




U.S. DEPARTMENT OF
ENERGY

Office of Science



ALICE-USA Computing Project PEAP 2024 Overview




ALICE-USA Computing Annual Meeting @ ORNL

October 3, 2023


Irakli Chakaberia



ALICE




Project Execution and Acquisition Plan 2024



ALICE-USA CPU and Storage Obligations


- Starting RRB23 accounting is in HepScores [HS23]
- The idea is that conversion between HS06 and HS23 is 1:1
- In 2025
 - Russian resource contribution will need to be shared among other FAs
 - ALICE-USA is expected to grow to 48 M&O-As
 - This puts some strain on our budget/plan
 - It might be prudent to plan ahead for this in 2024

| Year | FY2022 | FY2023 | FY2024 | FY2025 | FY2026 |
|-------------------------|--------|--------|--------|--------|--------|
| ALICE Requirements | | | | | |
| CPU (kHS06/HS23) | 1013 | 1164 | 1280 | 1475 | 1741 |
| Disk (PB) | 104 | 121 | 138 | 160 | 189 |
| ALICE-USA Participation | | | | | |
| (ALICE - CERN) M&O-A | 571 | 571 | 571 | 526 | 526 |
| ALICE-USA M&O-A | 46 | 46 | 45 | 48 | 48 |
| ALICE-USA/ALICE (%) | 8.06% | 8.06% | 7.88% | 9.13% | 9.13% |
| ALICE-USA Obligations | | | | | |
| CPU (kHS06) | 81.6 | 93.8 | 100.9 | 134.6 | 158.8 |
| Disk (PB) | 8.4 | 9.7 | 10.9 | 14.6 | 17.2 |



Resource Acquisition Plan

| Resource | Installed | FY2022 | FY2023 | FY2024 | FY2025 | FY2026 |
|---|-----------|------------|------------|----------|----------|----------|
| LBNL HW & Costs | | | | | | |
| CPU change (+/- kHS06) | | -2.0+10.0 | -7.0+11.0 | -6.0+9.0 | 0.0+18.0 | 0.0+12.0 |
| CPU Installed (kHS06) | 35 | 43 | 47 | 50 | 68 | 80 |
| Disk change (+/- PB) | | 0.0+2.0 | 0.0+0.0 | -1.6+1.4 | 0.0+2.0 | 0.0+1.3 |
| Disk installed (PB) | 2.75 | 4.75 | 4.75 | 4.55 | 6.55 | 7.85 |
| CPU (\$M) | | 0.100 | 0.105 | 0.081 | 0.153 | 0.102 |
| Disk (\$M) | | 0.110 | 0.000 | 0.063 | 0.080 | 0.052 |
| LBL ITD Hardware Maintenance (\$M) | | 0.032 | 0.032 | 0.053 | 0.053 | 0.053 |
| Project Management, Grid Operations, & Site Engineering | | 0.302 | 0.323 | 0.331 | 0.339 | 0.367 |
| ORNL HW & Costs | | | | | | |
| CPU change (+/- kHS06) | | -10.0+16.0 | -11.0+17.0 | 0.0+6.0 | 0.0+14.0 | 0.0+12.0 |
| CPU Installed (kHS06) | 35 | 41 | 47 | 53 | 67 | 79 |
| Disk change (+/- PB) | | -2.75+3.0 | 0.0+2.0 | 0.0+1.4 | 0.0+1.6 | 0.0+1.5 |
| Disk installed (PB) | 2.75 | 3.00 | 5.00 | 6.40 | 8.00 | 9.50 |




Currently Installed CPU & Plan


| Site | Purchase Date | Vender | Processor (Intel) | Clock (GHz) | Cores / node | Slots / node | HS06 / slot | HS06 / node | Number of Nodes Purchased | kHS06 | Node names |
|-------------|----------------|-------------|-------------------|-------------|--------------|--------------|-------------|--------------|---------------------------|--------------|----------------------------------|
| PDSF | 2015 Q3 | Cray | E5-2698 v3 | 2.3 | 32 | 40 | 16.7 | 668 | 32 | 21.38 | n0036.alice0-n0067.alice0 |
| CADES | 2015 Q3 | DELL | E5-2640 v2 | 2 | 16 | 27 | 11.8 | 318.6 | 56 | 17.8416 | alice-001 to alice-56 |
| CADES | 2015 | DELL | E5-2650 | 2 | 16 | | | 318.6 | 48 | 15.2928 | alice-060 to alice-108 |
| HPCS | 2017 Q4 | DELL | E5-2680 v4 | 2.4 | 28 | 40 | 14.5 | 578.6 | 4 | 2.31 | n0000-n0003 |
| ORNL | 2018 Q3 | Cray | E5-2695 v4 | 2.3 | 36 | 72 | 12.4 | 890 | 8 | 7.12 | cray011 - cray018 |
| HPCS | 2018 Q4 | DELL | C-6130 | 2.1 | 32 | 36 | 17.0 | 641 | 16 | 10.26 | n0004-n0018 |
| HPCS | 2019 Q1 | DELL | C-6130 | 2.1 | 32 | 36 | 17.0 | 641 | 16 | 10.26 | n0019-n0036 |
| ORNL | 2022 Q1 | S. Mech. | 5220R | 2.2 | 48 | 96 | 13.1 | 1258 | 4 | 5.03 | alice109-alice112 |
| ORNL | 2022 Q4 | S. Mech. | 5318Y | 2.1 | 48 | 96 | 12.5 | 1201 | 16 | 19.21 | |
| HPCS | 2022 Q4 | | 6330 | 2 | 56 | 112 | 10.2 | 1144 | 8 | 9.15 | n0068-n0075 |
| ORNL | 2023 Q4 | | | | | | | 1200 | 12 | 14.4 | |

| Year | 2015 | 2015 | 2018 | 2022 | 2022 | 2023 | Total |
|------|------|------|------|------|------|------|-------|
| ORNL | 18 | 15 | 7 | 5 | 19 | | 64 |
| | | | | | | | |
| Year | 2015 | 2017 | 2018 | 2019 | 2022 | 2023 | Total |
| LBNL | 18 | 2 | 10 | 10 | 9 | | 49 |

CPU Acquisition



- We are gradually retiring out of warranty hardware
- HS06 for the planned CPU acquisition is roughly estimated based on the measurements we made on similar nodes






Storage Acquisition



- Our storage utilization has been on a lighter side
- We keep evaluating the utilization and purchase the new storage based on it and optimized for the available budget





Cost trends for the CPU and Storage


- CPU unit cost per HS06 is based on extrapolating the CPU clock frequency
- We'll measure the new HS23 values as soon as the new hardware is deployed



Delivered Resources by Year

- The result of the above-described plan is shown in this table
- 2025 increase does not look outrageous, but the budget is larger “flat budget”

| Resource | FY2023 | FY2024 | FY2025 | FY2026 |
|------------------------------|--------|--------|--------|--------|
| ALICE-USA Obligations | | | | |
| CPU (kHS06) | 93.8 | 100.9 | 134.6 | 158.8 |
| Disk (PB) | 9.7 | 10.9 | 14.6 | 17.2 |
| ALICE-USA Plan | | | | |
| CPU (kHS06) | 102.0 | 111.0 | 143.0 | 167.0 |
| % CPU obligation | 108.8% | 110.0% | 106.2% | 105.1% |
| Disk (PB) | 9.8 | 11.0 | 14.6 | 17.4 |
| % Disk obligation | 100.0% | 100.7% | 99.7% | 100.7% |




PEAP 2024 Quarterly Milestones – Q1-Q2

Q1 FY2024

- Update Analysis Facility Proposal to DOE
- Restore HTTPS support in EOS at ORNL & LBNL
- ORNL: IPv6 support, upgrade perfsonar system and register with OSG
- Hold Fall project meeting at ORNL
- Receive and review 2024 ERCAP allocations at NERSC
- Switch to OSG HTCondor based job submission for OSG accounting

Q2 FY2024

- Switch to RHEL 8/9 at ORNL enabling ‘cgroupsv2’
- Attend EOS community workshop
- Deploy LBNL/HPCS perfsonar, register with OSG, and upgrade dashboard
- Attend OSG All Hands meeting
- **Initiate the larger usage of NERSC by the ALICE-USA physicists to run jobs with access to T2/AF storage**
- Hold annual ALICE-USA/CERN resource review meeting at LBNL
- Purchase and deploy new CPU for the 2024 RRB year as needed



PEAP 2023 Quarterly Milestones – Q3-Q4

Q3 FY2024

- Switch to RHEL 8/9 at LBNL enabling ‘cgroupsv2’
- Report status of ALICE-USA grid operations at annual ALICE T1/T2 Workshop
- Review updated ALICE requirements for 2025-2026





Q4 FY2024

- Update PEAP for FY2025
- Purchase and deploy remaining disk storage as needed depending on utilization
- Submit FY2025 ERCAP request for NERSC HPC CPU allocation
- Submit 2025 WLCG pledges

NERSC Allocation




- As part of the project, we request allocation on NERSC HPCs
- 2023 CPU initial award was only 12,000 node-hours
- Once workflow was deployed, we used it all in 1 week
 - This was time when Perlmutter was still not highly utilized so we were allocated a lot of nodes
- We requested additional resources and were awarded another 100,000
- Currently we are waiting for Perlmutter to come back with increased CVMFS open file parameter
- If we can use all this within a couple of weeks I might be able to get more during 2023
- GPU resources were requested as well again in 2024



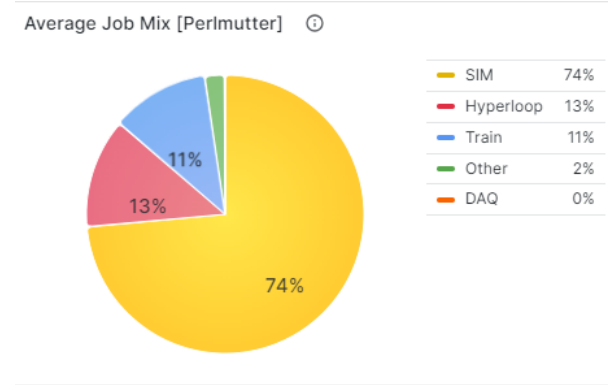

Perlmutter

Job Core Mix [Total CPU Count]






- Running on a VO box on Lawrenceium
- Problems setting up shared account that is able to run systemd
- 300 TB of scratch space
- Currently awaiting CVMFS number of open files increase

CPU Efficiency per Job Type



Core Mix





Summary

- Whole-node scheduling on both sites
- New EOS5 on both sites
- NEW AF[P] deployed
- Lawrenceium contribution is growing
- Perlmutter has become a stable site and should get even better



Questions