

Nuclear & Particle Physics Lattice QCD Computing Initiative

Jefferson Lab

Amitoj Singh

Thursday, April 18, 2024

 Jefferson Lab



Nuclear & Particle Physics Lattice QCD Computing Initiative

NPPLC Initiative started FY18

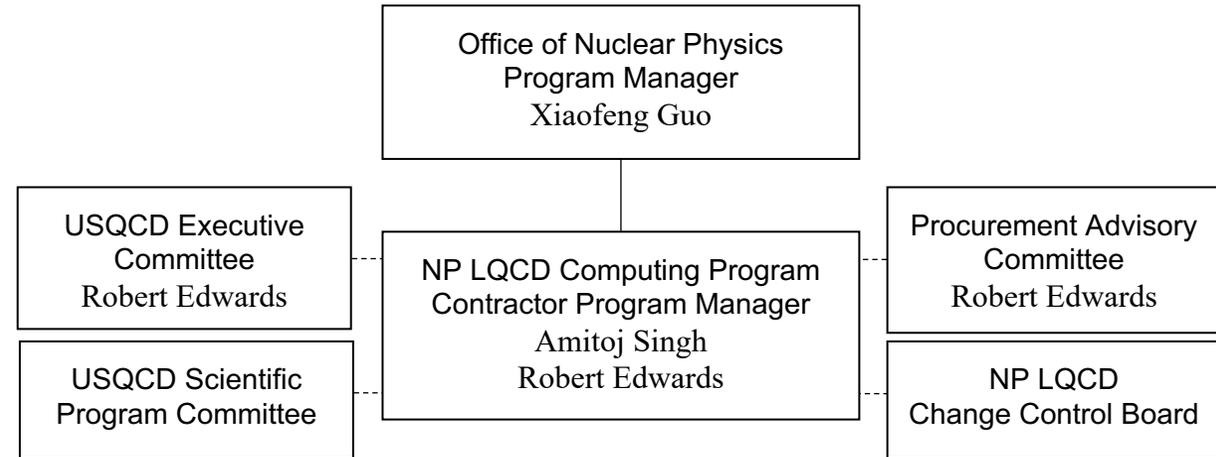
Mission - “*deploy & operate a significant dedicated computing resource for LQCD calculations.*”

NPPLCI organization structure

- Robert Edwards, PI
- Graham Heyes, Deputy
- Amitoj Singh, JLab LQCD site manager
- Edwards & Singh, JLab LQCD site architects

NPPLCI reporting structure

- Report annually to Research Division of DOE Office of Nuclear Physics
- First report (in Feb. 2020) included highlights of scientific program and future plans
- DOE responds to the reports; however, not under an annual panel review process
- NPPLCI 5-year renewal review May 29-30, 2024; this is significant



Vertical lines indicate reporting relationships. Horizontal lines indicate advisory relationships.

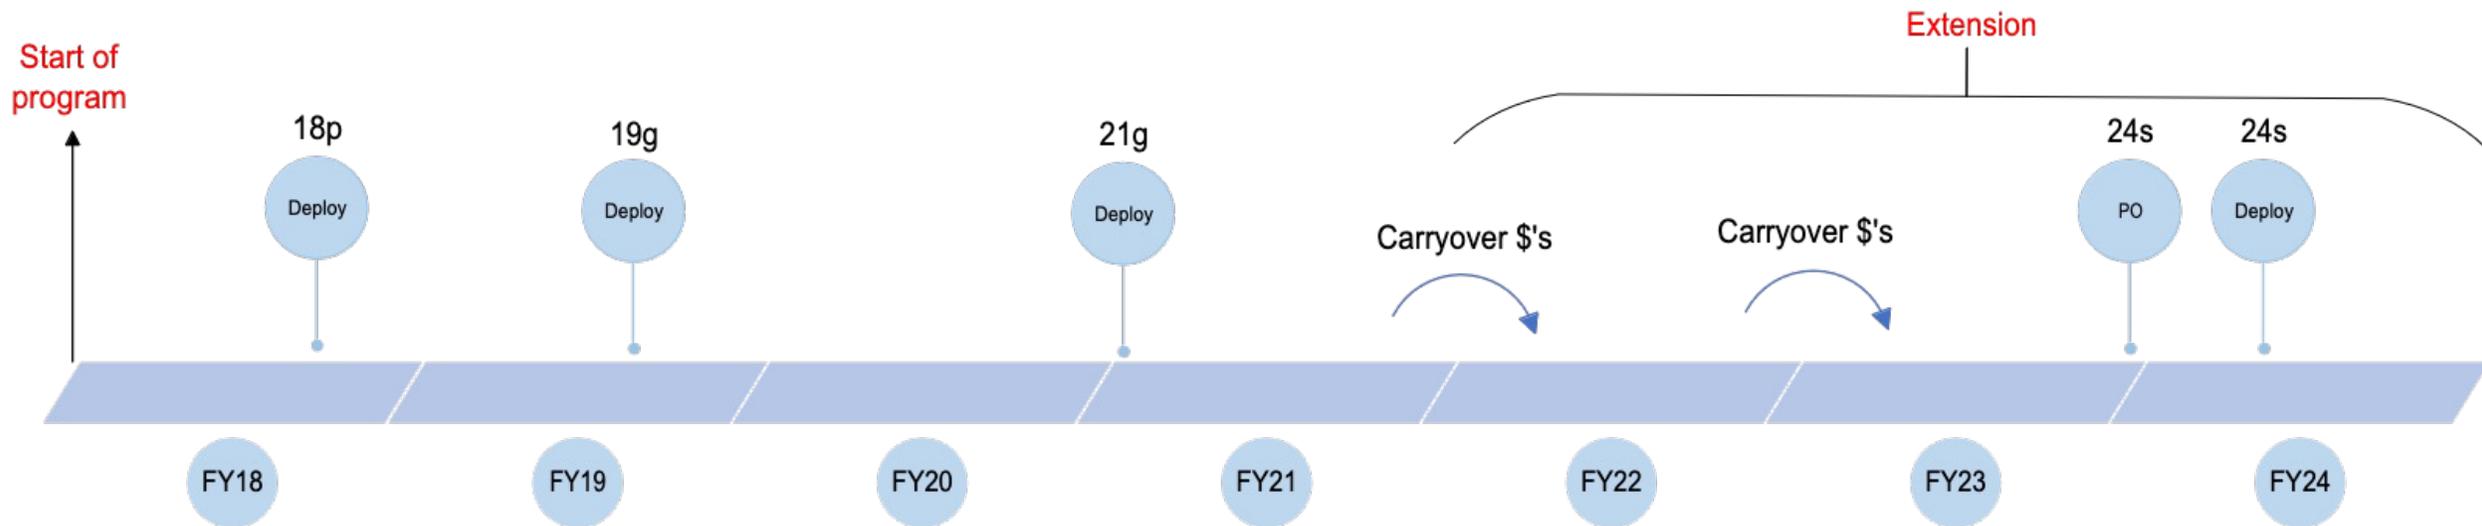
Nuclear & Particle Physics Lattice QCD Computing Initiative

Budget of ~\$1M/year - about 50:50 labor:hardware

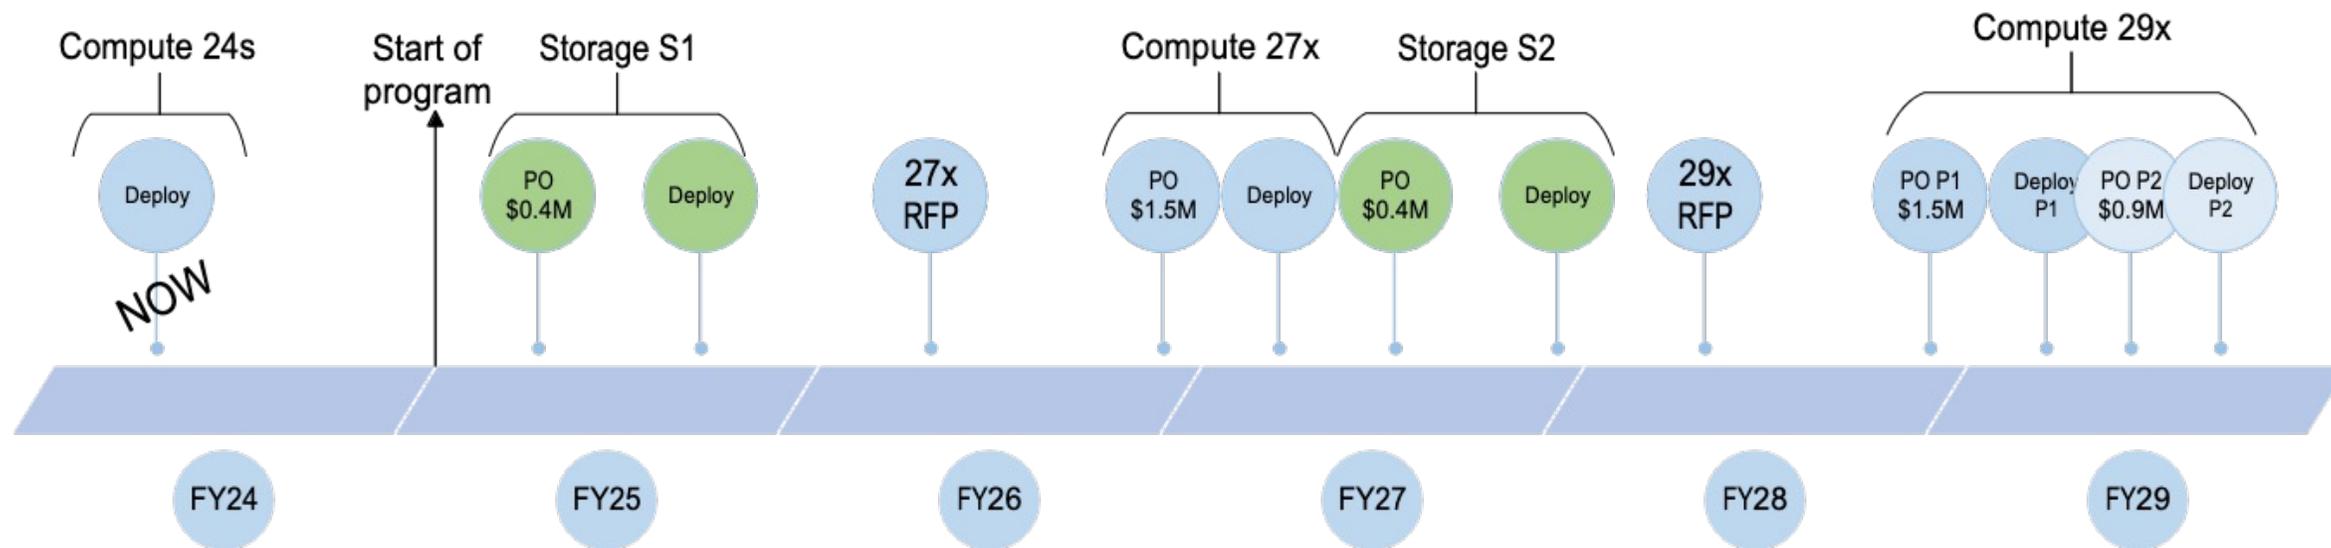
- Labor costs are for 2 FTEs
- Hardware acquisition over FY boundaries for a bigger purchase
- In FY23 we carried over hardware money from FY22 and made a \$1.3M acquisition named "24s"

History of past procurements:

- FY2018: "18p" upgraded Jlab's KNL resources
- FY2019: "19g" upgraded Jlab's GPU resources
- FY2021: "21g" AMD CPU + MI100 GPU system
- FY2024: "24s" new Intel CPU system



NPPLC Initiative – 5-year renewal business plan (subject to approval)

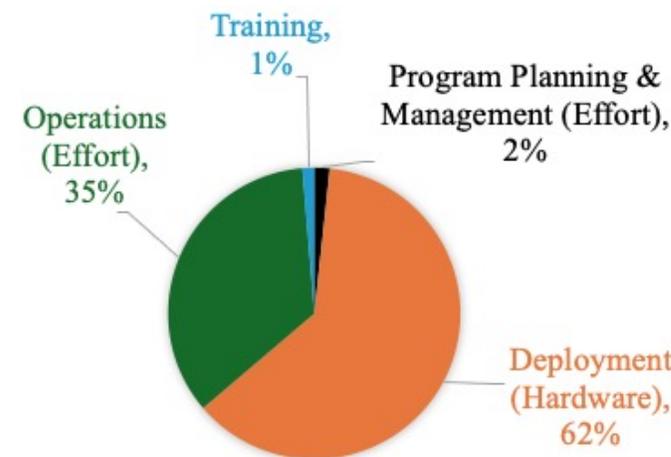


Budget of \$1.5M/year - about 50:50 labor:hardware

- Labor costs are for ~2 FTEs
- Hardware acquisition over FY boundaries for a bigger purchase

System	TFlop
19g	42
21g	63
24s	78
27x *	156
29x *	360

* Extrapolated from 24s. TFlop is based on average of DWF and Clover



Distribution of \$'s across the various funding categories for the program lifetime.

NP-LQCD & JLab-LQCD programs

Computing - Computational Sciences & Technology Division (CST)

- Provides expertise for cluster design, deployment and operations

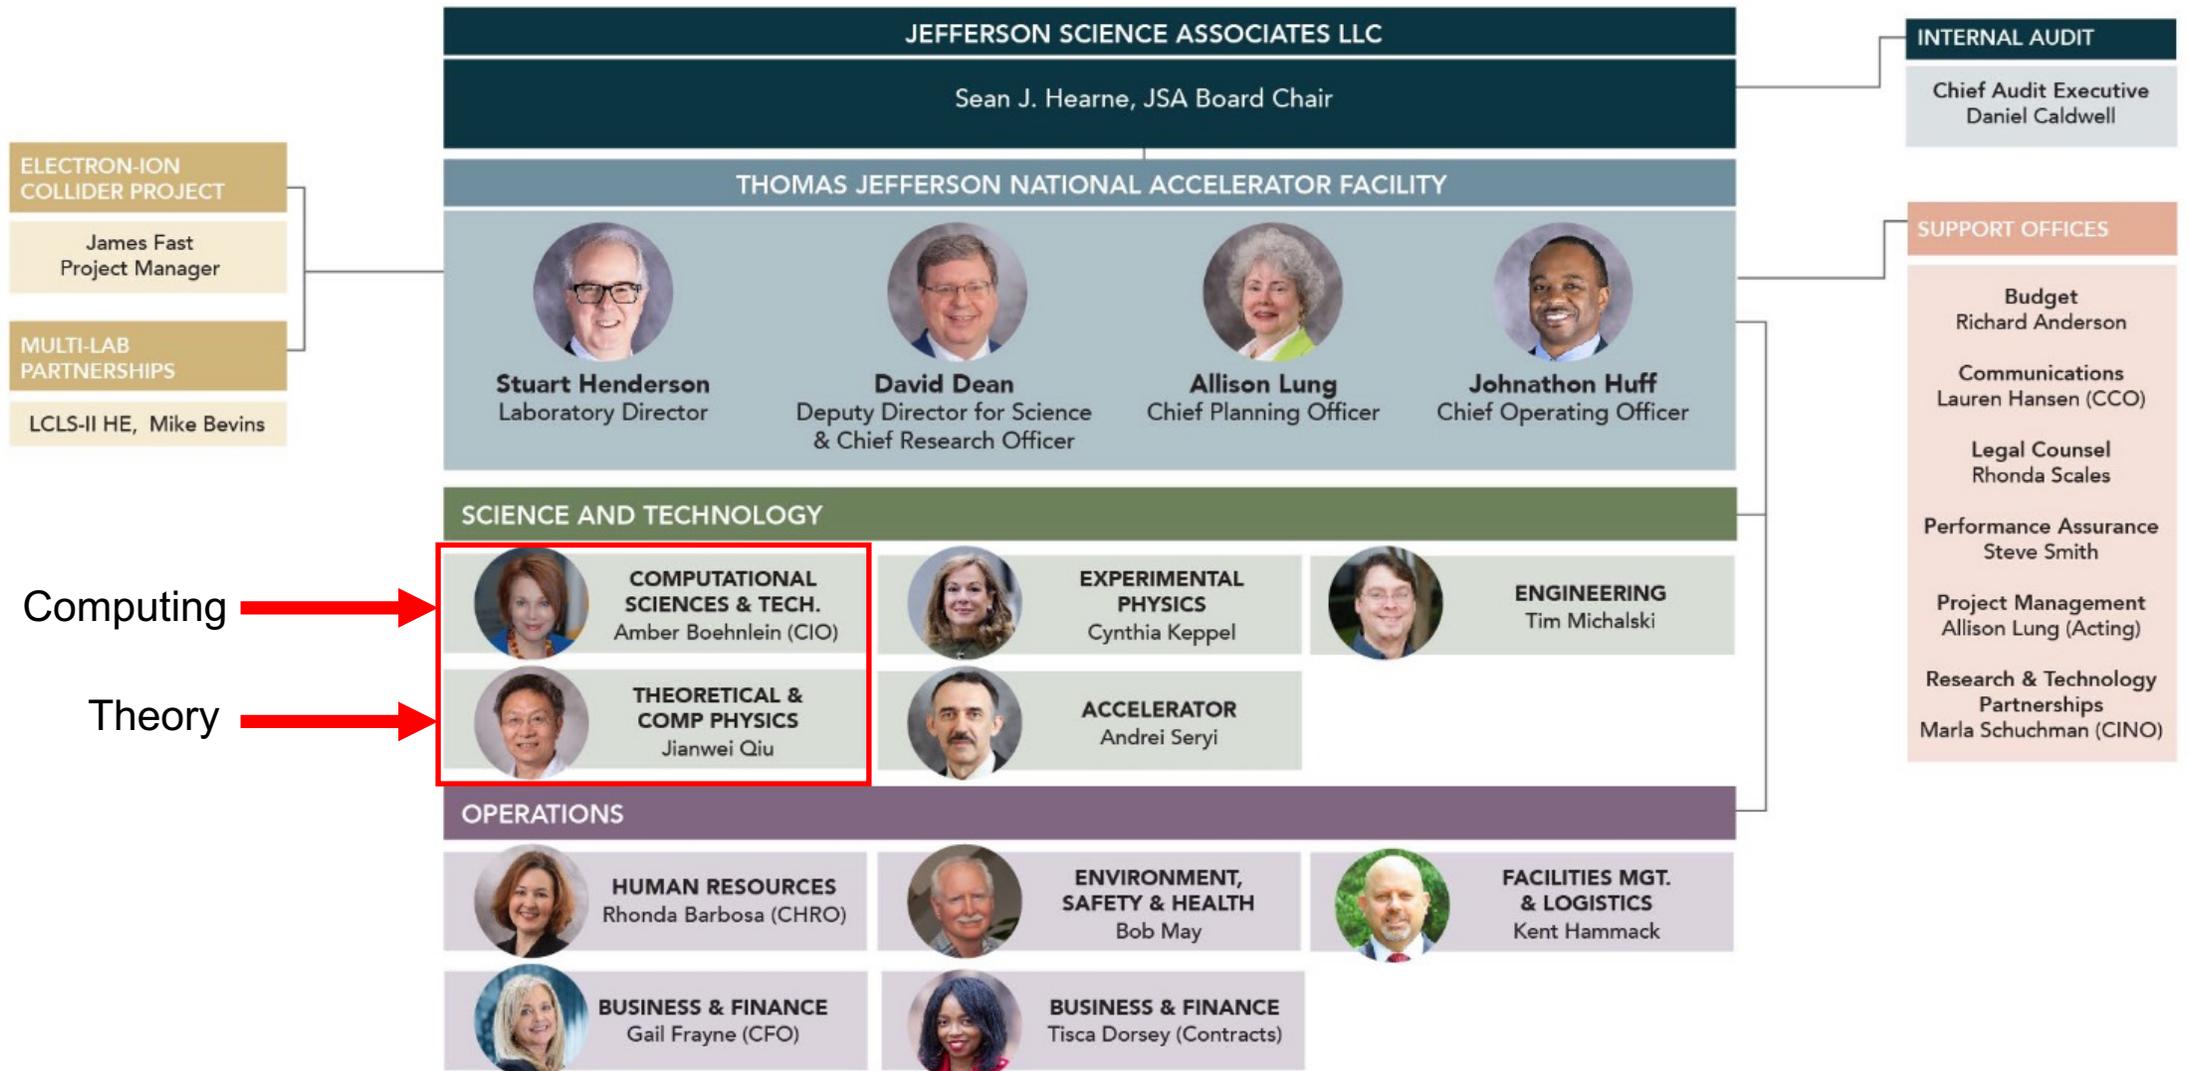
Theory – Theoretical & Computational Physics (Theory Center)

- LQCD@JLab started in ~FY01 - all funds directed through JLab/Theory
- Theory and CST group have worked closely over two decades
- CST division very responsive to the needs of Lattice QCD Theory
 - e.g., CST covers the cost of tape hardware, we only pay for media

Systems - based on computational requirements for Lattice QCD Theory and USQCD

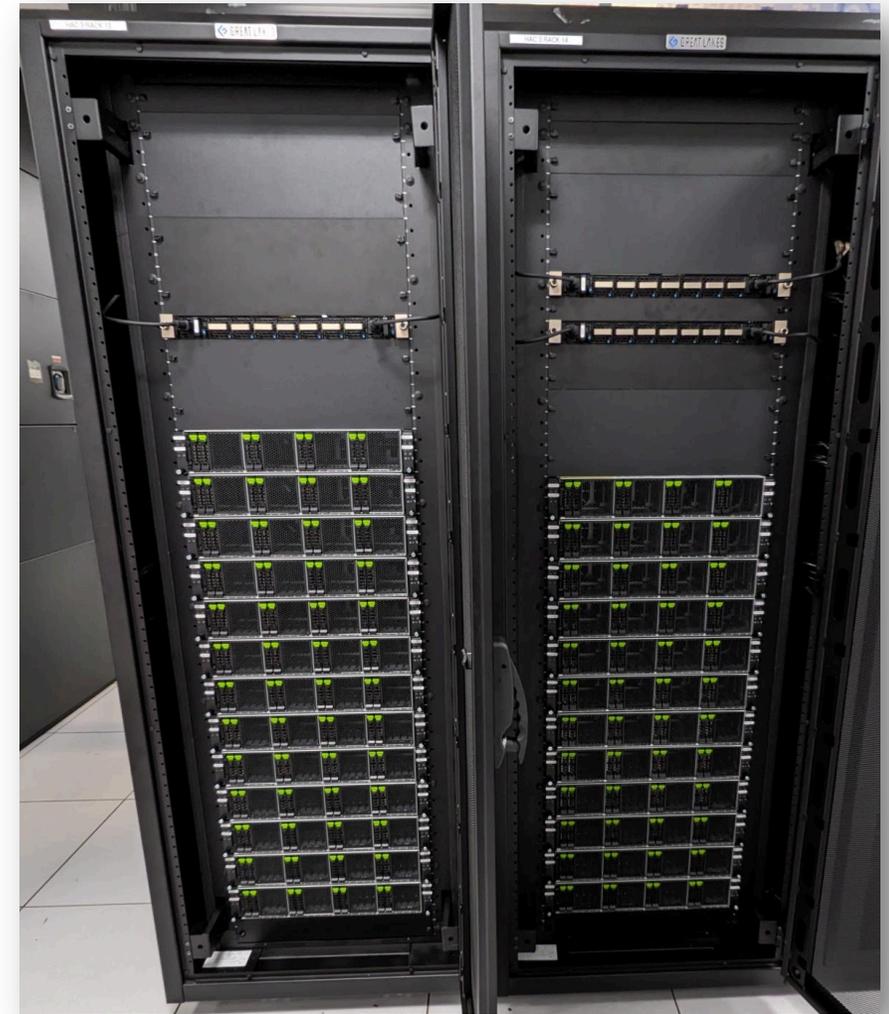
- Funding directly from DOE Office of Nuclear Physics
- Program is coordinated with USQCD Executive Comm. and Scientific Program Comm.
 - e.g., allocations taken from SPC

JLab Organization Chart



New Hardware Acquisition – 24s cluster

- Is a successor to KNLs
- Considered two main metrics to quantify price-performance
 - Wilson dslash - proxy for inverters
 - ZGEMM based contractions considering possible swapping to main memory
- 7 vendor responses to the RFP
- “Best value” evaluation criteria for bid responses
 - Performance/Price (details in next slide)
- Winning bid: Atipa Technologies
- Technical details in my JLab Facility talk



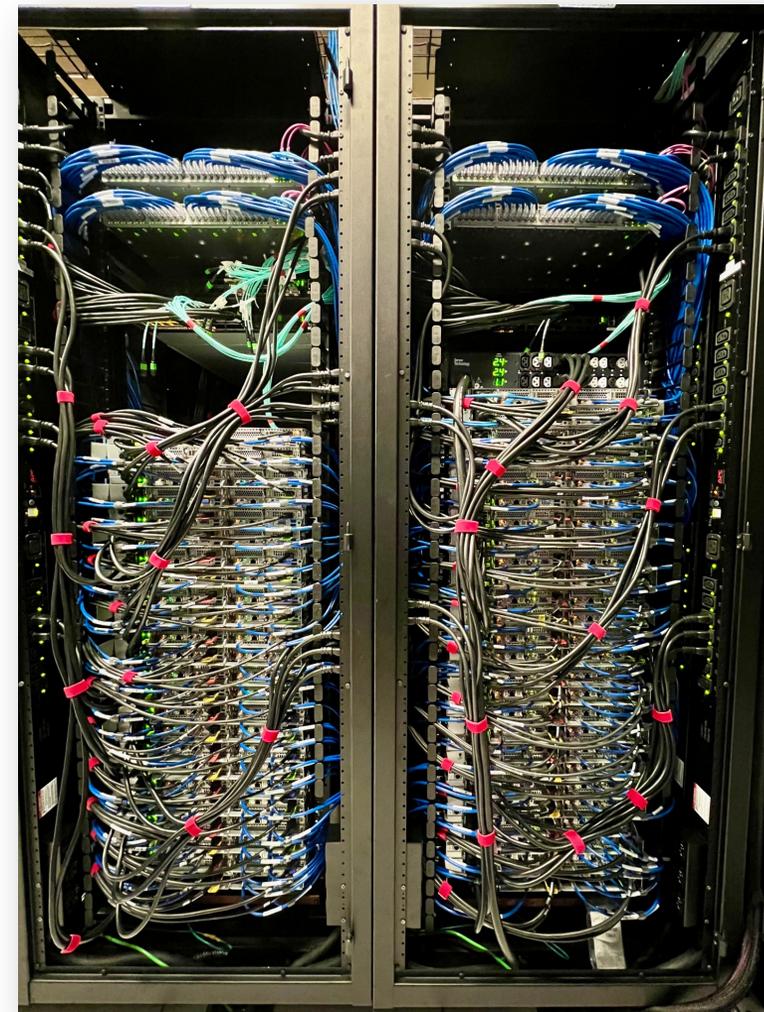
24s-front view

New Hardware Acquisition – 24s cluster

Acquisition Steps	Relative Task Days (Months)	Completion Date
Benchmarking on shortlisted hardware and configuration(s) finalized	-	06/16/23
RFP approvals completed and released for bids	31 (1)	07/17/23
RFP Responses due	61 (2)	08/16/23
PO Awarded – (4 months) Deadline as set in a DOE PEMP * to JLab	105 (4)	09/29/23
Delivery of equipment (5 months post Purchase Order award)	266 (9)	03/08/24
Installation of equipment	284 (9.6)	03/26/24
Burn-in and acceptance testing	319 (11)	04/30/24
Release to “friendly user” access	320 (11)	05/01/24
Release to production (Entire process takes about a year)	381 (13)	07/01/24

* PEMP = Performance Evaluation and Measurement Plan

Processor	Rank / Score			Overall
	Performance (higher is better)	Price (lower is better)	Performance/Price (higher is better)	
AMD 9334 “Genoa” 32-core	0.81	0.88	0.92	3
Intel 8462Y “Sapphire Rapids” 32-core	0.94	0.79	1.19	1
NVIDIA “Grace” 72- core (ARM)	1	1	1	2



24s-rear view

In Conclusion

Amitoj Singh

amitoj@jlab.org

- Jlab's support for NPPLCI is strong.
- JLab theory group and CST have decades of rich experience in LQCD software development and hardware deployments.
- Given our excellent past performance, we are well poised to make a strong case for a 5-year renewal of the program.
- Operations teams continues to make progress on getting 24s ready for production deployment.

