# EIC computing project at

EIC@LBL Meeting

30 September 2025

Irakli Chakaberia

# NERSC - Perlmutter

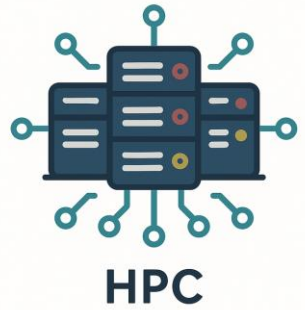


- *“NERSC provides High Performance Computing (HPC) and Storage facilities and support for research sponsored by, and of interest to, the U.S. Department of Energy (DOE) Office of Science (SC).”*

- Advantages:
  - “Free” resource
  - Low latency access from LBL to both data and WNs
  - Direct login.
  - SLURM Batch submission system.
  - SFAPI ready
  - Etc.


- Get an account: https://docs.nersc.gov/accounts/#obtaining-an-account

- Project Name: m3763 (eic)

# CPU Partition


| | |
|---|---|
| **Node Count** | **3,072** CPU-only nodes |
| **CPU per Node** | **2x AMD EPYC™ 7763** (Milan) Processors |
| **Cores per Node** | **128** (64 cores per CPU) |
| **System Memory (DRAM) per Node** | **512 GB** |
| **Interconnect** | **1x HPE Slingshot 11** NIC per node |
| **Primary Workloads** | Capability computing, traditional HPC, codes not easily ported to GPUs. |

# GPU Partition


| | |
|---|---|
| **Node Count** | **1,536** GPU-accelerated nodes |
| **CPU per Node** | **1x AMD EPYC™ 7763** (Milan) Processor |
| **Cores per Node** | **64** |
| **GPU per Node** | **4x NVIDIA A100 Tensor Core GPUs** (40GB or 80GB HBM2/HBM2e) |
| **GPU Interconnect** | **NVLink-3** connects the four A100 GPUs directly |
| **System Memory (DRAM) per Node** | **256 GB** |
| **Interconnect** | **4x HPE Slingshot 11** NICs per node |
| **Primary Workloads** | AI/ML, data analysis, and highly parallel, modernized HPC applications. |

# Energy Research Computing Allocations Process - ERCAP

- Annual allocation of resources at NERSC's HPC
  - Resources are allocated on an annual ERCAP proposal bases
  - Perlmutter is a valuable resource and have many users
  - These resources come for "free" to the group but are in high demand/competition
  - Anybody in this group can acquire NERSC account and use these resources for the EIC related work [project name "m3763"]
  - Project currently has 33 (14 in 2022) users

- Deadline for 2026 requests is October 6, 2025
  - The draft is ready
  - Feedback until this Friday
  - I will be submitting ERCAP over the weekend

# 2025 Resource Usage

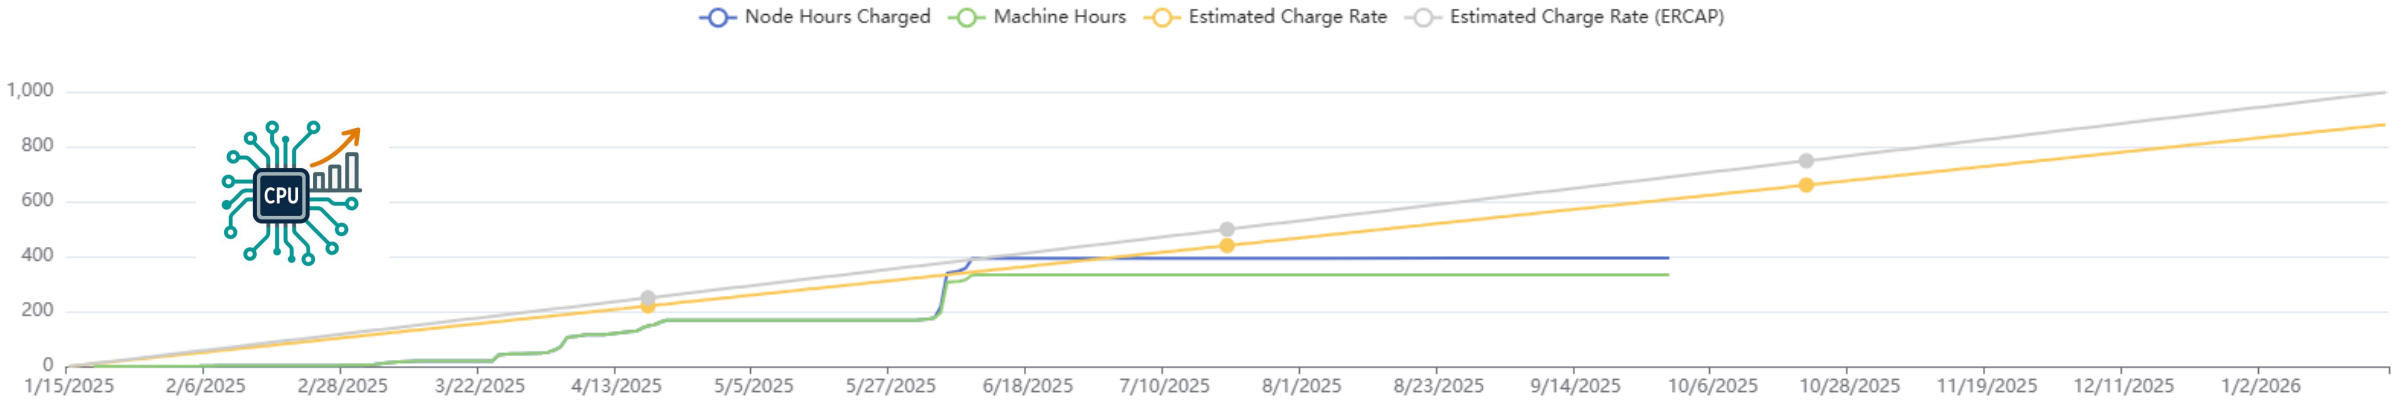


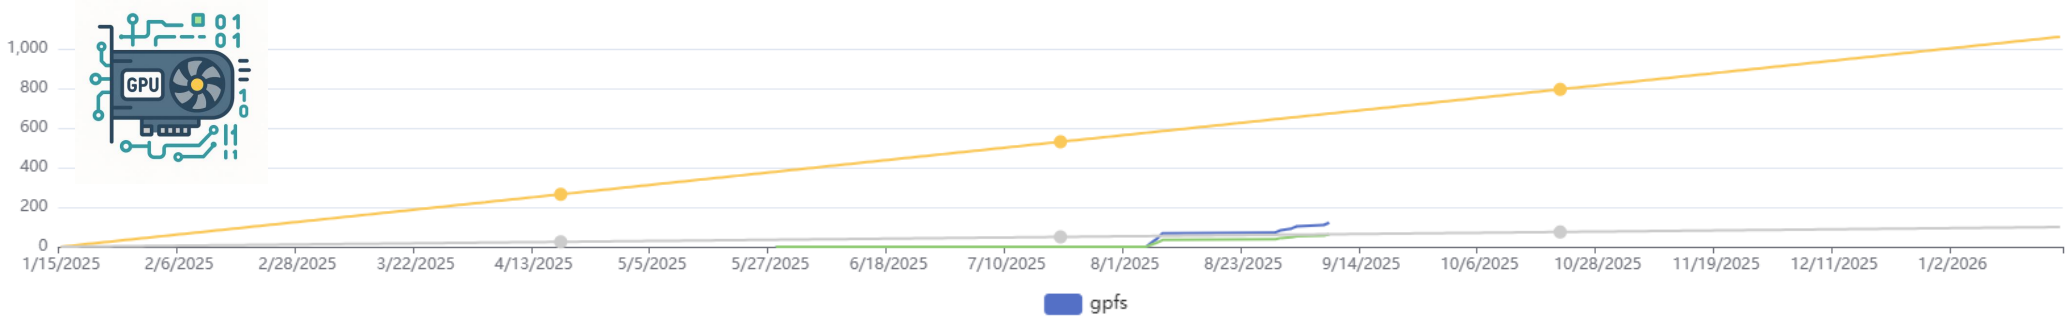


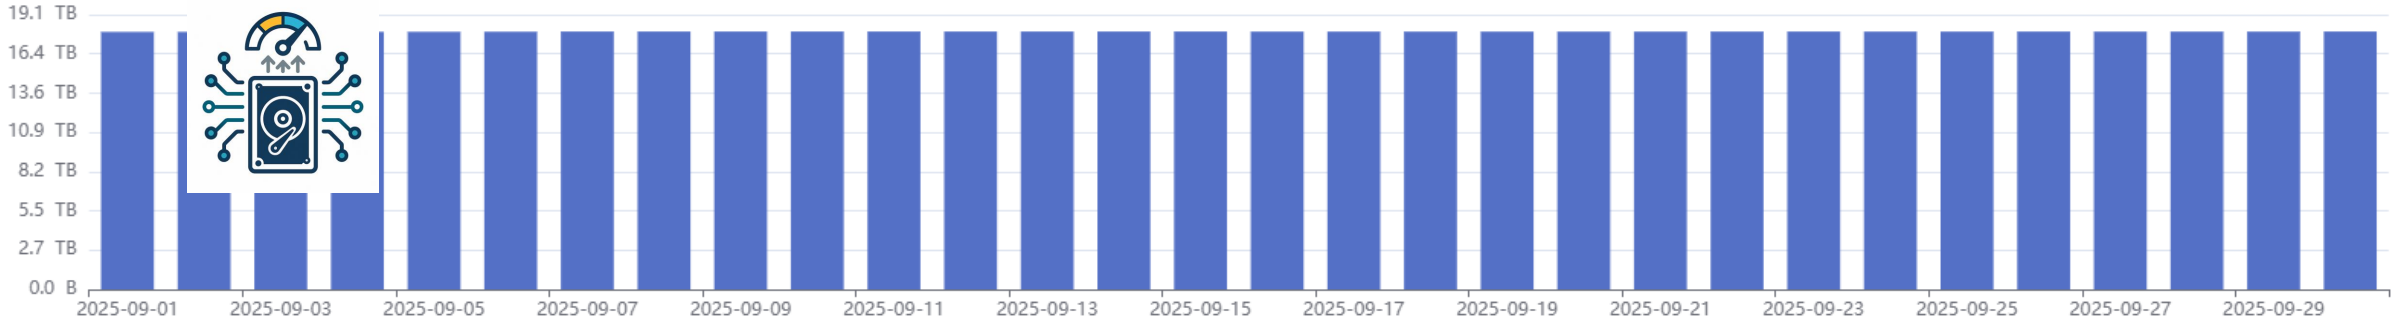

# ERCAP Allocation History

- EIC Project at NERSC started in 2021, with CORI HPC
- ERCAP allocations are based on
  - Justification of need
  - Past utilization
- During these years, more users have deployed their workflow to NERSC, but overall usage remains low
- This year we saw tests of workflow deployment on GPU
- Storage use has grown proportional to the simulations performed
- HPSS is not used as we have sufficient space on disk

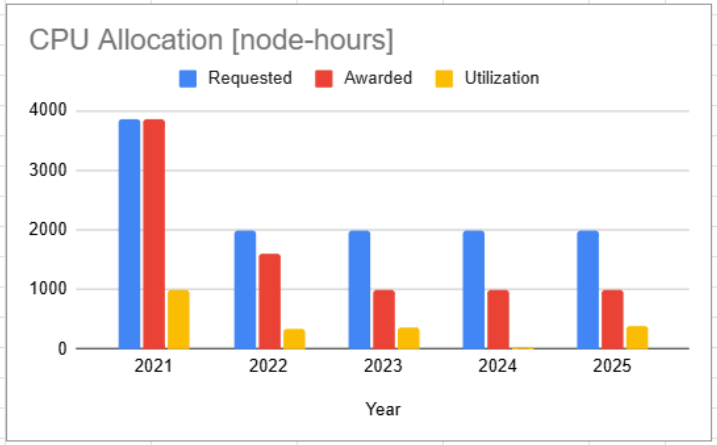


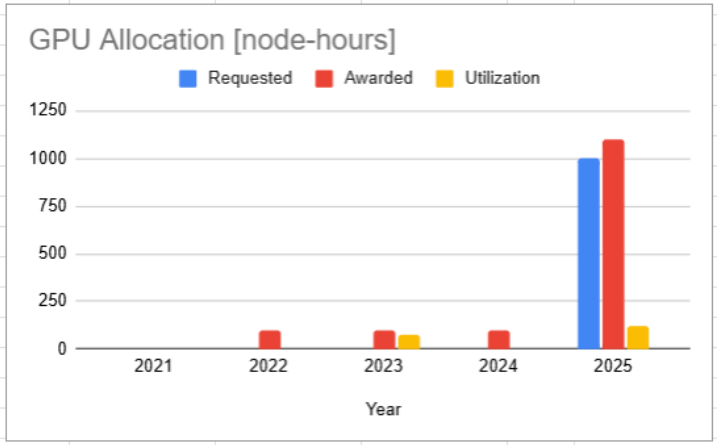


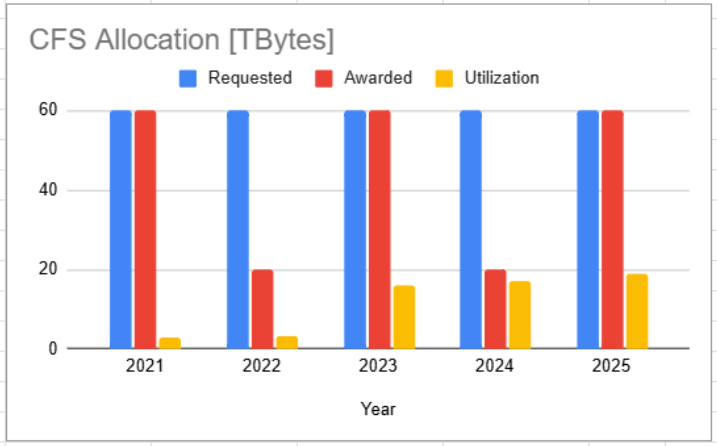


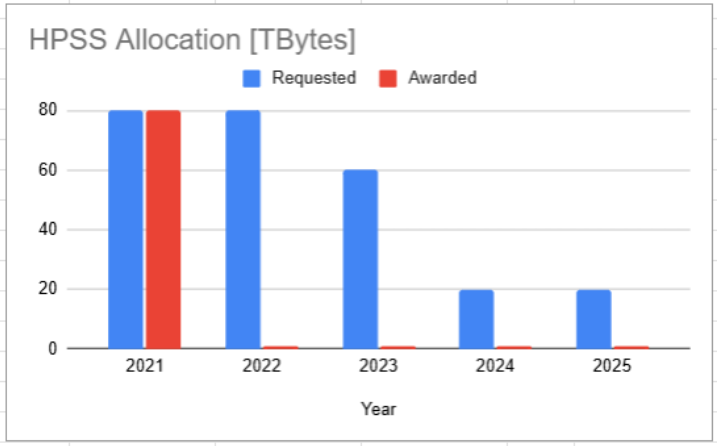

# 2026 ERCAP Justification

# CPU Allocation

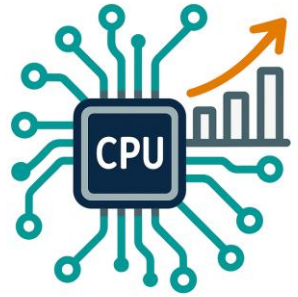


- Deployed [*Amir*] first scalable workflow at NERSC: https://epic-benchmark-lib.readthedocs.io/en/latest/
  - Used up 400 node-hours of CPU at NERSC
- Simulation package for the synchrotron radiation [*Andrii*] studies at EIC was developed
  - Tested on the login nodes
- Total request for the above for 2026: 5,000 node-hours
- Established another workflow [*Yu*] to run the test beam data analysis.
  - Using our condo resource on Lawrencium
- Based on initial tests, we estimate a need of 20,000 node-hours for the FY26 test beam data analysis
- In total we request allocation of **25,000** node-hours in 2026

# GPU Allocation

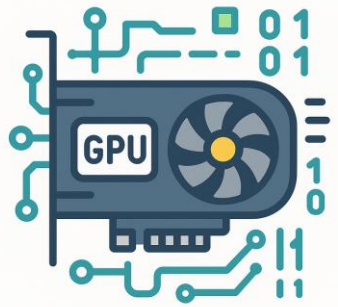


- Amir has tested the geometry simulation workflow on GPUs
- Naively, if all underlying packages are GPU portable, this is a “perfect” avenue to GPU acceleration
- These tests consumed entire “default” allocation of 100 node-hours
- We requested and were awarded another 1000 node-hours
- Plan is to continue testing
- For 2026: request **1000** GPU node-hours to continue porting of the simulation packages to GPU

# CFS Storage Allocation

- The FY26 test beam data analysis will be the biggest consumer of the CFS disk storage space
- We estimate the need of about 50 TB to host this data.
- Current usage: 15 TB of storage space for the simulated data
  - Assuming another 15 TB in 2026
- In total we would like to request **80** TB of CFS storage space

# HPSS Storage Allocation


- With the amount of simulation data this project plans to produce / analyze, HPSS is not critical

- However, in case we would like to vacate potential 50 TB of test beam data to make way for a new dataset I requested **50** TB of HPSS storage

- This is insignificant amount of storage for NERSC, and even if we are not allocated it right away can always be added during the need

# Summary

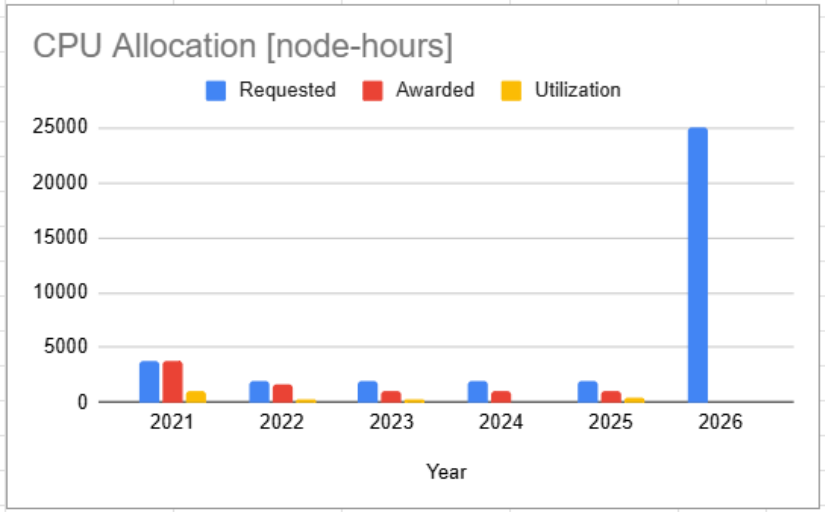
CPU Allocation [node-hours]
Requested
Awarded
Utilization
25000
20000
15000
10000
5000
0
2021
2022
2023
2024
2025
2026
Year

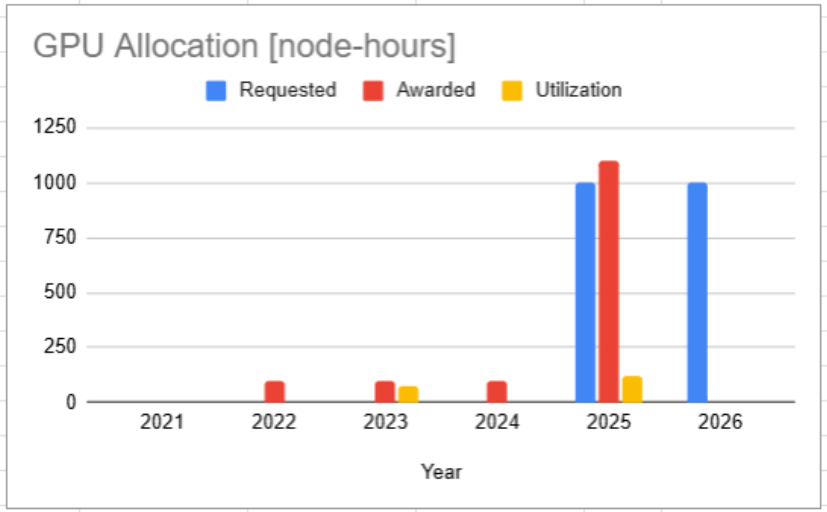
GPU Allocation [node-hours]
Requested
Awarded
Utilization
1250
1000
750
500
250
0
2021
2022
2023
2024
2025
2026
Year

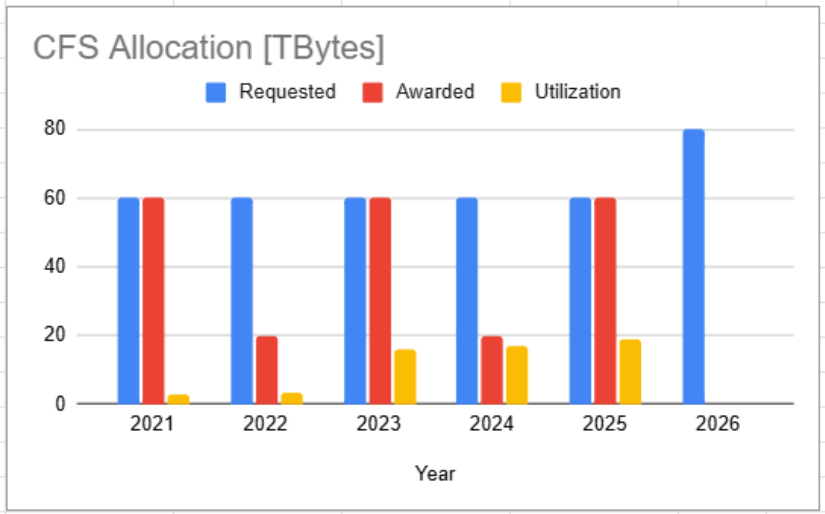
CFS Allocation [TBytes]
Requested
Awarded
Utilization
80
60
40
20
0
2021
2022
2023
2024
2025
2026
Year

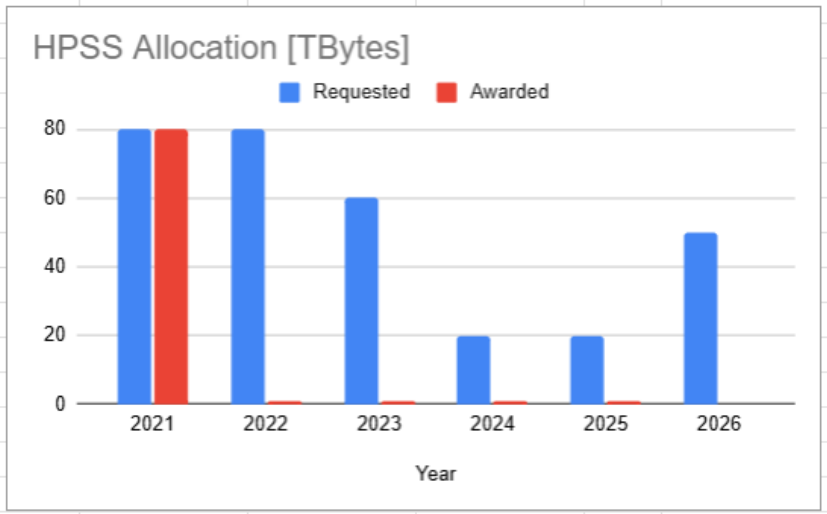
HPSS Allocation [TBytes]
Requested
Awarded
80
60
40
20
0
2021
2022
2023
2024
2025
2026
Year

# Outlook

- The next flagship supercomputer at the NERSC – Doudna
  - In honor of Nobel Laureate and Berkeley Lab biochemist Dr. Jennifer Doudna
  - Slated for arrival in **2026:** https://www.nersc.gov/news-and-events/news/doe-announces-new-supercomputer-powered-by-dell-and-nvidia-to-speed-scientific-discovery
  - Announced by DOE this May
- Built by Dell Technologies
- Powered by **NVIDIA's next-generation Vera Rubin platform**
- Interconnected by Quantum-X800 InfiniBand
- The system is designed to provide over **ten times the performance of the current Perlmutter machine:**
  - extreme-scale simulation
  - AI training