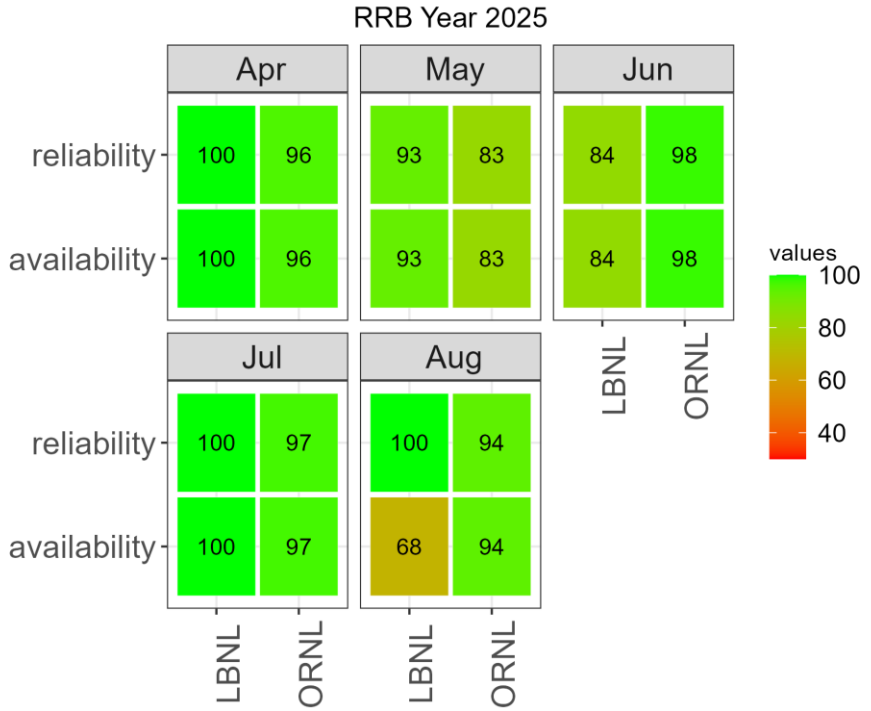

August 2025

- All sites had pretty stable load and handled it well
- Power work last two weeks of August at LBL
- VOBoxes vontinued to be up to Perlmutter kept getting jobs
- Dips in ORNL load are from the job cancellation issue that started in August

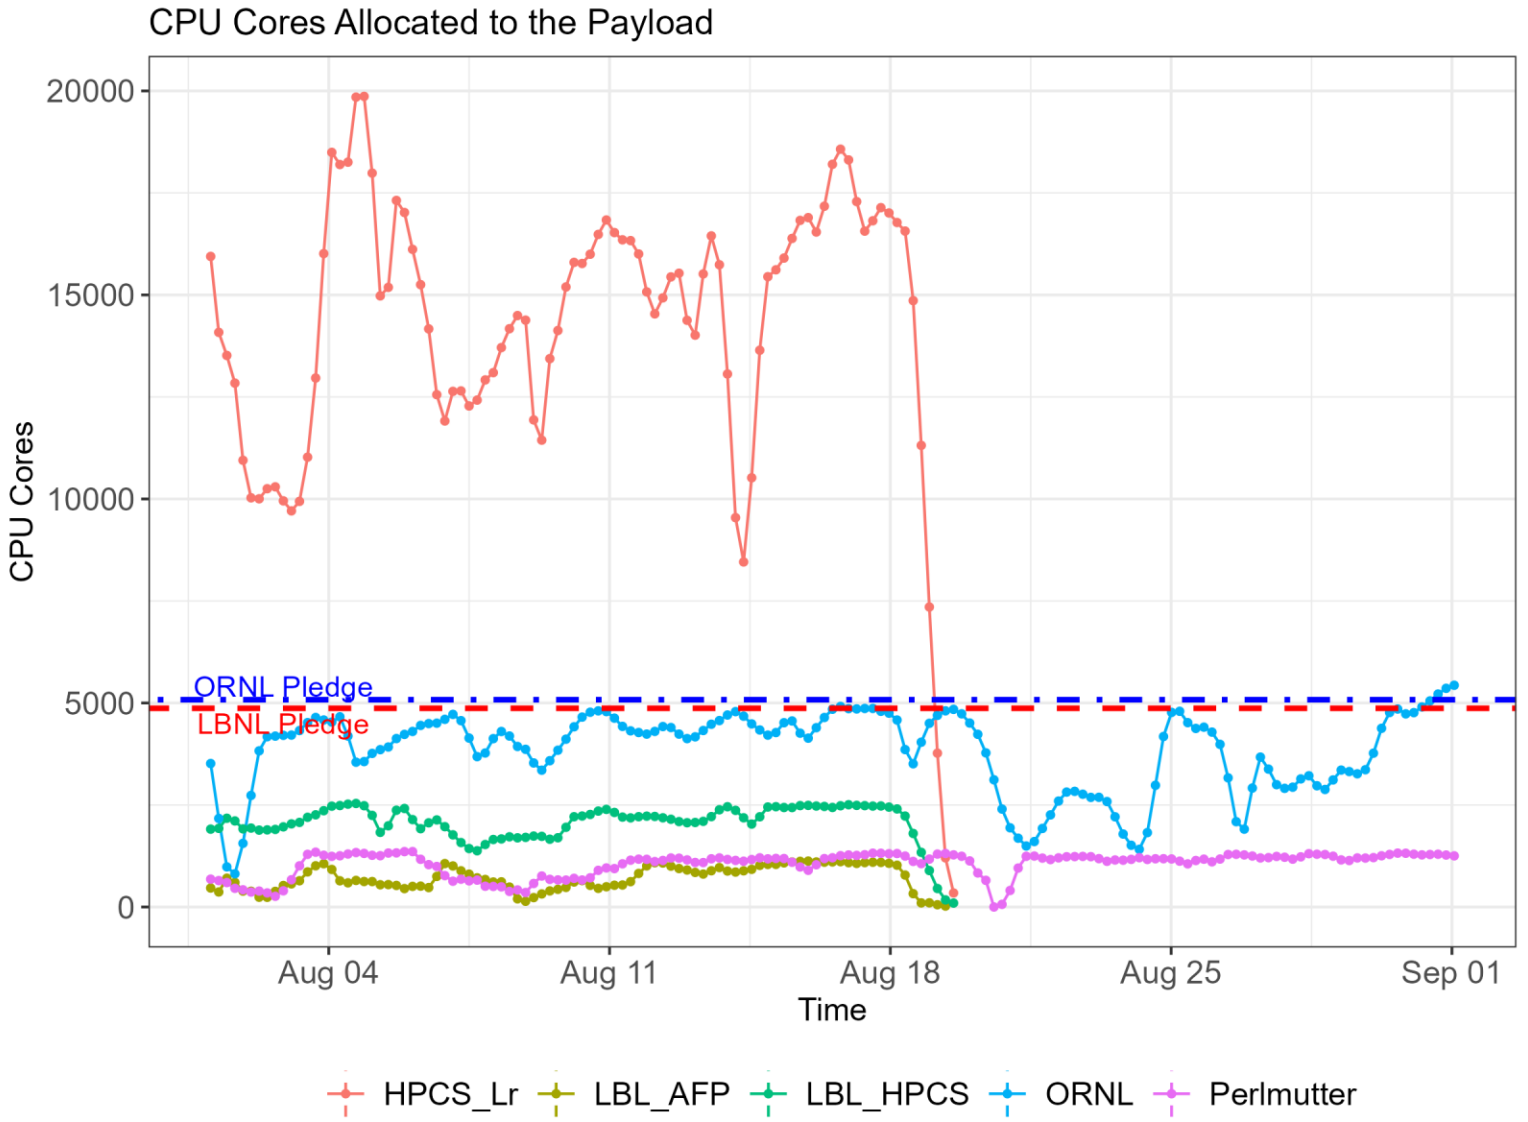

October 2024 to Now

Compute Elements [challenging year]

- **LBL-HPCS**
  - Steady performance
  - Power work interruptions twice
- **ORNL**
  - Steady performance
  - Grew by about 2700 cores at the end of last year
- **Lawrencium**
  - Very reliable and valuable resource
  - Some downtime after scratch system issues
  - Stabilized performance by moving away from the scratch
- **Perlmutter**
  - Steady allocation of nodes
  - Something throttles the increase in jobs… needs to be investigated
- **LBL_AFP**
  - Pretty steady performance

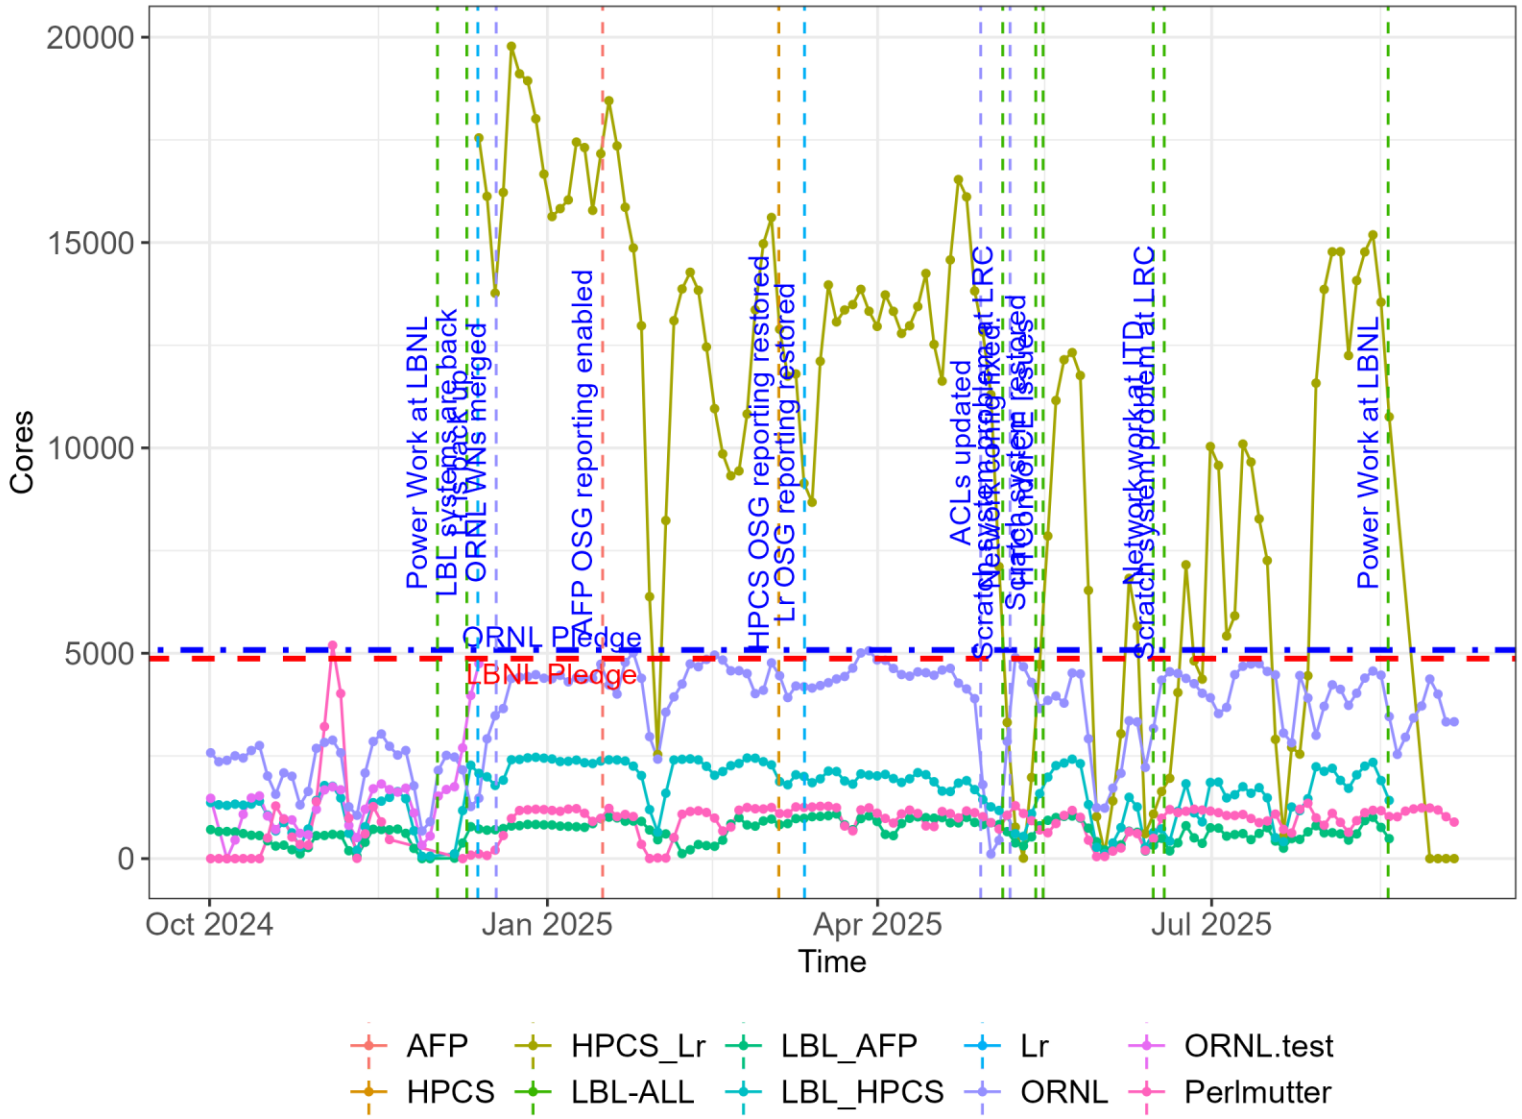

Efficiency: August 2025

Jobs efficiency (cpu time / wall time)

	Series	Last value	Min	Avg	Max
1.	HPCS_Lr	50.16	0.236	64.65	131.3
2.	LBL_AFP	73.65	1.875	85.72	158.2
3.	LBL_HPCS	72.35	0.892	79.96	192.3
4.	NIHAM	47.35	0	70.71	112.7
5.	ORNL	91.3	0	94.36	125.1
6.	Perlmutter	83.92	0.084	76.51	204
7.	UPB	79.32	1.652	80.71	295.8
Total		**71.15**		**78.94**	

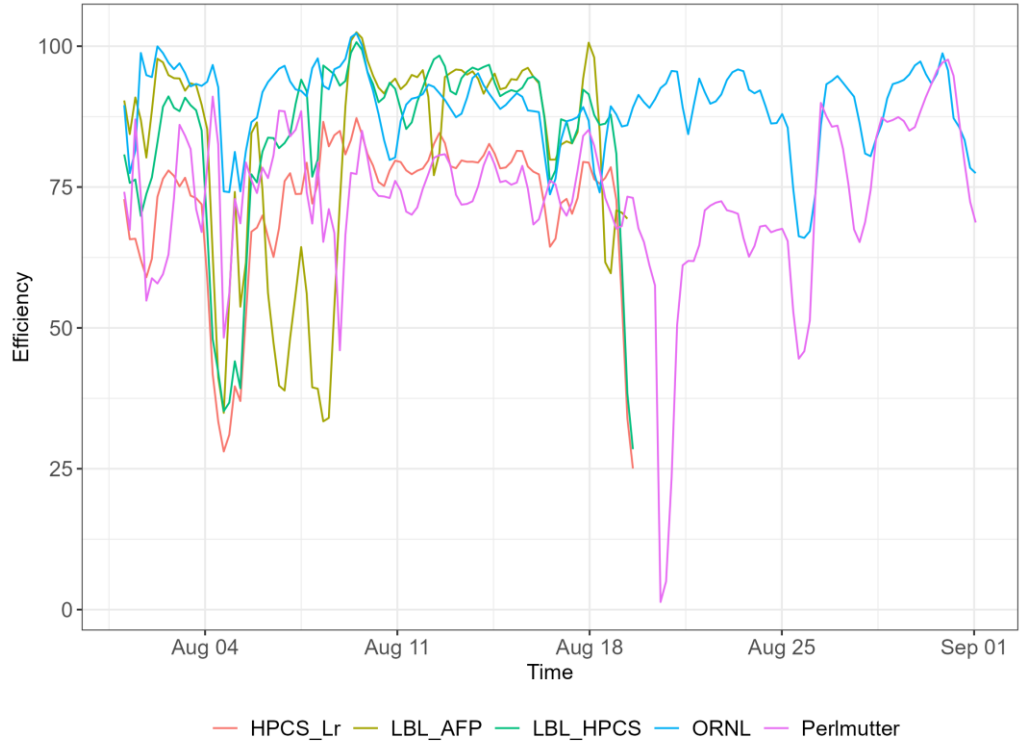

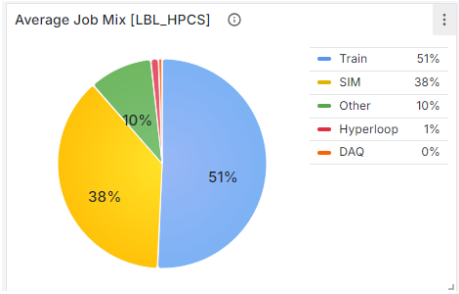

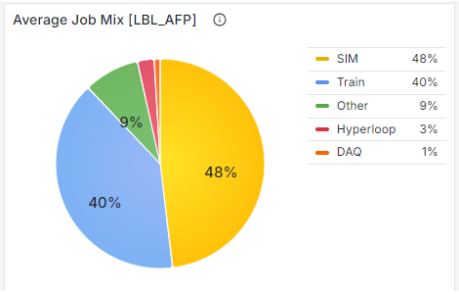

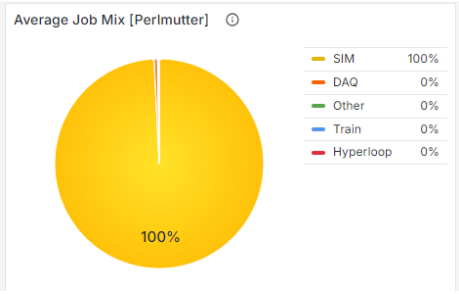

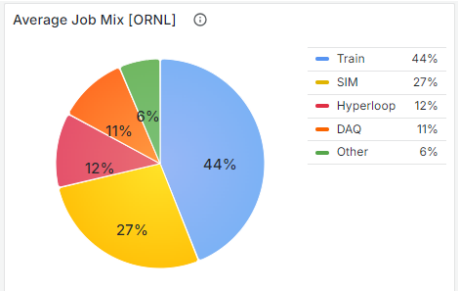

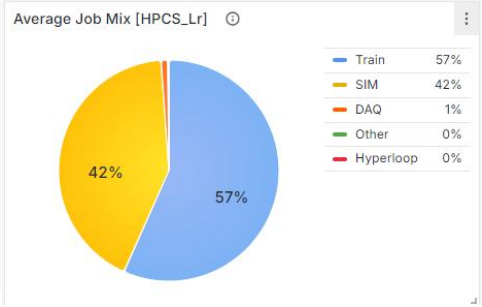

Accounting: August 2025

CPU Obligations	kHS06
ALICE-USA	100.9
LBNL HPCS T2	50.45
ORNL T2	50.45

T2 Site	Monalisa (khr)	WLCG (khr)	%diff
HPCS_Lr	6269.00	23800.00	68.26
LBNL	940.00		
AFP	344.00		
ORNL	2621.00	3250.00	19.35
Total	**10,174.00**	**27,050.00**	**62.39**

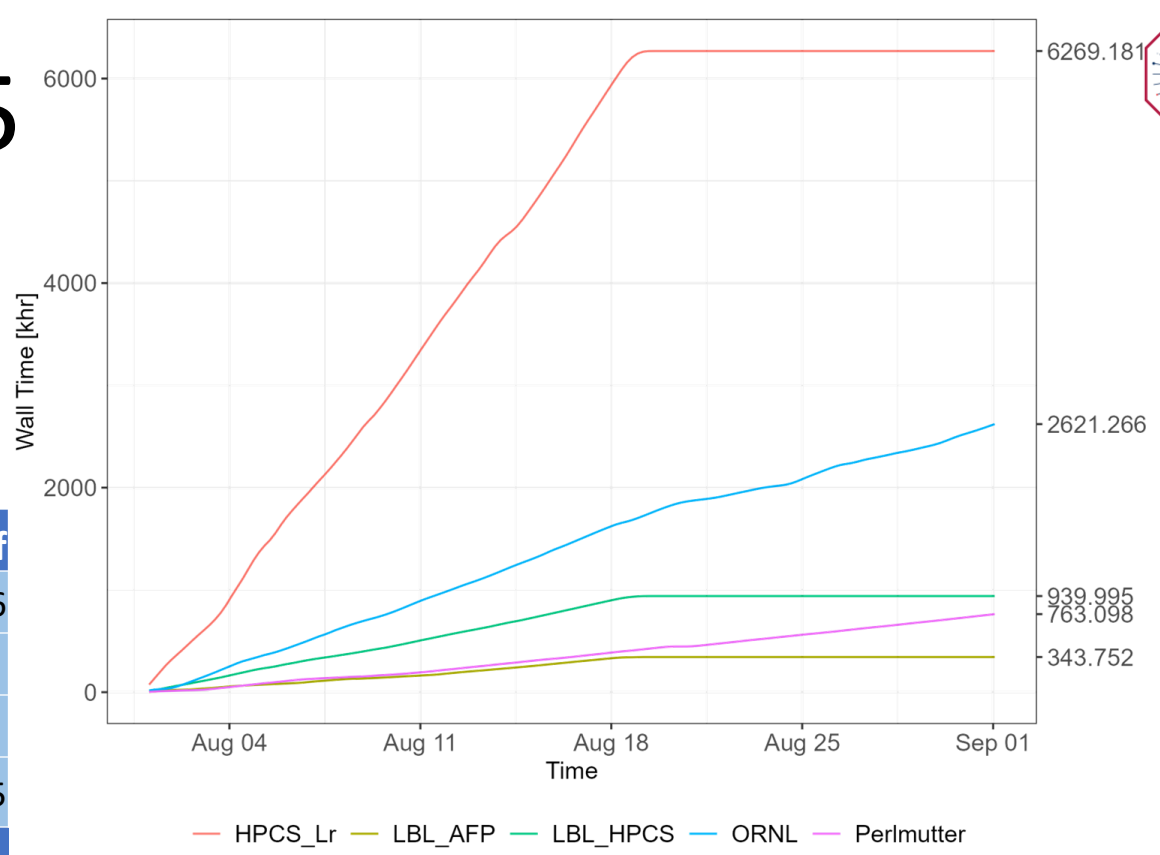

N.B. We do not run accounting on Perlmutter

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	62.69	238.00	144.3	547.8
LBNL	12.0		11.28		26.0	
AFP	12.0		4.13		9.5	
ORNL	11.5	43.45	30.14	37.38	69.4	86.0
Total		**86.90**	**108.24**	**275.38**	**124.56%**	**316.89%**

Accounting: October 2024 to Now

- This period covers 50/50 RRB204 and RRB2025
- Including Lr resources that we had in reporting to OSG we overdelivered both times
- In 2025 plan we put 20 kHS23 from Lr, which was very conservative estimate [*it was on the level of 70 kHS23 of pure payload*]
- In 2026 we keep 20 kHS23 from Lr to cushion the hardware budget a little

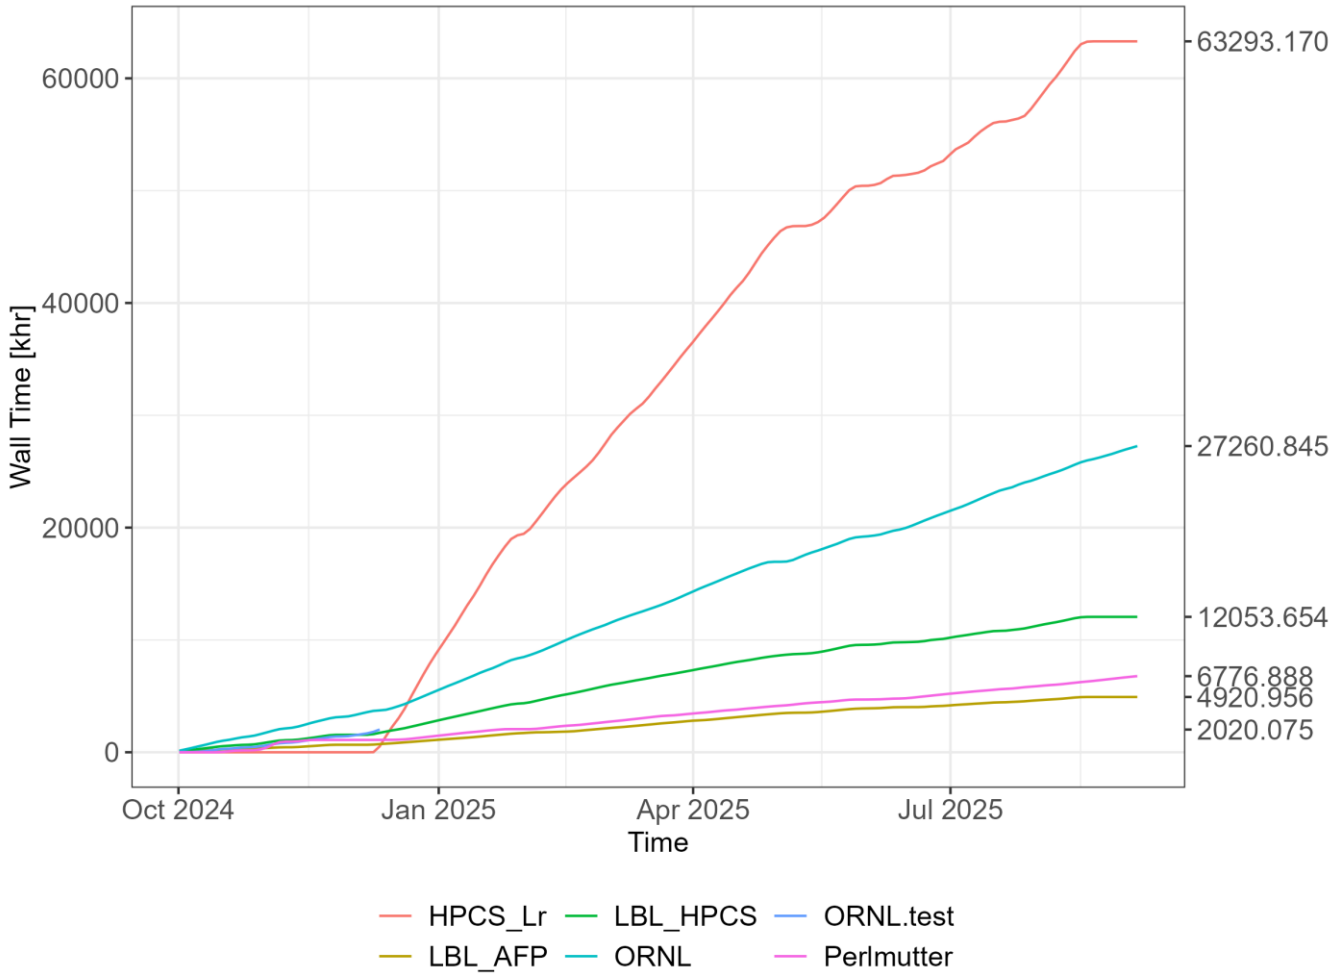

EOS: August 2025

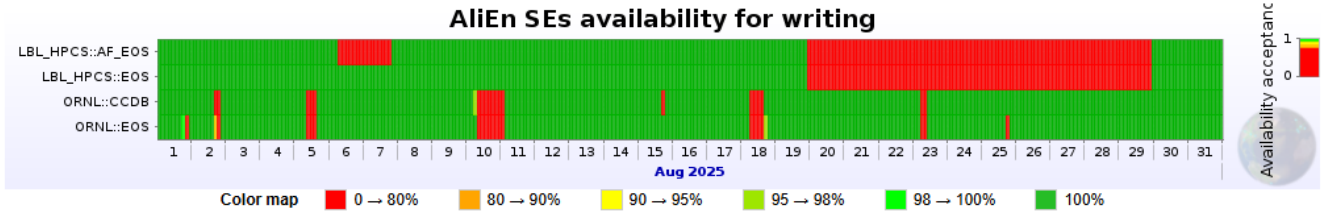

Statistics

Link name	Data		Individual results of writing tests			Overall
	Starts	Ends	Successful	Failed	Success ratio	Availability
LBL_HPCS::AF_EOS	01 Aug 2025 00:10	31 Aug 2025 23:36	473	277	63.07%	63.13%
LBL_HPCS::EOS	01 Aug 2025 00:19	31 Aug 2025 23:45	508	239	68.01%	68.02%
ORNL::CCDB	01 Aug 2025 00:08	31 Aug 2025 23:34	713	40	94.69%	94.71%
ORNL::EOS	01 Aug 2025 00:19	31 Aug 2025 23:46	707	41	94.52%	94.58%

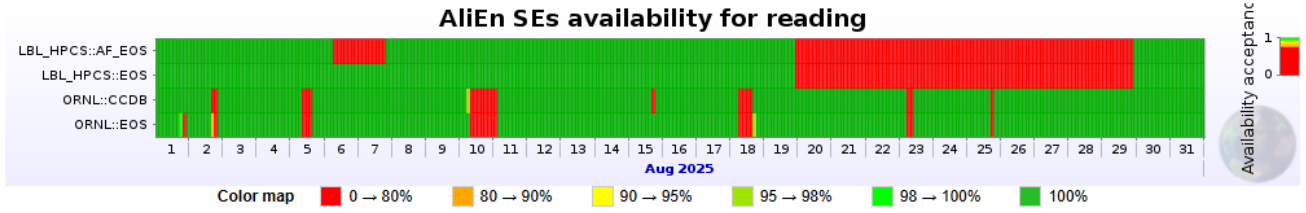

Statistics

Link name	Data		Individual results of reading tests			Overall
	Starts	Ends	Successful	Failed	Success ratio	Availability
LBL_HPCS::AF_EOS	01 Aug 2025 00:10	31 Aug 2025 23:36	473	277	63.07%	63.13%
LBL_HPCS::EOS	01 Aug 2025 00:19	31 Aug 2025 23:45	508	239	68.01%	68.02%
ORNL::CCDB	01 Aug 2025 00:08	31 Aug 2025 23:34	712	41	94.56%	94.58%
ORNL::EOS	01 Aug 2025 00:20	31 Aug 2025 23:46	707	41	94.52%	94.59%

	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 - AF_EOS	ALICE::LBL_HPCS::AF_EOS	2	1.092 PB	937.2 TB	181.4 TB	83.79%	1,010,150	FILE	1.083 PB	951.3 TB	158.2 TB	85.74%	Xrootd 5.7.1	
2. LBL_HPCS - EOS	ALICE::LBL_HPCS::EOS	2	6.313 PB	4.754 PB	1.558 PB	75.32%	86,659,460	FILE	6.295 PB	4.746 PB	1.549 PB	75.39%	Xrootd 5.7.1	
3. ORNL - CCDB	ALICE::ORNL::CCDB	2	100 TB	7.487 TB	92.51 TB	7.487%	19,909,558	FILE	-	-	-	-	Xrootd 5.8.3	
4. ORNL - EOS	ALICE::ORNL::EOS	2	6.586 PB	5.266 PB	1.32 PB	79.96%	66,529,886	FILE	6.586 PB	5.321 PB	1.265 PB	80.79%	Xrootd 5.8.3	5.3.15
Total			**14.09 PB**	**10.94 PB**	**3.146 PB**		**174,109,054**		**13.96 PB**	**11 PB**	**2.969 PB**			

EOS: October 2024 to Now

- SE availability
  - EOS problems at both sites
  - Similarity: mgm needs to restart from time to time
  - Difference: the errors seen are different
  - Action: EOS team is aware of the problem
  - Remedy: Manual/Automatic restarts at ORNL/LBNL
- IPv6 has been up at LBNL since early RRB2021
- IPv6 at ORNL is <u>still</u> looming on a horizon

	READ	WRITE
LBNL	84.6	83.4
ORNL	95.6	95.5

IPv4/IPv6

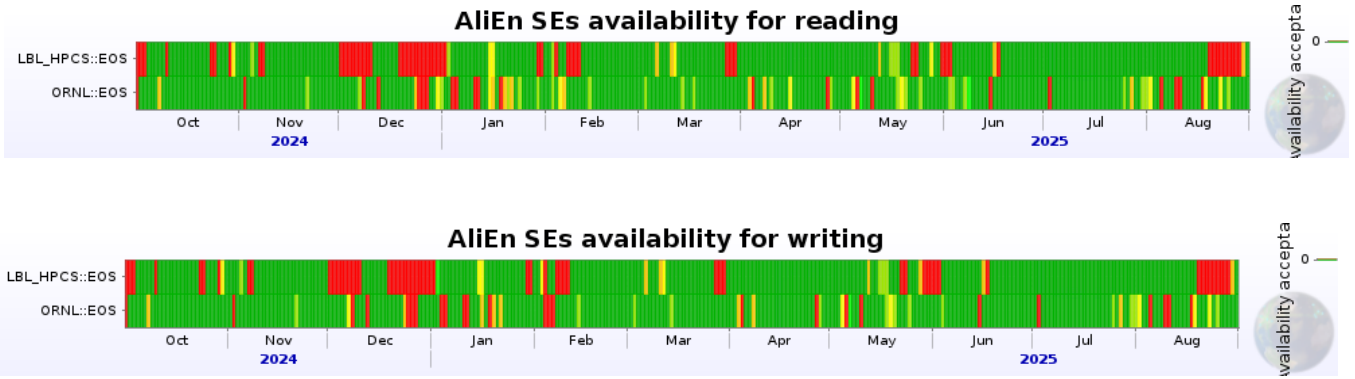

IPv6 Only

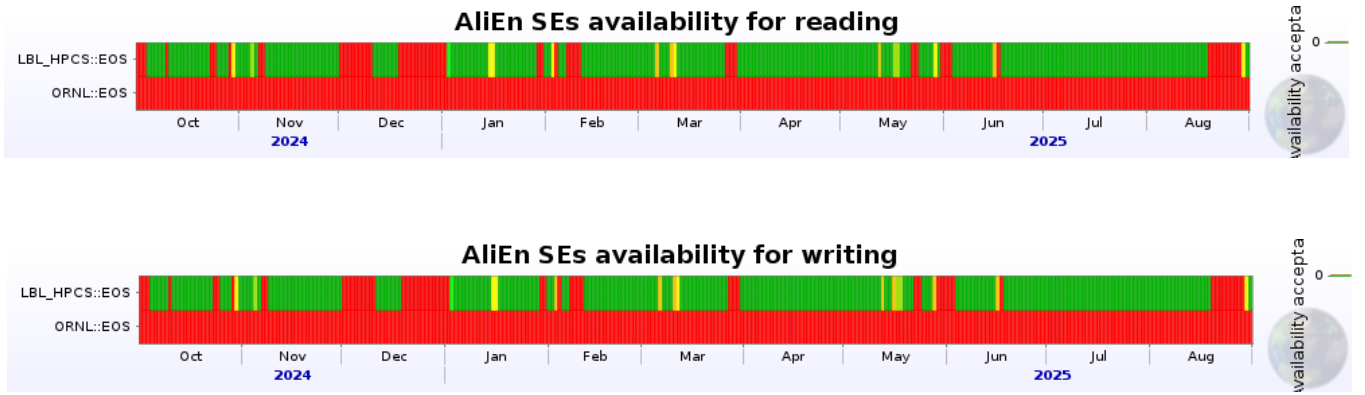

EOS: Success Ratio: August 2025

- We’ve had file accessibility drop to ~50% at both sites at one time or another
- The solution at ORNL came with one of the mgm restarts
- At LBNL we had a defective hardware that was tracked down to a particular drive
- The file access has been stable in past few months now

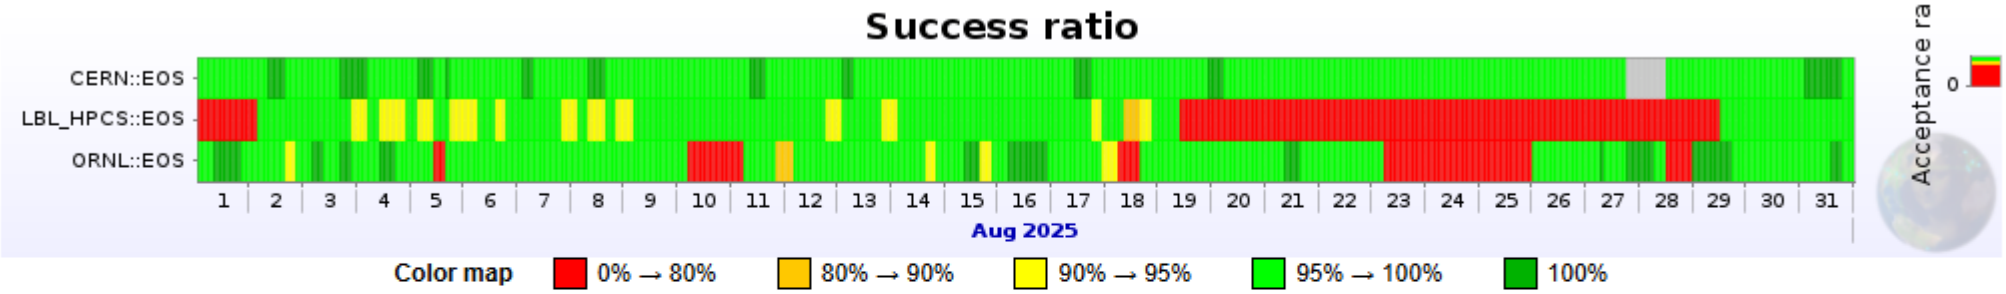

Averaged metrics for the selected interval

SE Name	Start	End	Success ratio	Corrupt ratio	Inaccessible ratio	Internal error ratio
CERN::EOS	01 Aug 2025 03:36	31 Aug 2025 22:06	99.80 %	0.01 %	0.14 %	0.05 %
LBL_HPCS::EOS	01 Aug 2025 03:37	31 Aug 2025 15:45	81.85 %	0.38 %	17.77 %	0.00 %
ORNL::EOS	01 Aug 2025 03:38	31 Aug 2025 22:12	87.92 %	0.10 %	11.96 %	0.01 %

Job Core Mix: October 2024 to Now

- Job core type mix is driven by the optimization of the resources on the nodes
- We also always keep an eye on the fraction of the allocated cores to the total number of cores on the sites

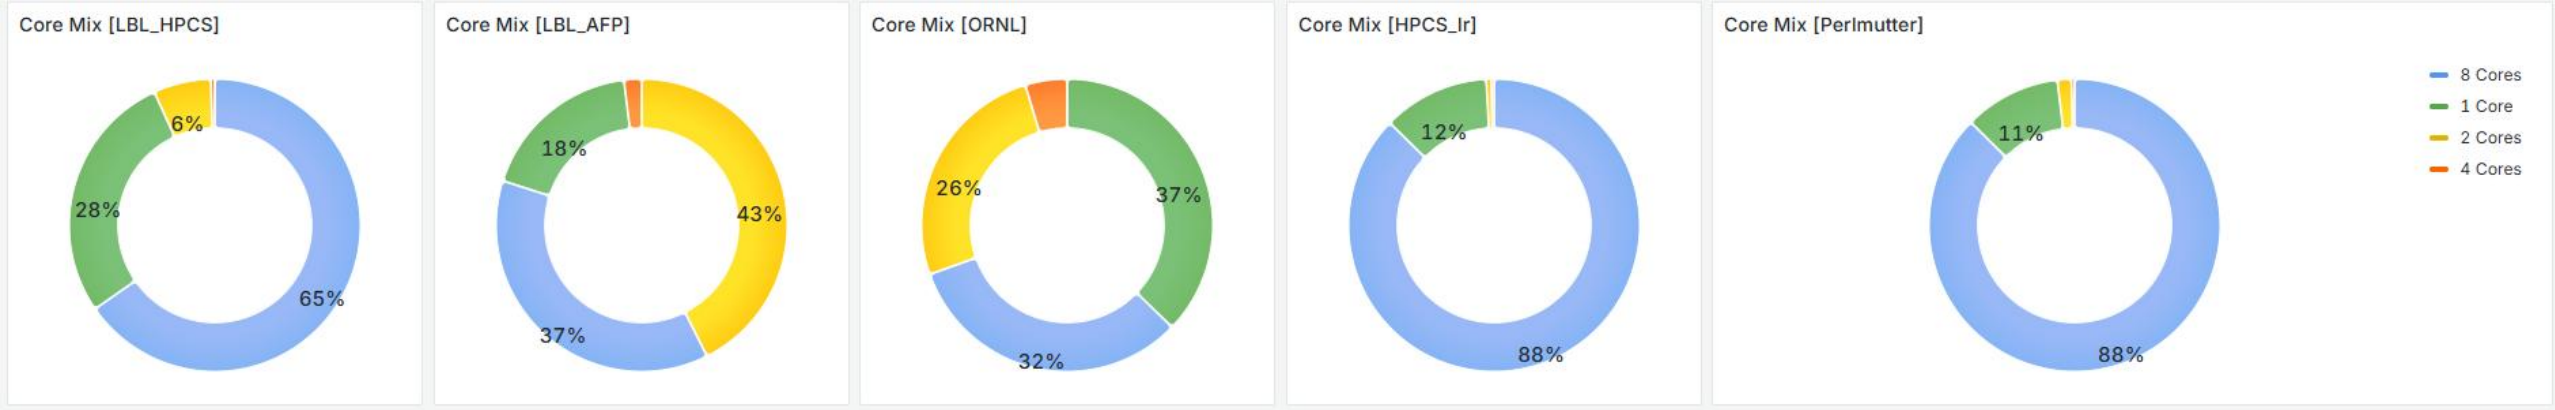

Current Issues

- Immediate issue
  - After we came back from the power work LBL sites seem to not be able to launch agents
- Issues Reported on Alimonitor
  - CVMFS is a little old on every site
  - IPv6 on WNs – has this become more important? If so how pressing?

Site name	VOBox and SE problems	Remarks
HPCS_Lr lrc-vobox.lbl.gov	Networking: No IPv6 public address CVMFS: old CVMFS version detected (2.11.5), please upgrade to 2.13.2.	Low long-term average job efficiency
ORNL vobox-alice.ornl.gov	CVMFS: old CVMFS version detected (2.13.1), please upgrade to 2.13.2.	
LBL_HPCS alice-vobox.lbl.gov	Networking: No IPv6 public address CVMFS: old CVMFS version detected (2.11.5), please upgrade to 2.13.2.	The site it not picking up jobs.
Perlmutter alice-sfapi.lbl.gov	Networking: low buffer size kernel parameters: tcp_wmem_max=4 MB only Networking: No IPv6 public address CVMFS: old CVMFS version detected (2.10.0), please upgrade to 2.13.2.	
LBL_AFP alice-afp.lbl.gov	Networking: No IPv6 public address CVMFS: old CVMFS version detected (2.11.5), please upgrade to 2.13.2.	
5 sites	**10 issues**	**2 remarks**

Some other issues/TODOs we are tackling

- Job cancellations at ORNL
  - We've tracked this down to old jobs being resubmitted over and over just to be cancelled due to their age
  - Needs final confirmation and removal of the resubmission flag
- LBNL
  - Perlmutter job log cleanup is still manual, needs to be automated

HTCondor

ALICE

- Both sites and all the voboxes are now working in the HTCondor CE mode
- This ensures OSG / WLCG reporting
- Only VObox / site not in the reporting is `alice-sfpi.lbl.gov` which launches jobs on Perlmutter. It is not trivial to include it in OSG reporting but also not necessary
- We need to add all the configs into the LBNL automatic provisioning system as after restarts some configurations are not set properly causing some delays

- At ORNL site we are considering dropping HTCondor to SLURM redirect and completely relying on HTCondor.
- Problem: we have a significant discrepancy between Lawrencium reporting to OSG/WLCG and what is seen in Alimonitor
  - This discrepancy does not seem to be easily explained by the whole-node underutilization only
  - I have access to the details reporting of CPU time of every job
  - I need more time to look into this to understand where the discrepancy is coming from

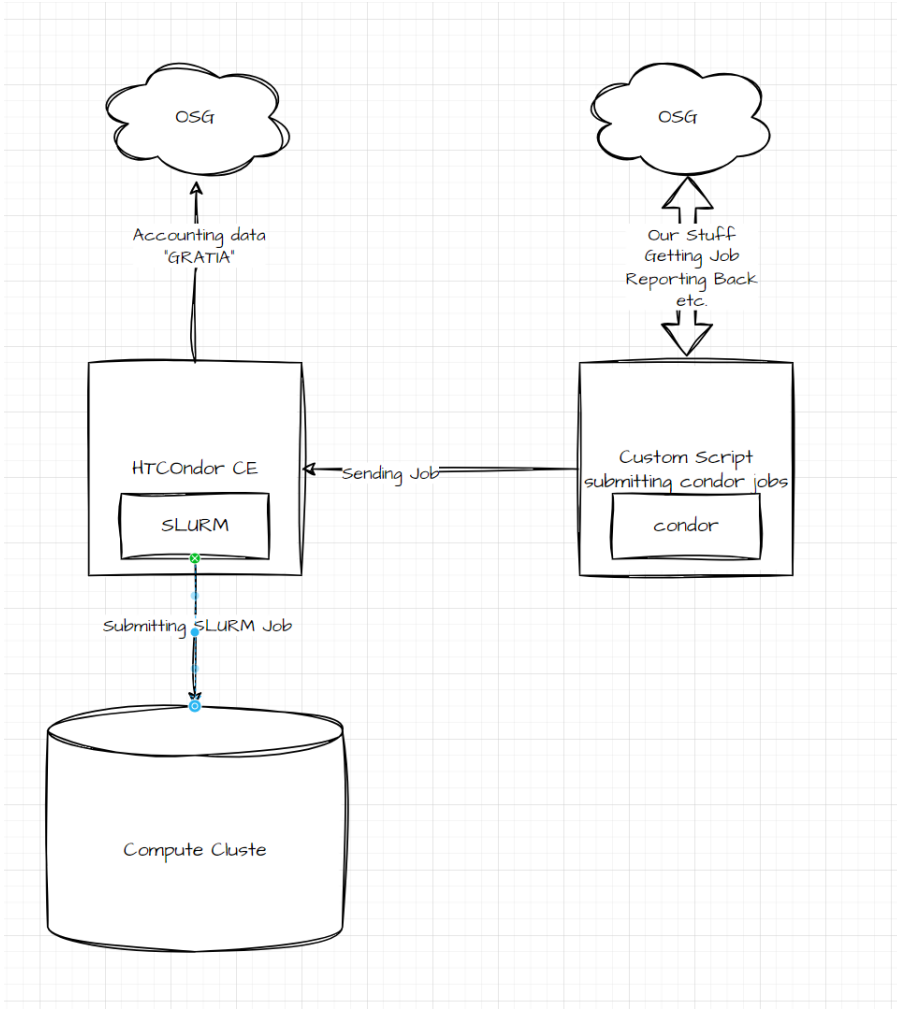

Perlmutter

- This year we switched from “regular” queue to “premium”, which stabilized our node allocation
- In ERCAP 2025 we requested 200,000 node-hours
- We were awarded 110,000
- I think we can request 200,000 node-hours in 2026
- I will also put in a request for GPU-hours

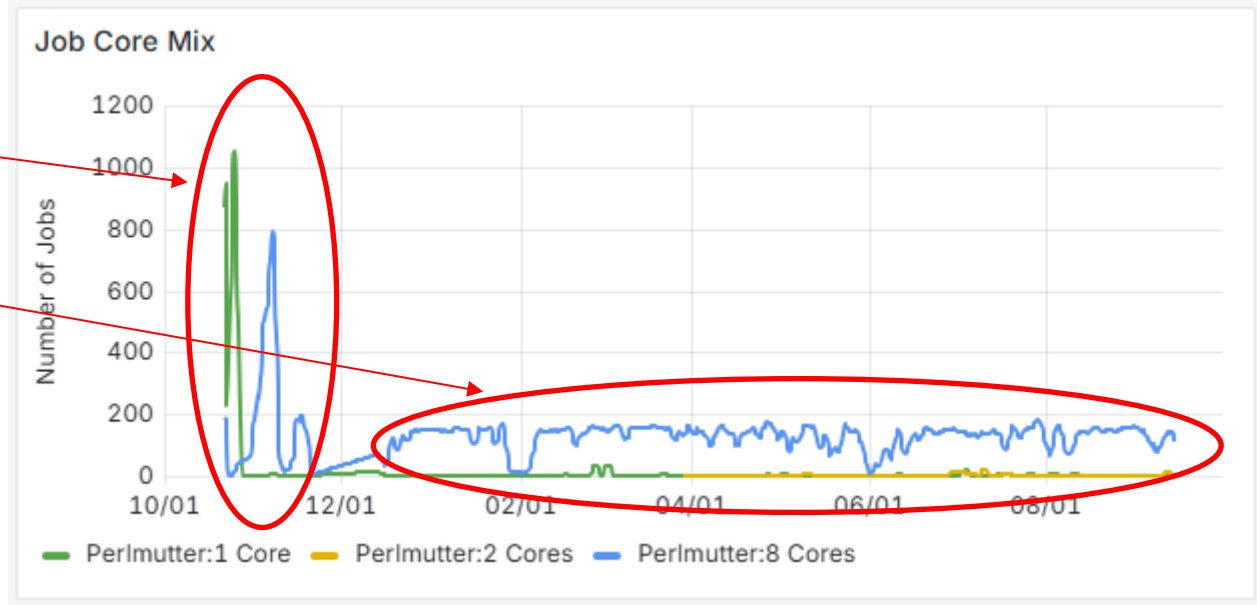

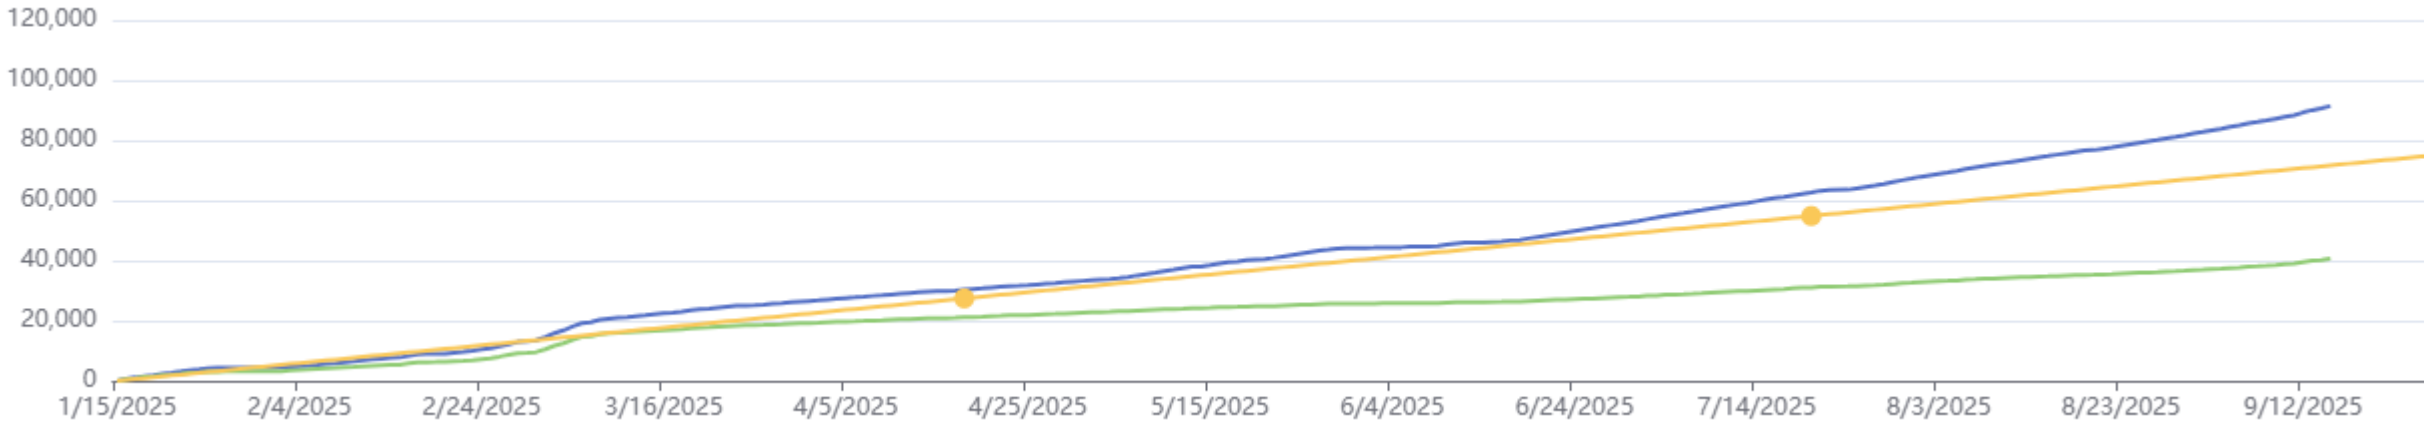

Summary

- The year was challenging and eventful
- We managed to keep sites running with minimal downtime
- We managed to bring up EOS stability considering persisting issues
- Plan for growth and update in the next presentation
- More detailed site reports will be given tomorrow