LQCD-ext III Program Management
USQCD All-Hands Meeting

Apr. 18-19, 2024

Alan G. Prosser, LQCD Contractor Program Manager
for the LQCD-ext III Project Team

Program Management Structure and Activities

LQCD Mission Need and Justification

LQCD "supports the mission of the DOE's SC HEP Program to explore and discover the laws of nature as they apply to the basic constituents of matter and the forces between them" (from Justification of Mission Need, Program Execution Plan)

To accomplish this support, the LQCD Program provides:

1. World-class computing hardware facilities and associated infrastructure to deliver reliable resources for the scientific mission

2. Computing professionals to plan, design, deploy, operate, and maintain these resources

The computing resources are located at and managed by:

Brookhaven National Laboratory (HEP)
Fermi National Accelerator Laboratory (HEP)
Thomas Jefferson National Accelerator Facility (NP)

LQCD Management Organization (Integrated Program Team)*

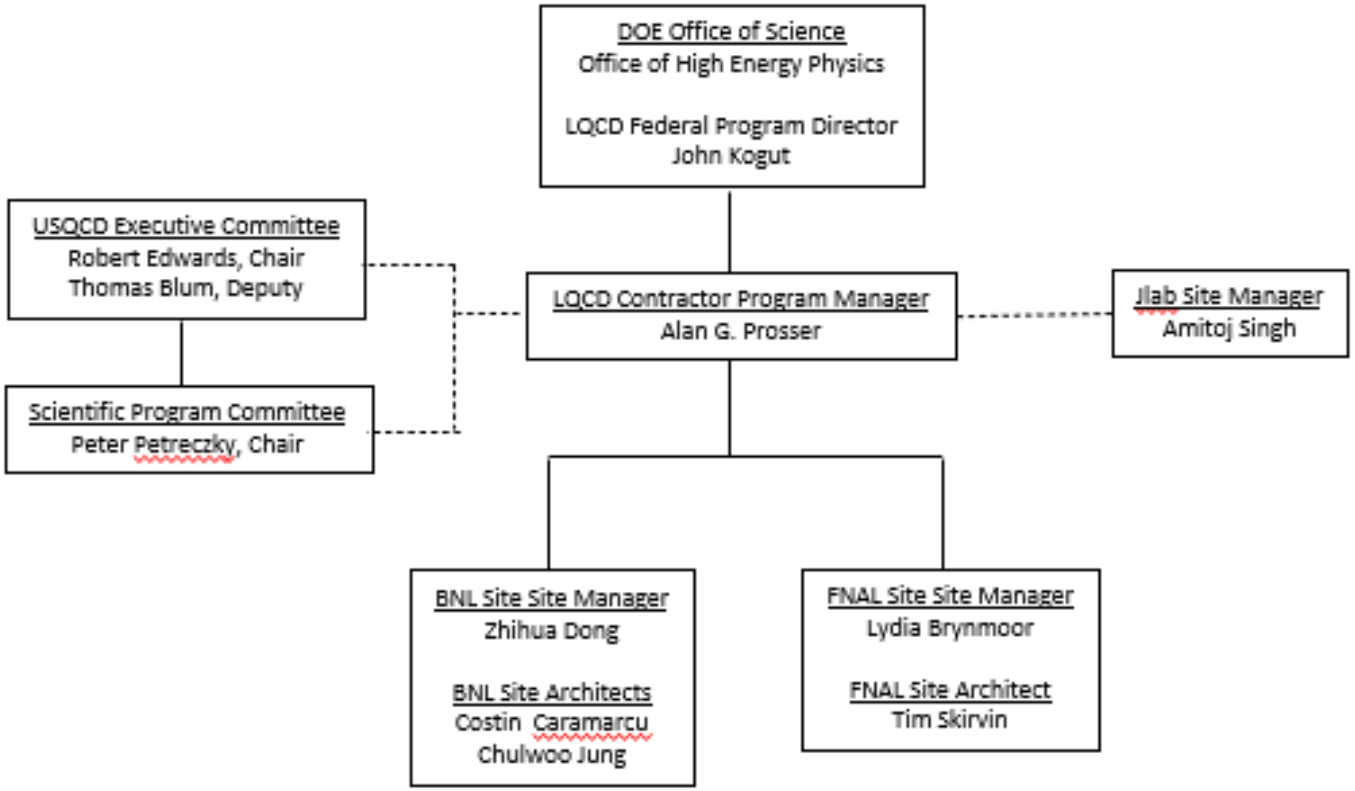

**from ext-IV Program Execution Plan (PEP)*

LQCD Management Organization: Federal Program Director

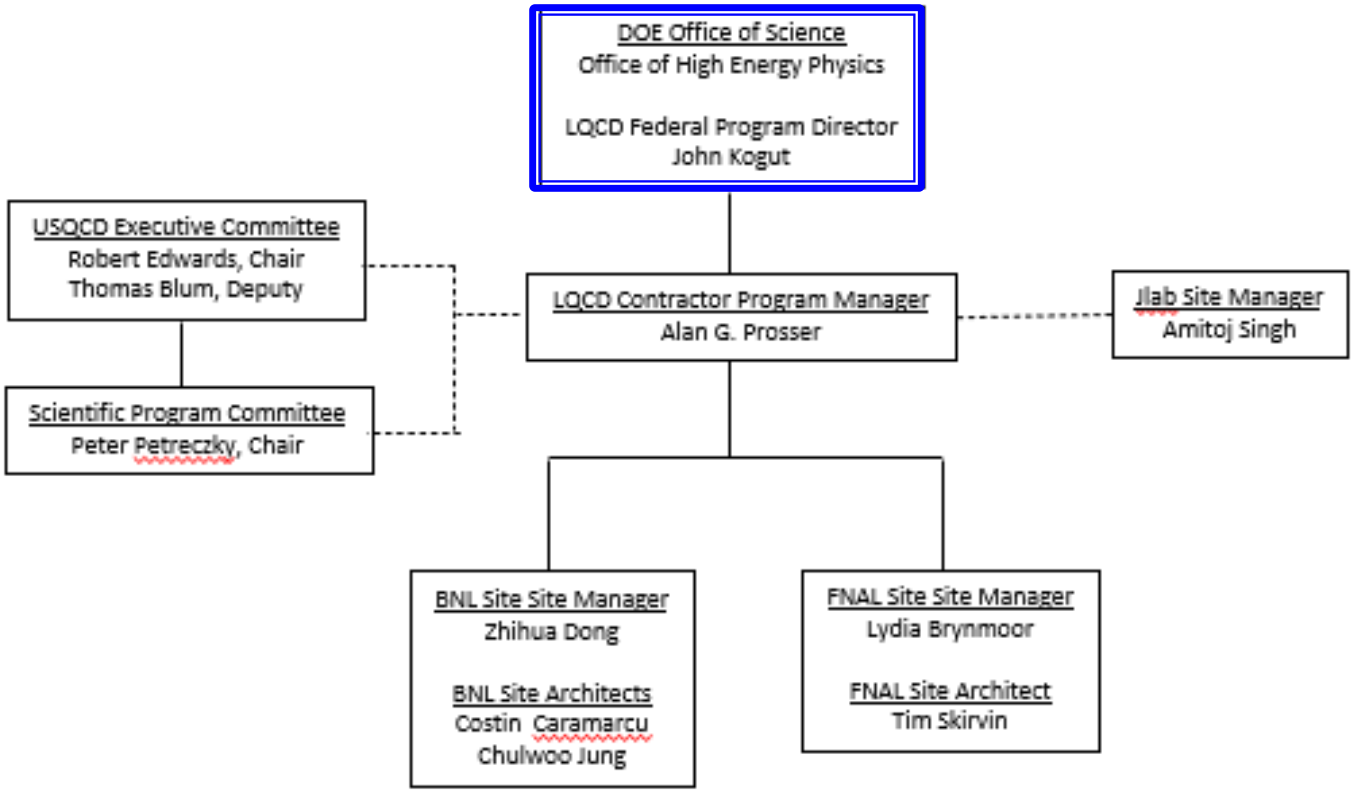

Responsibilities include:

- Program management direction
- Primary contact to DOE SC for LQCD matters
- Oversight of progress and review activities
- Budget and fund distribution management

LQCD Management Organization: Contractor Program Manager

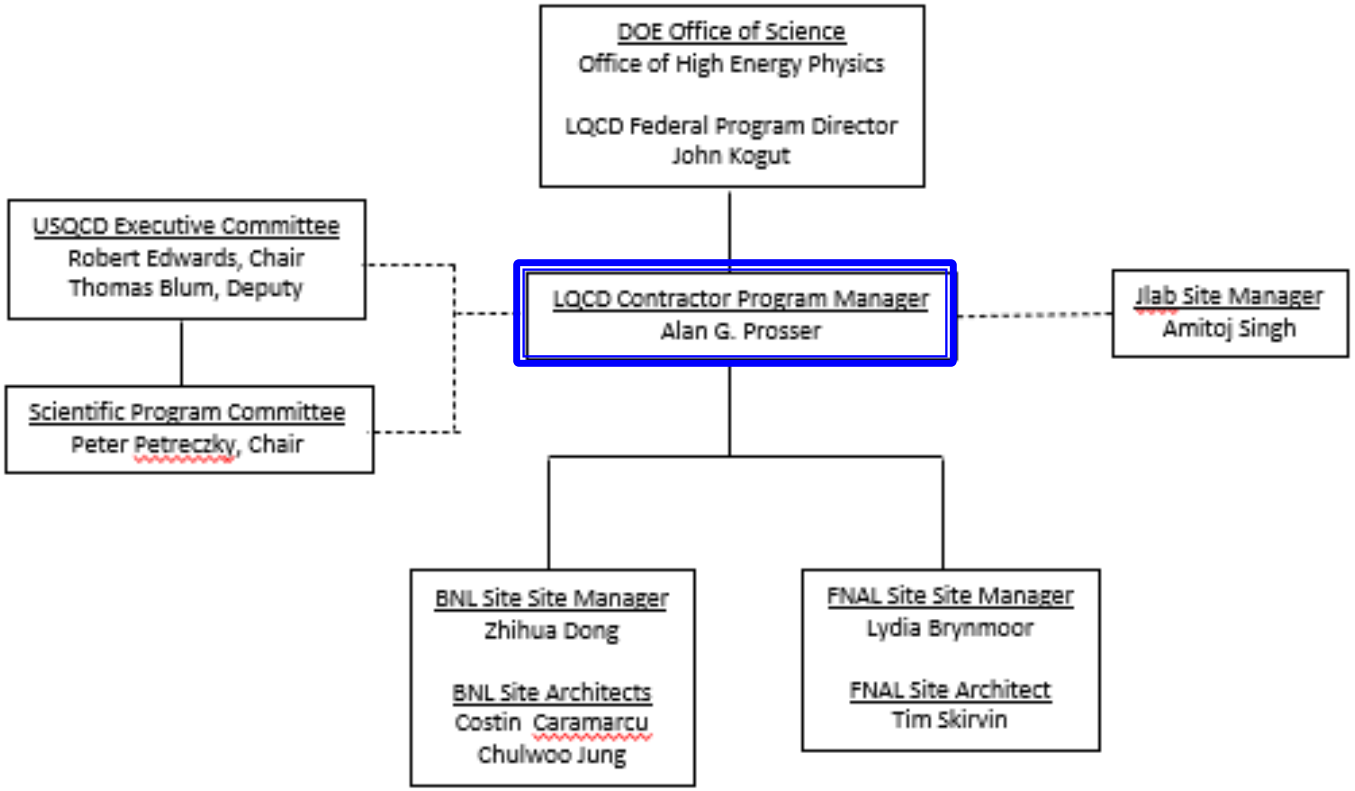

Responsibilities include:

- Oversight of planning and steady-state activities
- Documentation development and upkeep
- Establishment of MOUs with DOE laboratories
- Regularly scheduled monitoring and reporting of progress and issues

LQCD Management Organization: USQCD Executive Committee Scientific Program Committee

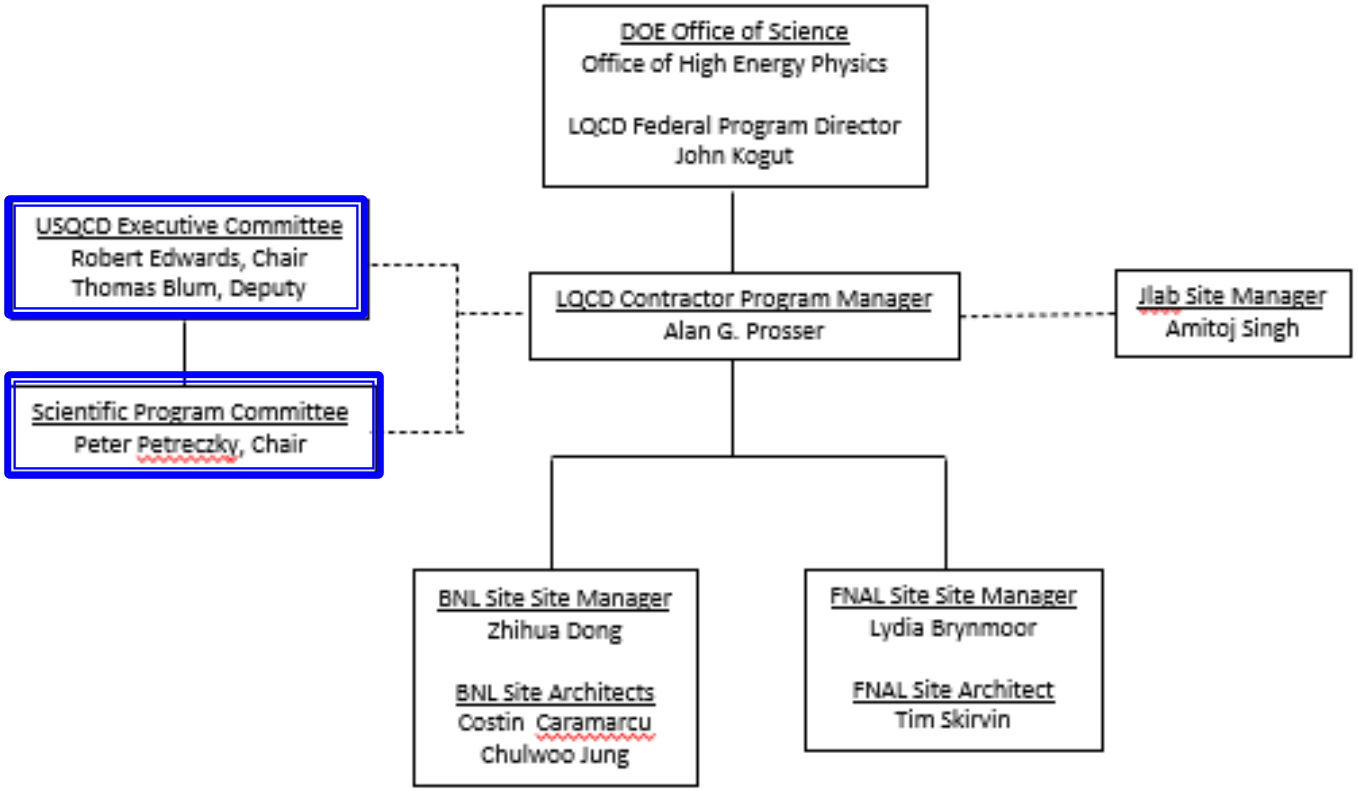

Responsibilities include:

- Establishment of scientific goals of the program
- Lead the definition of required computational and infrastructure
- Oversight of the implementation of resources into facilities
- Allocation of computational resources

LQCD Management Organization: Site Architects and Site Managers

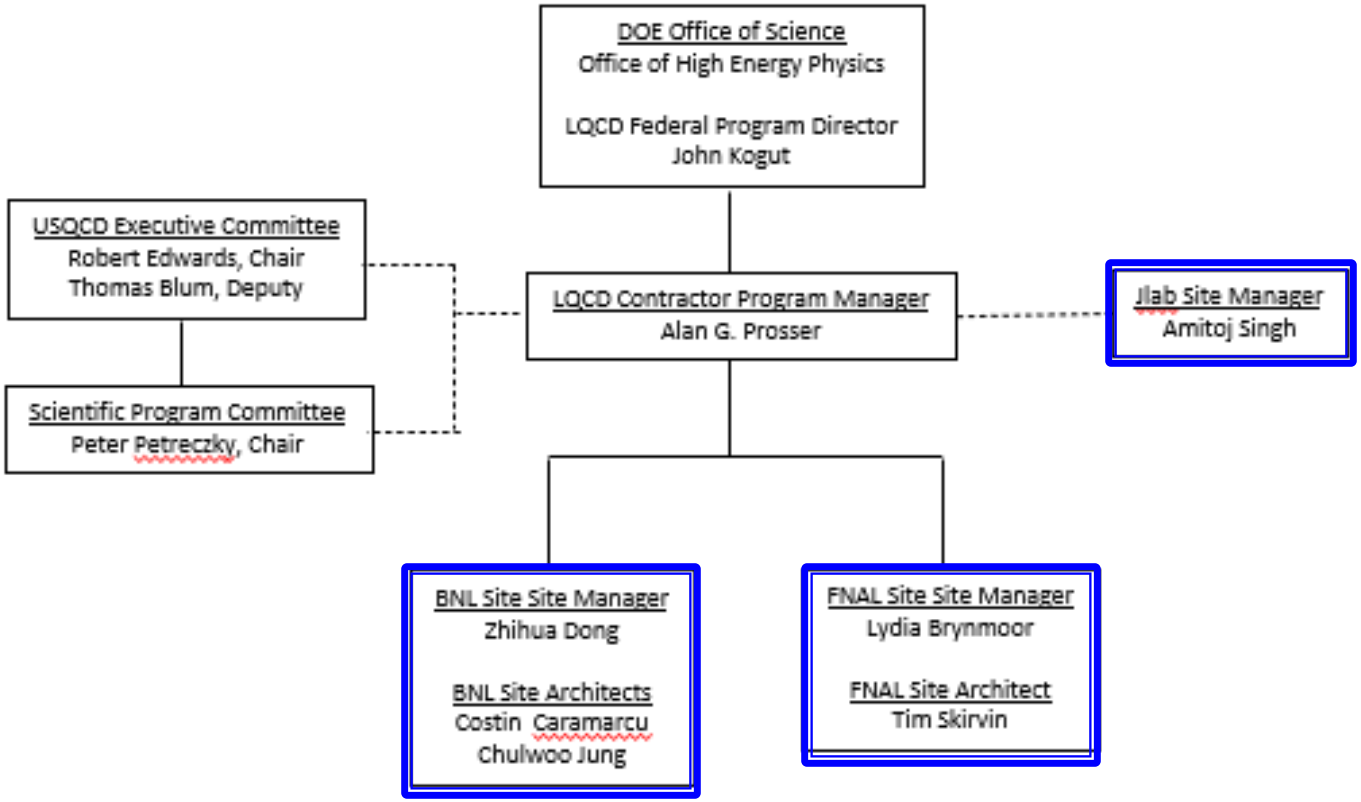

Responsibilities include:

- Day to day monitoring of computing resources
- Establishment of performance goals
- Technical oversight of computing resources
- Assist in budgeting and allocation
- Provide user support to LQCD community

LQCD Management Communications*

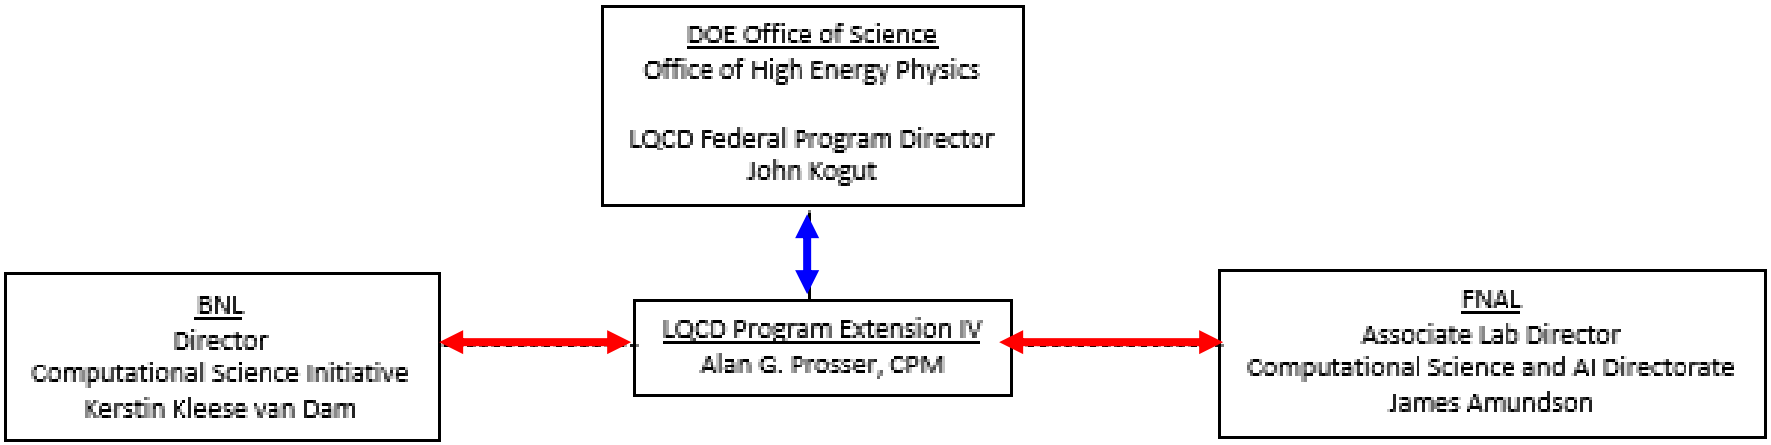

Additional Communication Activities include:

Early acquisition planning
Late acquisition planning
Early allocations process
Late allocations process

**from ext-IV Program Execution Plan (PEP)*

Program Management Documentation

Controlled Documents:

- Program Execution Plan
- Risk Management Plan & Risk Register
- Quality Assurance Plan
- Acquisition Strategy
- Certification and Accreditation Document
- Cyber Security Plan

Memoranda of Understanding (MOUs)

DOE Annual Review Reports

Performance Monitoring and Reporting

Performance Tracking and Reporting

Monthly meetings are held with the management team (incl. Federal Program Director)

CPM leads a review of the monthly performance and operational status of the sites (BNL* and FNAL) including:

Tracking of delivered computational resources and financial progress

Updates on hardware selection and procurement activities

Quarterly meetings include the same performance and status information from NP site (JLAB)

Biweekly site management meetings are held (HEP and NP represented)

**As of October, 2023, BNL allocatable clusters were retired (development nodes only are currently available) according to plan. FY24 will see the acquisition of new computing hardware for BNL (more on this later).*

Recent Conventional Systems Performance: LQ1

Conventional (CPU) Systems

- FNAL-LQ1

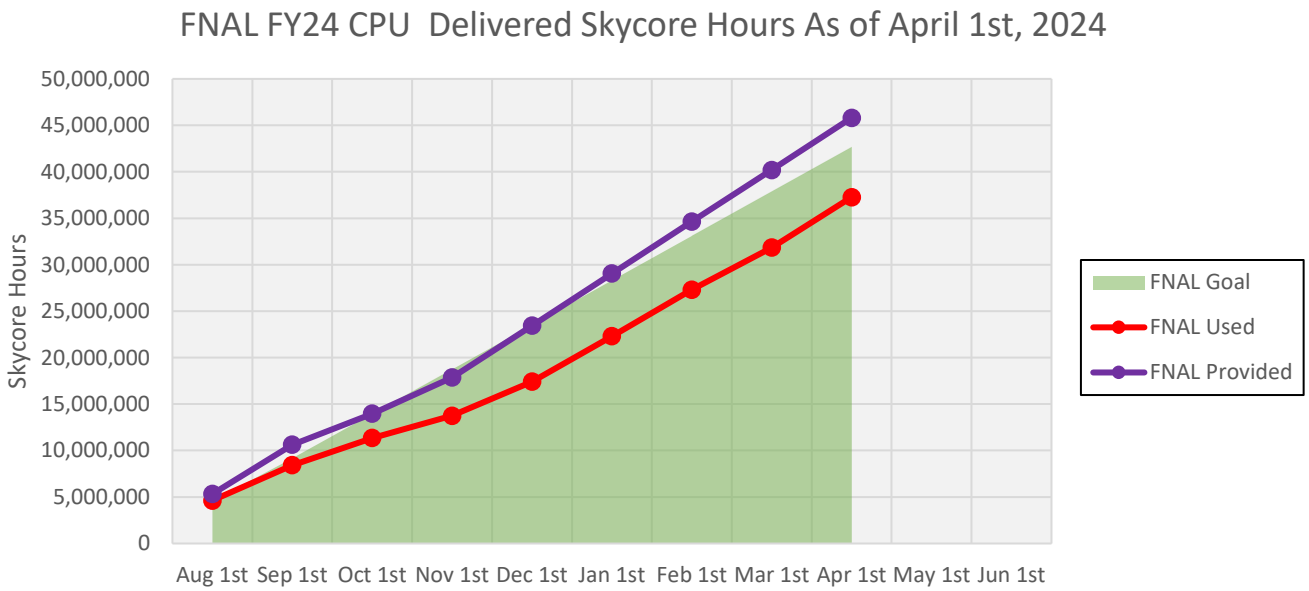

FY24	LQ1 Cumulative Performance (Skycore Hours)									
As of	Cumulative Pledged	Cumulative Provided	Cumulative Used	Cumulative % Used vs Pledged	Max possible Skycore hours	Monthly Pledged Skycore hours	Cluster Availability %	Monthly Pledged	Monthly Provided	Monthly Used
Aug 1st	4,577,580	5,319,357	4,615,256	100.82%	5,365,500	4,577,580	99.14%	4,577,580	5,319,357	4,615,256
Sep 1st	9,155,160	10,636,567	8,428,593	92.06%	5,365,500	4,577,580	99.10%	4,577,580	5,317,211	3,813,337
Oct 1st	13,942,001	13,969,371	11,352,817	81.43%	5,610,780	4,786,841	59.40%	4,786,841	3,332,803	2,924,224
Nov 1st	18,728,842	17,857,641	13,741,416	73.37%	5,610,780	4,786,841	69.30%	4,786,841	3,888,271	2,388,599
Dec 1st	23,515,682	23,445,978	17,409,658	74.03%	5,610,780	4,786,841	99.60%	4,786,841	5,588,337	3,668,242
Jan 1st	28,302,523	29,051,147	22,302,399	78.80%	5,610,780	4,786,841	99.90%	4,786,841	5,605,169	4,892,741
Feb 1st	33,089,364	34,656,316	27,317,731	82.56%	5,610,780	4,786,841	99.90%	4,786,841	5,605,169	5,015,332
Mar 1st	37,876,205	40,194,156	31,848,734	84.09%	5,610,780	4,786,841	98.70%	4,786,841	5,537,840	4,531,004
Apr 1st	42,663,046	45,804,936	37,276,263	87.37%	5,610,780	4,786,841	100.00%	4,786,841	5,610,780	5,427,529

Cluster Performance Metrics

Monthly and cumulative metrics are established to evaluate performance relative to

Pledged resource delivery:

This is a measure of how a computing site is committed by institution MOU to deliver computing resources

Provided resource delivery

This is a measure of what the computing site actually makes available for compute cycles

Used resources

This is a measure of how the provided resources are utilized by the user community

NOTE: The best possible world is one in which the sites deliver more than their allocated (pledged) resources and users take full advantage of that excess delivery

Monthly Performance and Utilization Reporting

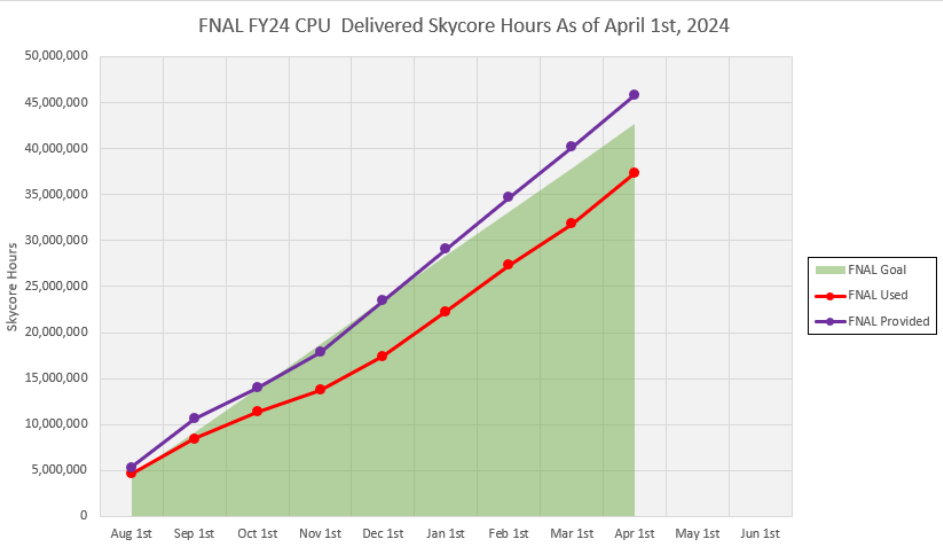

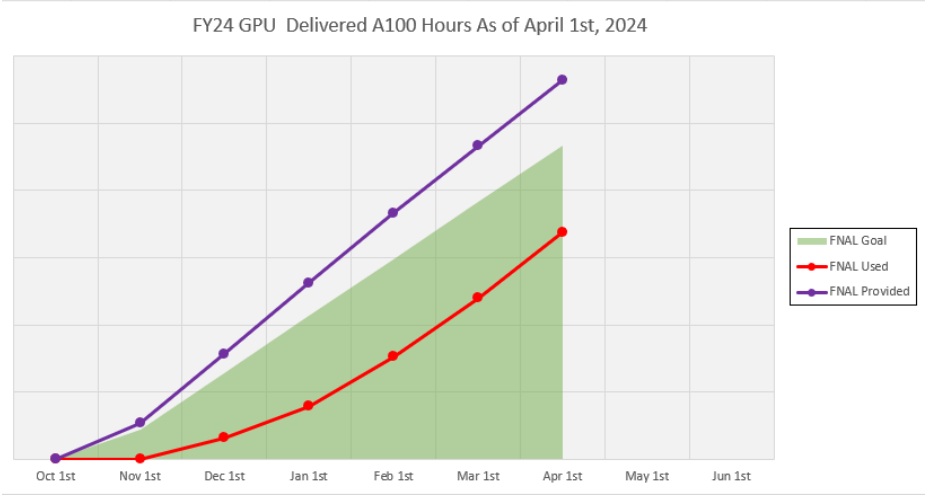

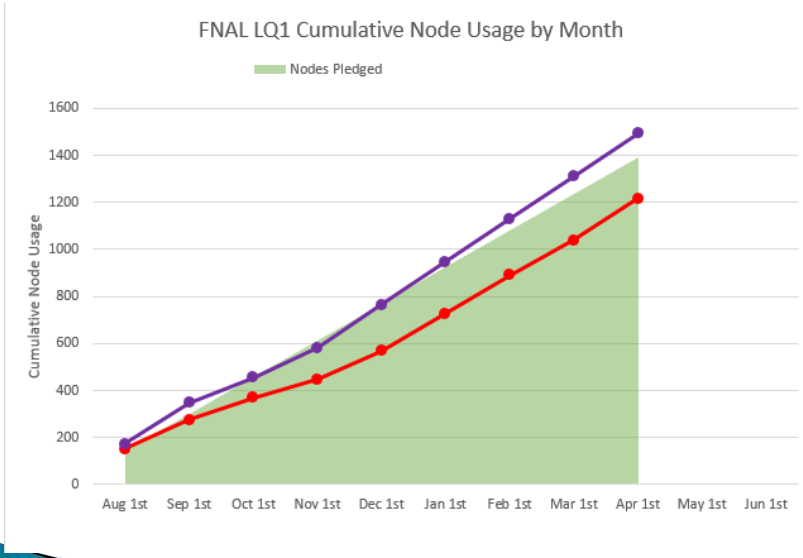

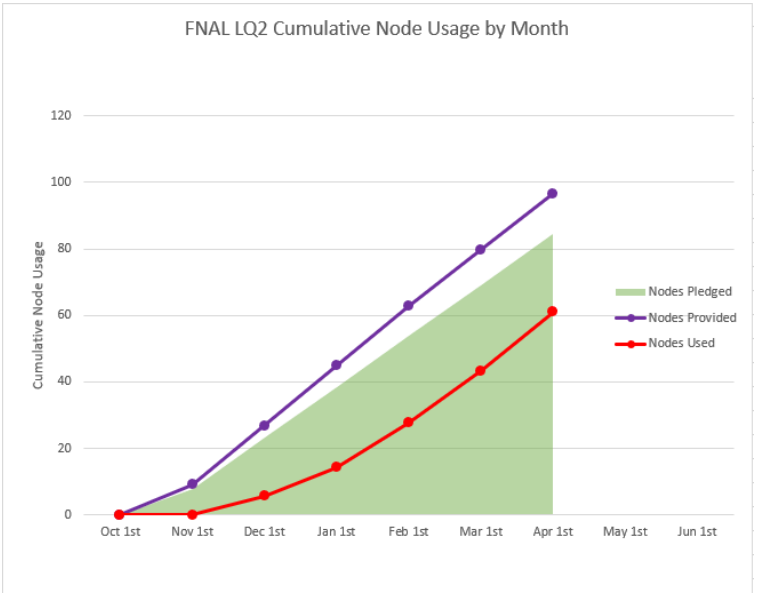

Utilization Metrics

USQCD LQ1 (FNAL) Average Node Usage for 31 days = 177 nodes

Monthly allocation = 175 nodes
Monthly Utilization (% of allocation) = 101 %
Monthly Utilization (% of pledged) = 113 %
Monthly Utilization (% of provided) = 97 %

LQ1	March				
Top User	Project	Sky-Core Hour	PI	Project Description	QOS
annah	su2betafn	1,132,011	Anna Hasenfratz	Novel phases and emerging fixed points in SU(2) gauge systems	normal
annah	fourpluseight	686,896	Anna Hasenfratz	Gradient flow renormalization scheme	normal
bazavov	mslight	610,974	Steven Gottlieb	New ensembles for precision light-meson decay constants.	normal
mlynch	stgmugm2	538,771	Michael Lynch	Muon g − 2 Hadronic Vacuum Polarization from four flavors of sea quarks using the HISQ Action	normal
bw2482	chiqcd	519,615	Bigeng Wang	Lattice calculation of nucleon energy-momentum tensor form factors using overlap fermions	normal
goodwill	lp3	513,305	Huey-wen Lin	Constraining the Bjorken-x Dependence of the Strange Distribution of the Proton Using Lattice Inputs	normal
dstewart	hadtensor	499,447	William Jay	A Lattice Calculation of the Hadron Tensor of the Pion	normal
atlytle	heavylight	343,123	Andrew Lytle	Semileptonic B decays with a vector final state	normal
trimisio	ahisq	188,951	Yannis Trimis	Novel anisotropic pure gauge simulations and the spectrum of anisotropic staggered quarks	normal
witzel	fourpluseight	127,453	Anna Hasenfratz	Gradient flow renormalization scheme	normal

Top LQCD Projects on LQ1 (March 2024)

Computing Resource Usage vs Allocation

Project Name	Cluster	SPC Original Allocation	Adjustments	SPC Adjusted Allocation	Project Used, as of April 1st, 2024	Progress against Adjusted Allocation	Remaining Allocation	30-day usage as of 04/01/2024	30-day burn rate as of 04/01/2024	Annual Pace
		(Sky-Core-Hours)	(Sky-Core-Hours)	(Sky-Core-Hours)	(Sky-Core-Hours)	(% of Alloc.)	(Sky-Core-Hours)	(Sky-Core-Hours)	(% of Alloc.)	YYYY-MM-DD
ahisq	FNAL-LQ1	3,500,000	0	2,989,157	2,034,389	68%	954,768	249,464	26.13%	2024-06-09
chiqcd	FNAL-LQ1	9,400,000	0	8,355,766	5,632,551	67%	2,723,215	519,614	19.08%	2024-08-09
fourpluseight	FNAL-LQ1	2,900,000	347,511	2,847,623	2,318,000	81%	529,623	843,212	159.21%	2024-04-28
hadtensor	FNAL-LQ1	5,600,000	0	7,669,393	4,781,074	62%	2,888,319	499,447	17.29%	2024-08-09
heavylight	FNAL-LQ1	3,800,000	0	4,999,069	3,521,213	70%	1,477,856	343,123	23.22%	2024-07-10
lp3	FNAL-LQ1	9,200,000	-80,848	3,812,758	2,190,423	57%	1,622,335	513,306	31.64%	2024-06-09
mslight	FNAL-LQ1	5,200,000	-10,799	6,075,933	3,539,794	58%	2,536,139	624,295	24.62%	2024-07-10
nptmd	FNAL-LQ1	3,500,000	-728,733	7,134,065	3,857,899	54%	3,276,166	0	0.00%	2025-07-10
stagscale	FNAL-LQ1	2,500,000	0	2,490,766	1,673,454	67%	817,312	64,871	7.94%	2025-03-11
stgmugm2	FNAL-LQ1	7,700,000	-233,206	7,044,260	4,011,983	57%	3,032,277	584,517	19.28%	2024-08-09
su2betafn	FNAL-LQ1	1,700,000	706,074	1,581,210	1,556,270	98%	24,940	1,156,762	4638.18%	2024-04-02
TOTAL	**FNAL-LQ1**	**55,000,000**	-	**55,000,000**	**35,117,050**	**63.85%**	**19,882,950**	**5,398,611**	-	-

Project Financials: FY24 Costs

Financial Performance through April 1, 2024

Cost Performance Summary		*FY completed:*		*50%*	*Period Begin/End:*	*10/1/2023- 9/30/2024*
Operating and Equipment Funds						
Lab	**FY24 Budget**	**FY24 YTD Costs**	**FY23 Open Commitments**	**FY24 YTD Open Commitments**	**FY24 YTD Balance**	**% Spent & Committed**
BNL	$1,876,000	$77,395	$0	$0	$1,798,605	4%
FNAL	$623,785	$329,633	$0	$0	$294,152	53%
Sub-total	$2,499,785	$407,028	$0	$0	$2,092,757	16%
Total	$2,499,785	$407,028	$0	$0	$2,092,757	16%

Program Extension (ext-IV) Budget Allocation Proposal

Yearly Budget Allocation by Expenditure Category ($k)

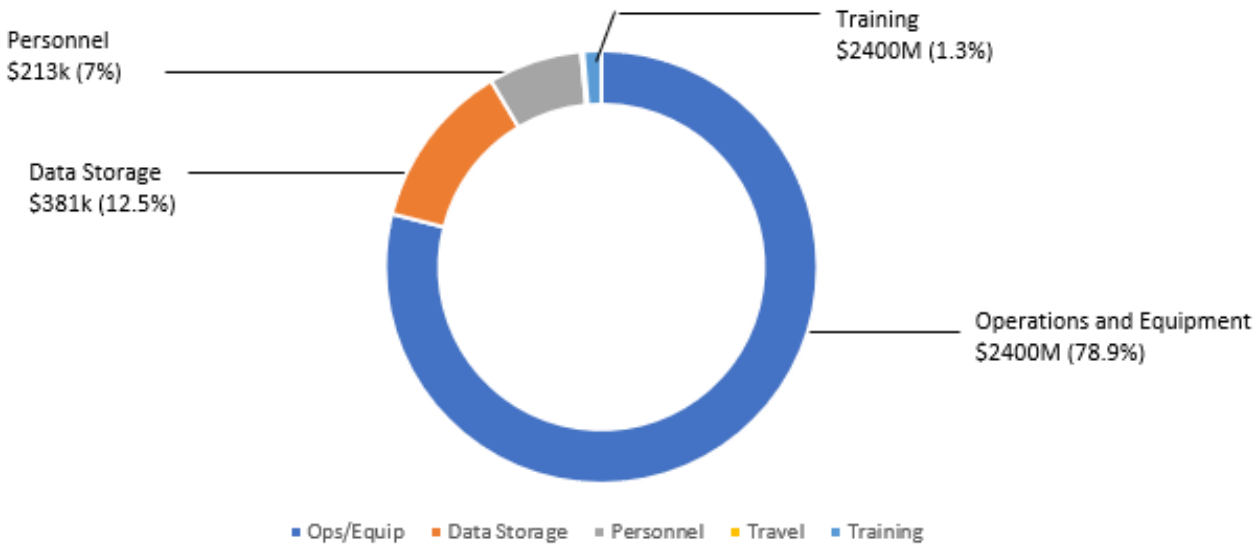

Expenditure Type	FY25	FY26	FY27	FY28	FY29	Total
Personnel	213	213	213	213	213	1065
Travel	6	6	6	6	6	30
Training	40	40	40	40	40	200
Operations and Equipment	2400	2400	2400	2400	2400	12000
Data Storage Services	381	381	381	381	381	1905
Total	3040	3040	3040	3040	3040	15200

Hardware Selection and Acquisition Activities

LQCD Cluster Portfolio Allocation Model

Hardware selection is based on an annual assessment of

- Program needs (current and projected)
- User input (technical and other criteria are considered)
- Technology roadmaps (including price points)
- Technology benchmarking

The 5-year hardware portfolio plan is updated in response to the findings of these assessments

Updates to the portfolio plan include:

- Recommendations for optimal mix of computing resources to be procured
- Establishment of procurement plans and timelines for procurement, installation, and commissioning

Presentations on the progress of the FY24 BNL acquisition as well as plans for future deployments at both BNL and FNAL are part of this meeting's agenda.

LQCD Ext-IV Capacity Planning (5-year projection)

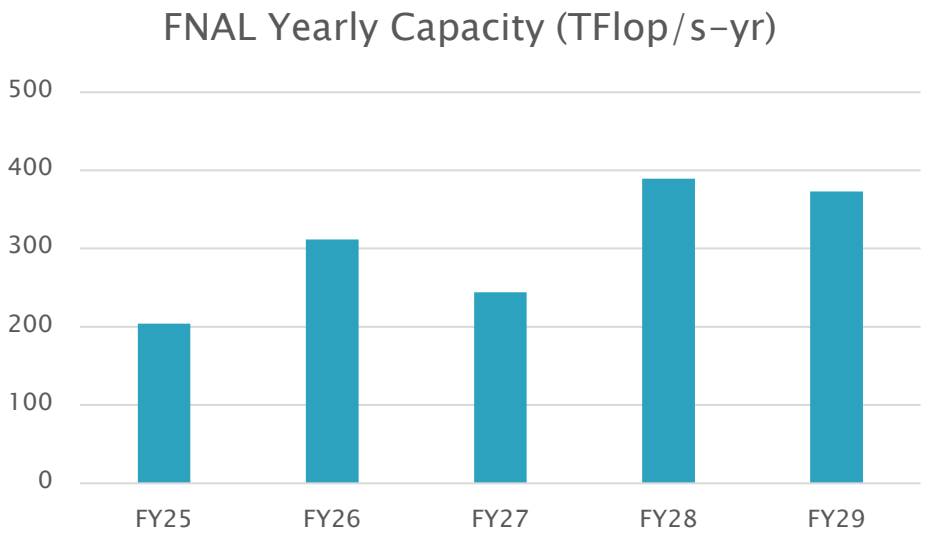

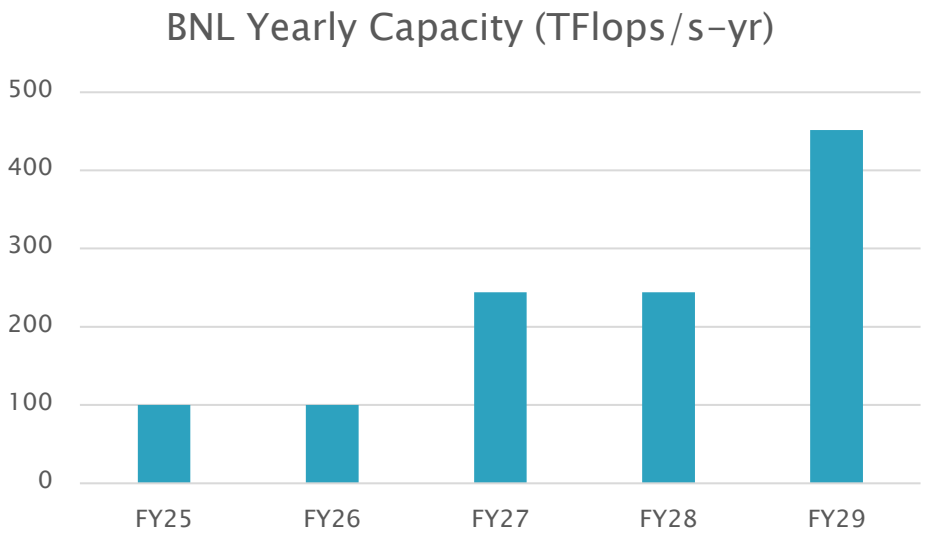

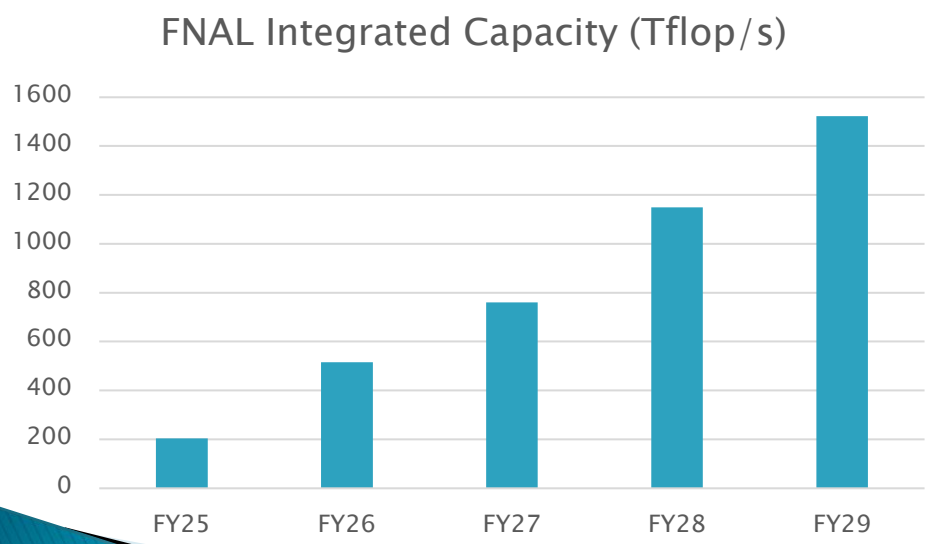

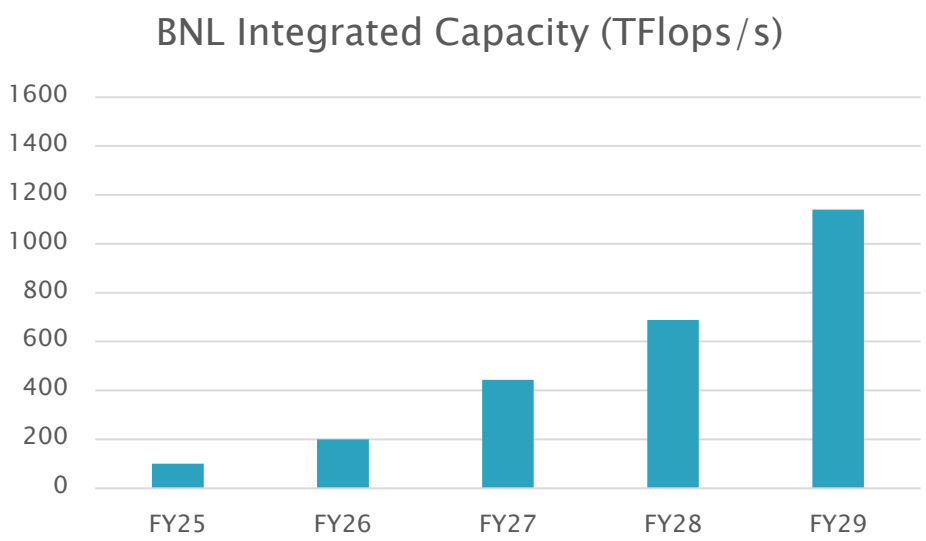

Hardware Acquisition Activities
Planning, Procurement, Testing and Delivery

Level 1 Milestones:

Computer Architecture planning for FY(n) hardware expansion complete and reviewed: *Q2 FY(n)*

Procurement of FY(n) Combined Resources: *Q4 FY(n)*

Target level of aggregate Combined Resource computing deployed and delivered: *Q2 FY(n+1)*

Level 2 Milestones
Preliminary System Design Document prepared
Request for Information (RFI) released to vendors
Request for Proposal (RFP) released to vendors
Request for Proposal (RFP) responses due
Purchase subcontract awarded
Approval of first rack
Remaining equipment delivered.
Successful completion of Acceptance Test Plan
Release to "Friendly User" production testing
Release to full production

US Lattice Quantum Chromodynamics

2023 Review of LQCD Extension III and NPPLC Initiatives (LQCD-ext III)

Date: May 29-30, 2024

In-person Meeting

Extra Slides

LQ1:

Monthly Maximum possible skycore hours

= #allocated nodes
x 40 skycores/node
x # cal-hours/month (24 x 365/12 = 730)
x skycore hour factor (1.05)

Monthly pledged skycore hours

= #allocated nodes
x 40 skycores/node
x #cal-hours/month
x skycore hour factor
x availability promise (95%)
x LQ1 Type A pledge factor (.9)

Monthly provided skycore hours

= Monthly maximum skycore hours x actual availability

LQ2:

Monthly maximum possible A100 core hours

= #allocated nodes
x 4 A100 cores/node
x #cal-hours/month

Monthly pledged A100 core hours

= #allocated nodes
x 4 A100 cores/node
x #cal-hours/month
x availability promise (95%)
x LQ2 Type A pledge factor (.85)

Monthly provided A100 core hours

= Monthly maximum A100 core hours x actual availability