Program Execution Plan
for
Lattice QCD Research Program Extension IV
(LQCD-ext IV)

Unique Program (Investment) Identifier: 019-20-01-21-02-1032-00

Operated at
Fermi National Accelerator Laboratory
Brookhaven National Laboratory

for the
U.S. Department of Energy
Office of High Energy Physics

Version 1

April 4, 2024

Version 1

Concurrences:

Alan G Prosser 5-24-2024

Alan G. Prosser — Date
LQCD Contractor Program Manager
Fermi National Accelerator Laboratory

Robert Edwards — Date
Chair, USQCD Executive Committee

5/10/2024

Kerstin Kleese van Dam — Date
Director
Computational Science Initiative
Brookhaven National Laboratory

James Amundson, UID:amundson

Digitally signed by James Amundson, UID:amundson
Date: 2024.05.09 12:49:20 -05'00'

James Amundson — Date
Associate Lab Director
Computational Science and AI Directorate
Fermi National Accelerator Laboratory

John Kogut 4/24/2024

John Kogut — Date
LQCD Federal Program Director
Office of High Energy Physics
Office of Science

Regina Rameika — 05.24.24

Regina Rameika — Date
Associate Director of Science
for High Energy Physics

LQCD-ext IV Program Execution Plan
Change Log

Version No.	Description / Pages Affected	Effective Date
1	New document for extension IV	4/3/2024

Table of Contents

1 INTRODUCTION

This document describes the program and related methodologies to be followed while executing the Lattice Quantum Chromodynamics Computing Research Program Extension IV (hereon referred to as LQCD) for the period FY2025 through FY2029. The official name is Lattice QCD Research Program Extension IV, and the Unique Program Identifier is 019-20-01-21-02-1032-00.

The LQCD research program is an extension of the LQCD-ext III Computing Project, which ends in September, 2024. LQCD will support the continued procurement of institutional computing hardware resources located at Brookhaven National Laboratory (BNL) and Fermi National Accelerator Laboratory (FNAL). BNL and FNAL provide facilities and infrastructure that deliver the mid-scale computing required by the LQCD Program. They will also provide computing professionals to plan, design, deploy, and operate the computing systems.

This plan has been prepared following the guidance in DOE Order O413.3B, *Program and Program Management for the Acquisition of Capital Assets* (dated 29-Nov-2010).

2 JUSTIFICATION OF MISSION NEED

LQCD directly supports the mission of the DOE's SC HEP Program "to explore and to discover the laws of nature as they apply to the basic constituents of matter and the forces between them". LQCD also supports the Scientific Strategic Goal within the DOE Strategic Plan to "Provide world-class scientific research capacity needed to: advance the frontiers of knowledge in physical sciences.; or provide world-class research facilities for the Nation's science enterprise."

Quantum Chromodynamics (QCD) is the fundamental theory of quarks and gluons that has far-reaching implications for understanding hadrons, nuclei, and other aspects of matter. The long-range interactions of QCD are essential for describing these phenomena, and the phenomenon of color confinement ensures that strongly interacting particles are detected as hadrons. This makes hadrons useful for experiments, where they are often used as beams or targets. Because of the importance of QCD, it is necessary to have a deep understanding of the theory at nonperturbative scales. Numerical lattice gauge theory is the only comprehensive, quantitative method for investigating nonperturbative QCD from first principles.

The purpose of these simulations is to attain a deep understanding of the physical phenomena that are encompassed by QCD. Additionally, these simulations aim to make precise calculations of the theory's predictions and test the validity of the Standard Model. Lattice simulations are indispensable for solving fundamental problems in high-energy and nuclear physics, which are central to the DOE's extensive experimental endeavors in these fields.

Major goals of the experimental programs in high energy and nuclear physics on which Lattice QCD simulations will have an important impact are to: 1) quark flavor physics and determination of Standard Model parameters, 2) lepton flavor physics including the muon anomalous magnetic moment, 3) neutrino physics, 4) strongly coupled quantum field theories beyond the Standard

Model, and (5) dark matter and the Cosmic Frontier. Lattice QCD calculations are essential to the research in all these areas.

3 PROGRAM DESCRIPTION

The purpose of the LQCD computing program is to provide the USQCD user community with the mid-scale computing resources required to meet the computational needs of the lattice quantum chromodynamics (QCD) research program for FY2025-2029. LQCD will purchase computing cycles on mid-scale computing systems and data storage capacity on institutional disk and tape drives located at BNL and FNAL.

Fermilab will provide the LQCD Program Manager, who will provide management and oversight of all program activities at the two host laboratories; the Program Office will be located at Fermilab. A detailed description of the roles and responsibilities of these and other key positions can be found later in this document. The budget associated with this work provides salary and travel support for Program management and computing professional staff. It also provides for the purchase of computing cycles and data storage capacity. Computing cycles are purchased in node-hours/month; storage capacity is purchased in terabytes/month.

3.1 Functional Requirements

Three phases of computing are done on lattice QCD machines. In the first phase, a Markov Chain Monte Carlo (MCMC) simulation of the QCD vacuum is carried out, and a sequence of gluon field configurations, which are representative samples of the vacuum, are generated and archived. An ensemble of several thousands of such configurations is typically needed to adequately characterize statistical errors. Collections of gauge ensembles for a fixed discretization of the QCD action, but with varying lattice spacing, sea-quark masses, and lattice volume are necessary to quantify systematic effects and take the continuum limit. The collections of gauge ensembles are useful for a wide range of physics projects. As such, they are of high scientific value and their generation costs are rapidly amortized by their high degree of reuse. Ensemble generation is serially dependent and represents a strong scaling computational problem. Ideally one would be able to use efficiently $O(10^4)$ computing nodes on $O(256^4)$ lattice sites, requiring leading capability-class computing resources beyond the scope of this project. Ensembles on $O(32^4)$ lattices, or smaller, however, are tractable using dedicated cluster resources. Computational performance on the largest scales is limited by halo-exchange communications since the local data bandwidths vastly exceed those of inter-node communication.

The second, hadron observables, phase of lattice QCD computes many thousands of valence quark propagators on the gauge configurations and assembles them into hadronic correlation functions. The expectation value of a given correlation function averaged over each gauge field in an ensemble yields the hadronic observable on the ensemble. Computational costs are dominated by the quark propagators which are the solutions of the inverse of the Dirac operator represented in lattice QCD by a sparse difference operator over the four-dimensional lattice domain. Iterative Krylov solver algorithms such as standard conjugate gradient or multigrid preconditioned conjugate gradient are used in propagator solves. The propagator solves and correlation function "tie-ups" for each gauge configuration are independent from the same computation on the other configurations, hence these operations are typically done in independent jobs requiring fewer nodes that needed for configuration generation. The large number of quark propagators permits opportunities to amortize the setup costs of advanced solver algorithms such as multigrid or eigenmode deflation. Moreover, the local-memory bandwidth requirements of calculating quark propagators can be amortized and the computational rate can be brought closer to the cache-bandwidth limit through use of multiple right hand side solvers and some trivial-level parallelism among computing nodes. In summary, computing hadronic variables is considerably more of a compacity workload in contrast to the capability requirements of configuration generation.

The third, analysis phase of a lattice QCD workflow gathers the hundreds of thousands of hadron correlation functions produced by many hadron observables campaigns for analyses leading to the physics results such as hadronic masses, resonance widths, and decay matrix elements in the continuum limit. A typical physics analysis involves doing upwards of a million chi-square minimization fits and statistical techniques such as bootstrap to refine the analysis and produce estimates with uncertainties. The management of many millions of data objects and analysis results is a concern in the analysis phase. The analysis phase requires interactive access to a multi-core workstation and the ability to take advantage of batch resources to offload some analysis tasks.

In each year of the program, the combination of institutional cluster resources that best accomplishes the scientific goals for LQCD calculations will be purchased. Determining the appropriate mix of conventional and GPU-accelerated resources occurs as part of the annual planning process and is based upon several factors, including user demand determined through the annual SPC resource allocation process, cost effectiveness, and software availability.

3.2 Computational Requirements

Lattice QCD is performed on structured Cartesian grids with a high degree of regularity and natural data parallelism. The fundamental kernels of both configuration generation and quark propagator solves are operations in SU(3) algebra. This algebra is represented by complex matrices (3x3) and vectors (3x1). SU(3) matrix-vector multiplication dominates the calculations. The four-dimensional space-time lattices used in lattice QCD calculations are quite large, and algorithms typically allow very little data reuse. With lattices spread over even thousands of processor cores, the local lattice volume still often exceeds the processor's cache sizes. Hence, the performance of these fundamental kernels is often not limited by the floating-point capability, but rather by either bandwidth to main memory, or by limitations of the network fabrics interconnecting the processors. Historically, the rate at which the network interconnects have improved lags improvements in memory performance. A balanced system design takes into consideration such system limitations. A fruitful way to characterize the performance of algorithms is to consider their arithmetic intensity, or the ratio of bytes accessed to the number of floating-point operations. Table X lists the arithmetic intensities of the Dslash kernels for several discretizations of the quark action.

Action	**Flops**	**Bytes**	**Bytes / Flops**
Staggered	1146	1560	1.36
Wilson	1320	1440	1.09
DWF	L_s x 1320	L_s x 864	0.65
Wilson (mult rhs)	N_{rhs} x 1320	N_{rhs} x 864	0.65

Staggered (mult. rhs)	N_{rhs} x 1146	N_{rhs} x 408	0.36

Table X Arithmetic intensities of the Dslash operator. Kernels with higher intensities are more susceptible to bandwidth limitations. The quantity L_s is the length of the fifth dimension for domain wall fermions. Algorithms allowing multiple right-hand sides (mult. rhs) in solves have higher reuse of data.

From a roofline performance analysis and the intensities in table X, staggered Dslash is expected to have a lower maximum flops-per-second performance rate than DWF for a fixed limited memory bandwidth. We find that algorithms taking advantage of multiple right-hand sides are able to mitigate the bandwidth limitation through higher data reuse. Note that optimized propagator solvers frequently take advantage of mixed precision by doing initial iterations with lower float precision. Lower precision floats require fewer bytes hence reducing the arithmetic intensity and raising the achievable flops-per-second rate. It is desirable to hide network latencies by overlapping off-node communications and computations within kernels. This sets a minimum aggregate bandwidth requirement on the network. Reference [?] considers the actual DWF Dslash performance on quad GPU worker nodes with lattice local volume 32^4 per node. The measured performance is 10 Tflop/s per node; hence, the effective cache bandwidth is B_{cache} = 10000 GFlops/s x 0.65 Bytes/Flop = 6500 GB/s. As shown in the reference, the network aggregate bandwidth needed to overlap communications with computation in the DWF kernel is about $B_{network} \sim 1/L \times B_{memory}$ = 200 GB/s where L=32. Similar predictions can be done for the other actions, although, the prefactor (1/L) is different for other actions.

3.3 I/O and Data Storage Requirements

During vacuum configuration generation, data files specifying each representative configuration must be written to storage. These files are at least 10 GBytes in size, with a new file produced every two hours. Thus, the average I/O rate required for configuration storage is modest at only 1.4 Mbytes/sec. However, higher peak rates of at least 100 Mbytes/sec are desired, to minimize the delays in computation while configurations are written to or read from external storage.

During the next stage, hundreds of configurations must be loaded into the machines to calculate the propagation of quarks on each configuration. This requires the numerical determination of multiple columns of a large sparse matrix. The resulting "propagators" are combined to obtain the target measurements. Propagator files for Clover quarks, for example, are 16 times larger than the corresponding gauge configuration. Often, dozens or hundreds of propagators are calculated for each gauge configuration. To minimize the time for writing to and subsequently reading from scratch storage space, the sustained I/O rate for each independent analysis job may be as high as 300 Mbytes/sec for a fraction of the duration of the job. The mix of jobs on a given cluster may be manipulated using the batch system to preclude saturation of the I/O system.

Sometimes, instead of computing propagators directly, many eigenvectors of the Dirac matrix (or similar) are stored, especially for problems that require numerous propagators. The eigenvectors are reused many times. On current lattices and ensembles, the required long-term storage can be several PB. These storage needs must be subjected to cost-benefit analyses to ensure that storage is not more expensive than re-computation.

The final stage, in which statistical and systematic uncertainties are estimated, imposes a more modest burden on storage.

3.4 Data Access Requirements

The bulk of configuration generation is performed at the DOE Leadership Computing Facilities and other capability facilities. Archival storage of ensembles of these configurations utilizes robotic tape facilities at BNL and FNAL, in addition to tape storage available at facilities producing configurations. LQCD maintains services to provide facile movement of data sets among sites involved with the generation and analysis of gauge ensembles. The aggregate size of data moved between sites is at least 200 terabytes per year and although the 200 terabytes number is difficult to verify, it is considered a lower bound.

3.5 Cluster Portfolio Allocation Model

The LQCD Program Office maintains a 5-year hardware portfolio plan that defines current and planned mid-scale computing assets available to the research program from the host laboratories. Based on experience, the production lifecycle for suitable hardware clusters is typically 5 years.

On an annual basis, the Program Office and USQCD leadership meets with computing leadership from the host laboratories to review past performance, LQCD computational needs, vendor, and Leadership Class Facility (LCF) roadmaps, institutional hardware roadmaps, and future expansion plans. The purpose of these meetings is for all parties to update each other on current and future and needs, and to ensure alignment of objectives. Outcomes from the meetings include updating the 5-year hardware portfolio plan, determining the optimal mix of computing resources that will be procured from the host laboratories in the next fiscal year, and formulating plans and timelines for expansions of existing computing and storage resources.

When a host laboratory announces its intention to procure a new compute cluster based on an architecture not currently in the existing cluster portfolio, a joint evaluation committee is formed to evaluate options. These joint committees consist of subject matter experts (SMEs) representing the needs of LQCD and other laboratory user groups who may use the system. The committee charge includes gathering user requirements, identifying potential solutions, and formulating software benchmarks to use in measuring performance. The committee evaluates potential solutions against anticipated usage and performance objectives. Potential vendors are asked to provide test hardware so LQCD codes can be run on the new system. Performance is benchmarked and compared against scientific requirements and planned milestones. An alternatives analysis is performed to determine the most cost-effective solution for a given year. The committee makes a recommendation to laboratory computing leadership and the LQCD Contractor Program Manager (CPM) regarding the preferred solution. The recommendation will be included in a written report that summarizes and present the results of the committee's work.

The CPM reviews the recommendation and forwards a copy of the written report to the USQCD Executive Committee (EC) for consideration. If appropriate, the CPM approves the recommendation with the concurrence of the EC. Host laboratory computing leadership is informed that the recommended option is acceptable to LQCD. If also acceptable to the host lab, the procurement process is initiated by the host lab. If the recommended option is not accepted by both LQCD and the host laboratory, further discussions occur between LQCD, the host laboratory, and the evaluation committee with the intent of identifying a solution acceptable to all parties.

Once an acceptable solution is identified, the host laboratory initiates the procurement process, which typically begins with the preparation and execution of a Request for Proposals (RFP). The host laboratory is responsible for procuring, installing, commissioning, and deploying new systems. A deployment schedule is created and communicated to the LQCD Program Office, where it is tracked to completion.

3.6 Operations

LQCD operations associated with the use of the institutional clusters include user support, system administration, system performance monitoring (e.g., capacity utilization and system availability), configuration management, cyber security, data storage, and data movement.

Archival storage of physics data utilizes tape robots and hierarchal mass storage systems at BNL and FNAL. Tape media used to store program data is procured using program funds.

On a periodic basis, USQCD collaboration members apply to and receive from the Scientific Program Committee allocations of computing time at one or more of the two sites. Specific physics programs may utilize both sites to take advantage of the specific characteristics of each. For this reason, efficient movement of physics data between sites is essential.

3.7 Major Interfaces

As previously noted, BNL and FNAL are the primary participating laboratories. Memoranda of Understanding (MOU) are established between LQCD and each host laboratory that define the relationships and expectations between both parties.

3.8 Key Stakeholders

Key stakeholders include the DOE Office of Science, the DOE Office of High Energy Physics, and the laboratories hosting LQCD computing facilities. Members of the USQCD collaboration are key customers of the LQCD computing facilities. These include laboratory and university researchers, as well as post-docs and students. Their feedback will be provided throughout the Program through the USQCD Executive Committee and spokesperson.

4 MANAGEMENT STRUCTURE AND INTEGRATED PROGRAM TEAM

This section describes the management organization for LQCD and defines roles and responsibilities for key positions. The management structure is designed to facilitate effective communication between the management team and key stakeholders. The LQCD organization chart for management and oversight is shown in Figure 1. Solid lines indicate reporting relationships. Dashed lines represent advisory relationships with the USQCD Executive and Scientific Program Committees, and a collaborative relationship with Jefferson Lab (JLab).

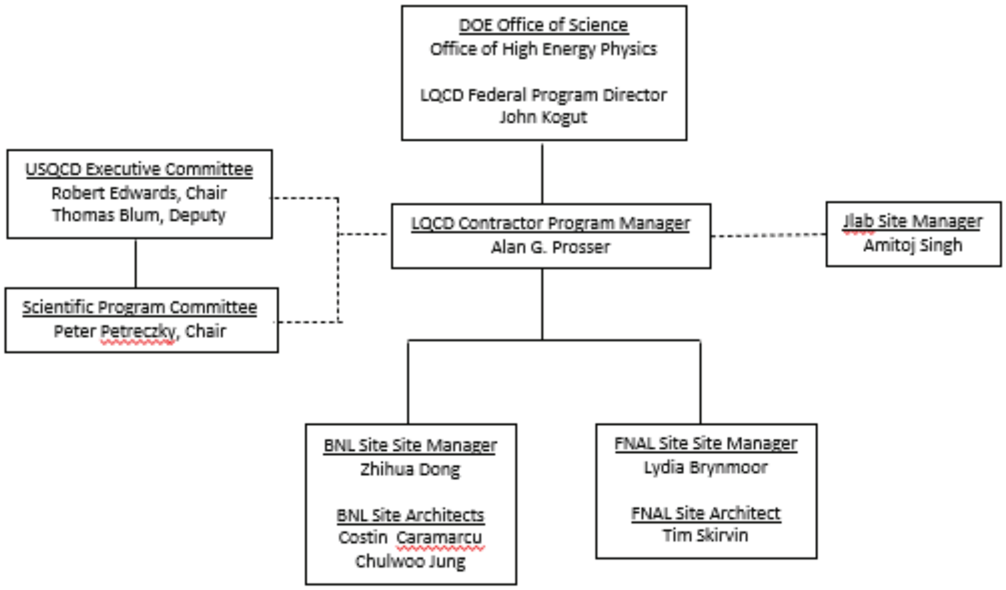

Figure 1: LQCD Management Organization Chart

4.1 Roles and Responsibilities

4.1.1 LQCD Federal Program Director

Overall management and oversight are provided by the DOE Office of Science, through OHEP. The LQCD Federal Program Director is appointed by OHEP. The LQCD Federal Program Director is John Kogut; he is a certified DOE Level 1 Qualified IT Project Manager

Specific responsibilities of the Federal Program Director include the following:

- Provide Program management direction for the LQCD Program.
- Serve as the primary point of contact to DOE SC headquarters for LQCD matters
- Oversee LQCD progress and help organize reviews as necessary
- Budget and manage the distribution of funds for LQCD

4.1.2 Contractor Program Manager

The LQCD Contractor Program Manager (CPM) is responsible for the overall management of the Program. This person is the key interface to the Federal Program Director for financial matters, reporting, and reviews. The CPM has significant budgetary control and is in the approval chain for all major Program commitments and procurements. The Contractor Program Manager is Alan Prosser from Fermilab.

Specific responsibilities for the Contractor Program Manager include the following:

- Provide management and oversight for all planning and steady-state activities associated with program execution.
- Ensure that critical program documents exist and are kept up to date, such as the Program Execution Plan, Risk Management Plan, Acquisition Plan, Alternatives Analysis, and Certification & Accreditation Documentation.

- Develop and maintain a work breakdown structure (WBS) with tasks defined at a level appropriate to successfully manage the program, and that can be externally reviewed. The WBS should include program milestones at a level appropriate to track progress.
- Establish and maintain MOUs with the DOE laboratories providing LQCD computing resources.
- Gather and summarize financial information for the monthly progress reports to the LQCD Federal Program Director.
- Present monthly progress reports to the LQCD Federal Program Director. These reports cover cost and schedule performance, performance against established key performance metrics, review of annual acquisition strategies and progress against deployment plans, and other significant issues related to execution as appropriate.
- Prepare and submit to DOE annual operating budgets and financial plans consistent with the program plan and performance objectives and manage costs against the approved budget.
- Provide final approval of all major (> $50K) procurements
- Provide internal oversight and reviews, ensuring that funds are being expended according to the program plan and identifying weaknesses in the execution of the plan that need to be addressed.
- Prepares and manages other technical and controlled documents as requested
- Monitors and reports on activities related to performance assessment
- Develops and maintains program-management-related communications including the Program web site and the repository of program documents, etc.
- Leads the annual user survey process, which includes preparing the survey, analyzing, and reporting on survey results, and preparing annual user survey reports

Interactions of the Contractor Program Manager:

- Reports to the LQCD Federal Program Director.
- Serves as the primary point of contact with DOE SC, through the LQCD Federal Program Director, on matters related to budget and schedule for all funded activities.
- Interacts with host laboratory senior management regarding program-related matters.
- Provides direction and oversight to LQCD Site Managers on program-related matters.
- Interacts with the Chair of the USQCD Executive Committee and the Chair of the Scientific Program Committee to ensure collaboration needs are being met.

4.1.3 Site Managers

Steady-state operations at each host laboratory are monitored by a designated Site Manager (SM) who is located at that site. The SM is responsible for developing and executing the corresponding components of the WBS and making sure that appropriate commitments by the host laboratory are obtained and carried out. The SM is the primary interface between the CPM, the host laboratory, and the individuals associated with the work to be performed at that host laboratory.

Specific site manager responsibilities include the following:

- Provide day-to-day monitoring of the LQCD computing resources being used at his/her site.

- Provide user support to the USQCD community
- Ensure that funds are being expended according to the program plan and identifying weaknesses in the execution of the plan that need to be addressed.
- Obtain necessary resources and approvals from laboratory management and coordinate resources contributed by the laboratory
- Provide technical oversight of the LQCD computing resources at the host site including the monitoring and reporting of system performance metrics such as uptime and usage.
- Implement and monitor user allocations as determined by the Scientific Program Committee.
- Participate in the hardware selection process for deployments at his/her site, representing LQCD computing needs.
- Assist in the annual budget planning and allocation process, and in the preparation of detailed planning documents, including the WBS and performance milestones at a level appropriate for external review.
- Track progress of site-specific performance milestones.
- Prepare and submit monthly status reports, including expenditures and effort, to the CPM
- Prepare materials for external oversight and reviews and participate in external review activities, as necessary.

Interactions of the Site Manager:

- Reports to the CPM
- Works closely with the CPM and other Site Managers both to assist in defining milestones and infrastructure deployment schedules, and to ensure a high level of coherency.

4.1.4 Site Architects

The Site Architect (SA) is responsible for representing the technical design and architecture needs of the LQCD Computing Program at their host site. The Site Architect assists the Site Manager on strategic issues, monitoring, and reviews, but does not have day-to-day operations responsibilities.

Specific site architect responsibilities include the following:

- Participates in hardware selection activities at their host laboratory, working with the SM and host laboratory management.
- Establishes performance goals and benchmarks for LQCD systems located at or to be located at their site.
- Assist the SM in the monitoring and assessment of actual performance versus planned performance.
- Assists the CPM in documenting and communicating:
  - Hardware selection information for acquisition planning (target audience is USQCD Executive Committee)
  - Performance goals and benchmarking information for allocation process (target audience is Scientific Program Committee)

Interactions of the Site Architect:

- Reports to the CPM

- Works closely with the CPM, Site Managers, and other Site Architects both to assist in defining milestones, and to ensure a high level of coherency across the program.
- Works closely with technical staff at the host laboratory in communicating LQCD computing needs and participating in the design and selection of new institutional clusters.

4.1.5 Integrated Program Team

The LQCD Integrated Program Team (IPT) is composed of the LQCD Federal Program Director, CPM, Site Managers, Site Architects, USQCD Executive Committee Chairperson and Deputy, and USQCD Scientific Program Committee Chairperson and Deputy. The LQCD Federal Program Director chairs the IPT. The current membership of the IPT is given in Appendix A.

The full IPT meets on an as-needed basis, however subsets of the IPT meet on a regular basis. For example, monthly meetings are held between the Federal Program Director and CPM to review progress against goals and milestones. The CPM, and Site Managers meet bi-weekly to discuss operations and review performance on a more detailed, technical level. These meetings often involve planning for subsequent deployments and sharing lessons learned. Site Architects participate in these meetings when they involve acquisition planning, architectural design, or other Site Architect responsibilities.

4.1.6 USQCD Executive Committee

The charter of the USQCD Executive Committee is to provide leadership in developing the computational infrastructure needed by the United States lattice gauge theory community to study Quantum Chromodynamics (QCD), the theory of the strong interactions of subatomic physics. The Executive Committee is responsible for setting scientific goals, determining the computational infrastructure needed to achieve these goals, developing plans for creating the infrastructure, securing funds to carry out these plans, and overseeing the implementation of all the above. The Executive Committee advises the CPM regarding scientific priorities and the computing resources needed to accomplish them. The Executive Committee appoints the Scientific Program Committee, which allocates the Program's computational resources.

Members of the Executive Committee rotate at the rate of around one per year. Around half of the members of the Executive Committee are expected to remain during the lifetime of the Program. If a vacancy occurs, it is filled by a vote of the remaining members of the Executive Committee. Appendix B contains a list of the current members of the Executive Committee.

Responsibilities

- Sets the scientific goals and determines the computational infrastructure needed to achieve them
- Establishes procedures for the equitable use of the infrastructure by the national lattice gauge theory community
- Arranges for oversight of progress in meeting the scientific goals
- Arranges regular meetings of the national lattice gauge theory community to describe progress, and to obtain input
- Oversees the national lattice gauge theory community's SciDAC grants and provides coordination between the work done under those grants and in the current Program

- Appoints the members of the Scientific Program Committee

4.1.7 Scientific Spokesperson

The Chair of the Executive Committee serves as the Scientific Spokesperson for the LQCD Research Program.

Responsibilities

- Determines scientific goals and required computational infrastructure together with the USQCD Executive Committee
- Chairs the USQCD Executive Committee

Interactions of the Spokesperson:

- Principal point of contact to DOE on scientific matters related to the Program
- Presents the Program's scientific objectives to the DOE, its review committees, and its advisory committees
- Liaison between the Executive Committee and the CPM, relating the Executive Committee's priorities to the CPM, and transmitting the CPM's progress reports to the Executive Committee

4.1.8 Scientific Program Committee

The charter of the Scientific Program Committee (SPC) is to assist the Executive Committee in providing scientific leadership for the LQCD infrastructure development efforts. This committee monitors the scientific progress of the effort and provides leadership in setting new directions.

The Scientific Program Committee is charged with allocating time on the integrated hardware resources provided through the LQCD computing program. This committee has instituted the following allocation process. Once a year, proposals are solicited for the use of computational resources that are available to the user community during the allocation period July 1 to June 30. The Committee reviews the proposals and makes preliminary allocations based on its reviews. An open meeting of the user community is then held to discuss the proposals and the preliminary allocations. The Committee makes final allocations for each site following this meeting. The LQCD Site Managers are responsible for executing these allocations. The objective of this process is to achieve the greatest scientific benefit from the computing resources through broad input from the community. The committee is also charged with organizing an annual meeting of the user community to review Program progress and hardware roadmaps, scientific progress achieved with the infrastructure, and provide input on future computing needs.

Members of the Scientific Program Committee are appointed by the Executive Committee. The committee chair rotates every two years. Current members have staggered terms of four years. When a vacancy occurs, the open slot is filled by the Executive Committee. The current membership of the SPC is shown in Appendix B.

4.2 Program Communications

In addition to the interactions defined under Roles and Responsibilities, the following formal communications touchpoints are to occur annually, as appropriate:

Table 1. Program Communications Activities

Touch Point and Timing	Attendees	Actions and Goals
Early Acquisition Planning	CPM, Executive Committee	CPM leads discussion of acquisition planning, timeline. Goal: Concurrence on scope, non-technical considerations as input.
Late Acquisition Planning	CPM, Site Architects, Executive Committee	CPM presents acquisition plan. Goal: Concurrence on proposed acquisition plan.
Early Allocations Process	CPM, Site Architects, Scientific Program Committee	CPM presents performance benchmarks, deployed capacity. Goal: Address questions from SPC related to their Allocations process.
Late Allocations Process	CPM, Site Managers, Scientific Program Committee	SPC presents allocations including expectations for Class B, C allocations in coming year. Goal: Address questions from Site Managers related to their monitoring of allocations.

4.3 Interaction with Host Laboratory Management

Line management within the two host laboratories (BNL and FNAL) provides support to the program in several ways, including management and infrastructure support. Management authorities for DOE and senior management of the laboratories are shown in Figure 2. The primary flow of communication regarding LQCD program matters between the DOE Federal Program Director and laboratory management is through the LQCD Program Office.

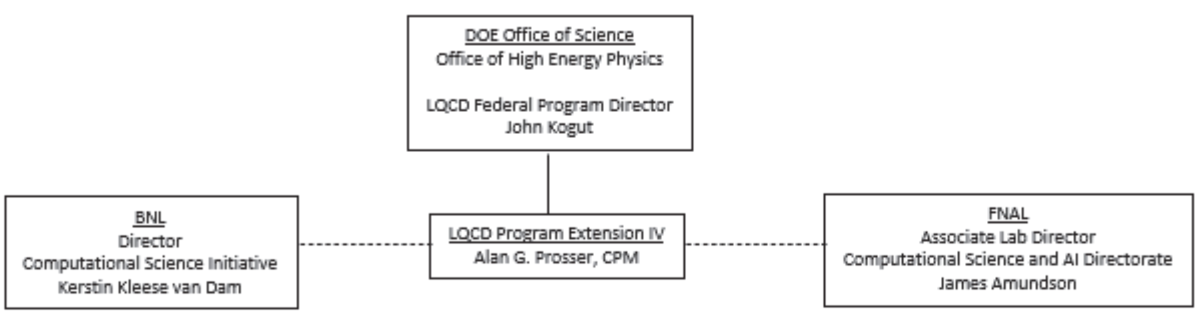

Figure 2: Communication flow between LQCD and Laboratory Computing Management

5 COST AND SCHEDULE MANAGEMENT

5.1 Work Breakdown Structure

The LQCD computing program is categorized as an OMB Exhibit 53 mixed life-cycle investment. Program work is organized into a Work Breakdown Structure (WBS) for purposes of planning, managing, and reporting activities. Work elements are defined to be consistent with discrete increments of work and the planned method of control. The LQCD program plan has two major WBS Level 2 components based on the work performed at each participating laboratory (BNL & FNAL). Under the Level 2 components are the following Level 3 components:

> **Steady-State Operations**: Includes all activities associated with the procurement, allocation, use, and monitoring of computing resources from the two host laboratories. The budget associated with Operations It supports the procurement of computing cycles and storage capacity provided by the host laboratories. It also supports Site Manager and Site Architect activities as defined above, along with a modest level of travel support.
>
> **Program Management**: Includes all activities associated with program management and oversight, as described above. The budget associated with Program Management supports salary costs for the Contractor Program Manager and Associate Contractor Program Manager, as well as a modest amount for travel and miscellaneous Program Office expenses.

Before the beginning of each fiscal year, a WBS is developed for the work to be performed in the coming year, with bases of estimates derived from past purchase records and effort reports. The WBS is developed with the concurrence of the Site Managers. Once defined, the WBS is baselined and a process for reporting status against the baseline is initiated. The WBS is developed and maintained using Microsoft Project.

Program milestones are defined in the WBS. Site Managers report the status of completion for each milestone to the CPM monthly. Any significant changes to milestone schedules are processed according to the change control procedure described later.

5.2 Program Milestones

Table 2 shows the Level 1 Program milestones that are tracked by the DOE Federal Program Director and Program Monitor. These milestones are also defined and tracked in the WBS. The target levels for new computing capacity deployed and aggregate computing delivered are defined in Appendix D - Computing Facility Performance Metrics.

Table 2: Level 1 Milestones

No.	Level 1 Milestone	Fiscal Year
1	FNAL Computer architecture planning for FY25 hardware expansion complete & reviewed	Q2 FY25
2	FNAL Procurement of *Combined Resources* in FY25	Q4 FY25
3	FNAL Target level of aggregate *Combined Resource* computing deployed & delivered in FY25	Q2 FY26

4	BNL Computer architecture planning for FY26 hardware expansion complete & reviewed	Q2 FY26
5	BNL Procurement of *Combined Resources* in FY26	Q4 FY26
6	BNL Target level of aggregate *Combined Resource* computing deployed & delivered in FY26	Q2 FY27
7	FNAL Computer architecture planning for FY27 hardware expansion complete & reviewed	Q2 FY27
8	FNAL Procurement of *Combined Resources* in FY27	Q4 FY27
9	FNAL Target level of aggregate *Combined Resource* computing deployed & delivered in FY27	Q2 FY28
10	BNL Computer architecture planning for FY28 hardware expansion complete & reviewed	Q2 FY28
11	BNL Procurement of *Combined Resources* in FY28	Q4 FY28
12	BNL Target level of aggregate *Combined Resource* computing deployed & delivered in FY28	Q2 FY29
13	FNAL Computer architecture planning for FY29 hardware expansion complete & reviewed	Q2 FY29
14	FNAL Procurement of *Combined Resources* in FY29	Q4 FY29
15	FNAL Target level of aggregate *Combined Resource* computing deployed & delivered in FY29	Q2 FY30

In addition to these Level 1 milestones, the WBS contains lower-level milestones that provide the means for tracking progress at a more granular level. Table 4 contains an example of the type of Level 2 milestones contained within the WBS that are associated with project tracking of host laboratory procurement activities. Tracking of these activities is necessary to ensure laboratory procurement activities stay on schedule, since the SPC allocates usage on new computing systems based on schedules provided by the host laboratories. Slips in procurement schedules negatively impact the quantity of computing resources provided to USQCD researchers. Reductions in computing resource availability has the potential to reduce scientific output.

Table 3: Example of Level 2 Milestones in the WBS
Associated with Hardware Procurement Tracking Activities

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	

Progress against all milestones is tracked and reported by the LQCD Program Office. Site Managers at each host laboratory report the status of completion for each milestone to the Program Office monthly. Progress against Level 1 and Level 2 milestones is discussed with the DOE Federal Program Director during monthly progress conference calls.

5.3 Planned Funding Profile

The total funding commitment for LQCD is $15 million. The program is supported by the DOE Office of High Energy Physics. The HEP planned funding profile for LQCD is shown in Table 4.

Table 4: Planned Funding Profile for LQCD (in millions $US)

FY25	FY26	FY27	FY28	FY29	Total
3.04	3.04	3.04	3.04	3.04	15.2

Program funds are used to procure compute cycles and data storage capacity and provide labor support for steady-state operations (e.g., site management, system administration, hardware support, and deployment of LQCD software) and program management. In addition, funds will be put in place to support yearly training activities for students. Software development is not in the scope for the LQCD Program.

Each host site will continue to contribute support to the Program in the form of infrastructure facilities and equipment. Each host site also provides administrative and technical support and services to the program in areas such as environment, safety, and health (ESH&Q), cyber security, disaster planning and recovery, networking, procurement, financial management services, and administrative support. The Program contributes to the pool of funds at each site used to cover these costs, through the assessment of overhead charges by each host site in accordance with standard laboratory policies.

5.4 Operations Budget

5.4.1 Operations Budget Profile

The LQCD operations budget comprises new funding as described in Section 5.3 and unspent funds from LQCD-ext III that were carried forward into LQCD-ext IV. Unspent funds were the result of delays in the deployment of new institutional clusters. Costs for institutional cluster computing services begin when new systems are released to production. Accordingly, LQCD plans to start incurring costs based on planned release dates. When a new system release is delayed, it delays when computing services are available for use and when computing service costs are incurred. This situation occurred late in the LQCD-ext IV lifecycle. Unspent funds were carried forward and used to increase the quantity of compute services purchased in FY20-21.

The hardware project will also support a LQCD-ext IV traineeship program with access to HPC resources through the “junior investigator” proposal track and through facilitating mentoring connections outside the trainee’s home institution as appropriate. Funds are requested to support visiting scholarships for USQCD graduate students (beyond those involved in the LGT4HEP program) to intern at one of the two HEP supported LQCD project labs over the summer. This will allow for additional mentorship of USQCD students and will help establish their careers in this field.

In total, four summer training sessions per FY on the use of the LQCD clusters are requested. The trainees will intern at FNAL or BNL, wherever is the focus on the training in that particular year.

Table 5 shows the LQCD operations budget profile in terms of commonly recognized expenditure types, by fiscal year. The personnel budget provides salary support for Site Managers, Site Architects, and program management. All labor cost estimates are based on fully loaded average labor rates at the host laboratories and have been inflated using an annual escalation rate of 3%. The travel budget covers costs for the Program Office, site architects and site managers to participate in annual DOE reviews and the USQCD All-hands Meeting. The compute services budget covers the cost of computing cycles delivered from the institutional clusters. The data storage services budget covers the cost of disk and tape storage. Indirect charges will be applied according to agreements established between the Program and the host laboratories and documented in approved MOUs.

Table 5: Operations Budget Profile by Expenditure Type ($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

Figure 4 shows the proportional cost breakdown by expenditure type. Approximately 90% of the total budget will be allocated to new compute and storage hardware. The level of personnel support is based on past operating experience. Program funds allocated to support travel have been kept to a minimum, with budgeted levels based on and consistent with past operating experience.

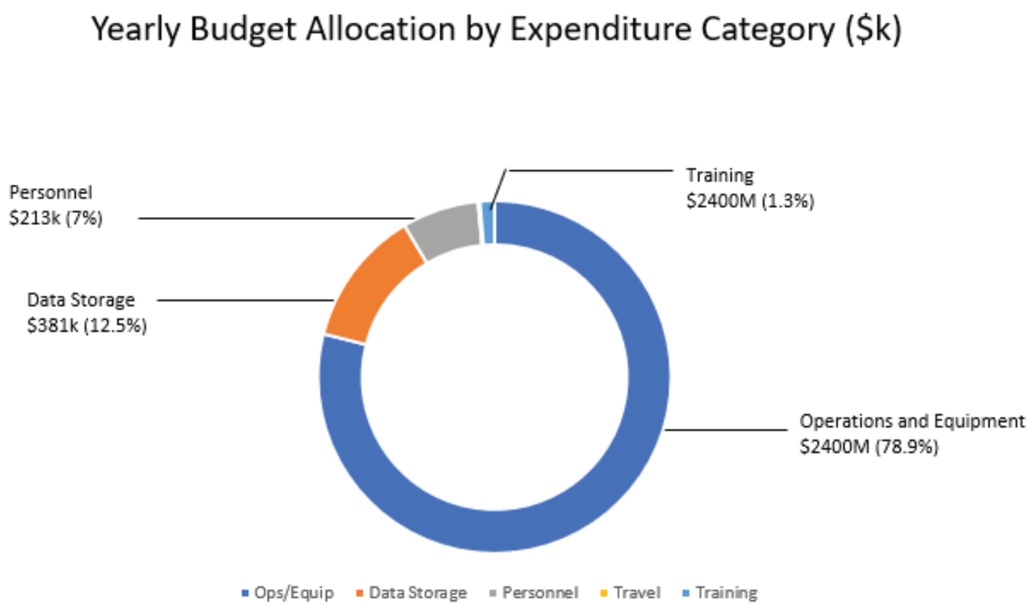

Figure 4: LQCD Yearly Program Budget Fraction by Expenditure Type

Figure 5 shows in graphical form the data presented in Table 5. The budgets for personnel and storage services are relatively flat, with a small upward trend due to inflation at a planned rate of 3% per year. The budget for travel is not shown because it is not visible given its small size. The budget for computing services in the first two years is greater than subsequent due to the carryover discussed above. Carryover funds were used to purchase additional computing services.

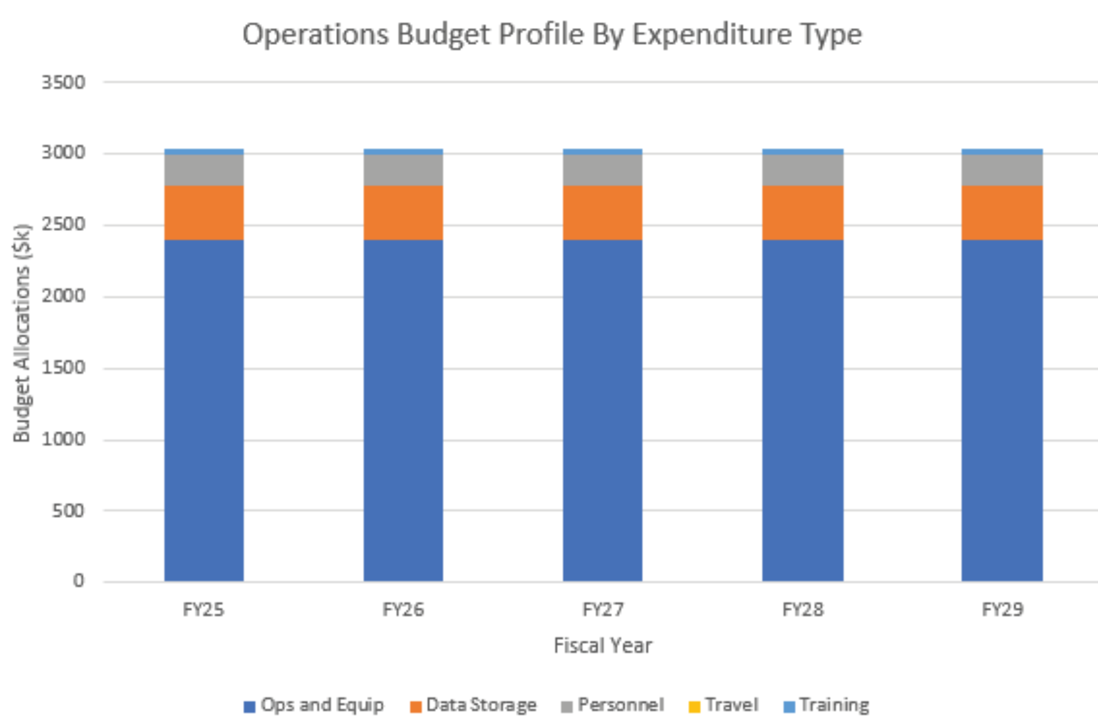

Figure 5: Operations Budget Profile by Fiscal Year

5.4.2 Management Reserve

No funds have been set aside for management reserve. The program is managed to fit within the budget guidance. Any unplanned cost increases in any category will be offset by decreasing the budget in other categories. The categories with the most built-in contingency are the computing and storage services budgets. Given that decreasing budget in these categories directly impacts the amount of computing and storage delivered to the science program, all costs are carefully managed.

5.4.3 Deployment Performance Planning

In each year of the LQCD program, the LQCD team selects the most cost-effective mix of computing and storage services that meets the needs of the science program. The selection of "node-hrs delivered" from the available cluster portfolio and the volume of data storage capacity contracted is constrained by the available budget.

All hardware procurements executed by the host laboratories utilize firm fixed-price contracts and are "built-to-cost" in accordance with approved budgets. Given fixed budgets, the precise number of processors procured is determined by the purchase price of systems and network equipment in that year. Variation in purchase price of these components, from the estimates used in the budget, results in greater or lesser computing capability from the estimated value. Variation in performance of the components from the estimates will also result in greater or lesser computing capability. The resulting performance risk is managed by the fact that the scope of the Program is fluid; small negative variances in available computing capability and/or capacity may result in schedule delays in completing scientific computing Programs. Large negative variances will prevent the achievement of computing goals; these may trigger review and modification of the USQCD scientific program, such as through changes or elimination of allocations of computing resources to specific science projects.

The risk of large performance variances is minimized using conservative projections in the performance of each future system development. Allocations of computing resources, and the planning of the USQCD scientific program, will be based upon these conservative estimates.

Table 6: Yearly Planned Delivered Computing Performance

	FY 2025	**FY 2026**	**FY 2027**	**FY 2028**	**FY 2029**
Planned delivered performance for FNAL (Tflop/s-yr)	204	312	244	390	373
Planned delivered performance for BNL (Tflop/s-yr)	100	100	244	244	451
Combined Planned delivered perfprmance (Total) (Tflops/s-yr)	304	412	488	634	824

Table 7: Integrated Planned Delivered Computing Performance

	FY 2025	**FY 2026**	**FY 2027**	**FY 2028**	**FY 2029**

Planned integrated performance for FNAL (Tflop/s)	204	516	760	1150	1523
Planned integrated performance for BNL (Tflop/s)	100	200	444	688	1139
Combined Planned Integrated perfprmance (Total) (Tflop/s)	304	716	1204	1838	2662

Performance of New System Deployments, and Integrated Performance (DWF+HISQ averages used). Integrated performance figures assume an 8300-hour year. The delivered performance figures (conventional/accelerated) shown in each year reflect the total across all sites.

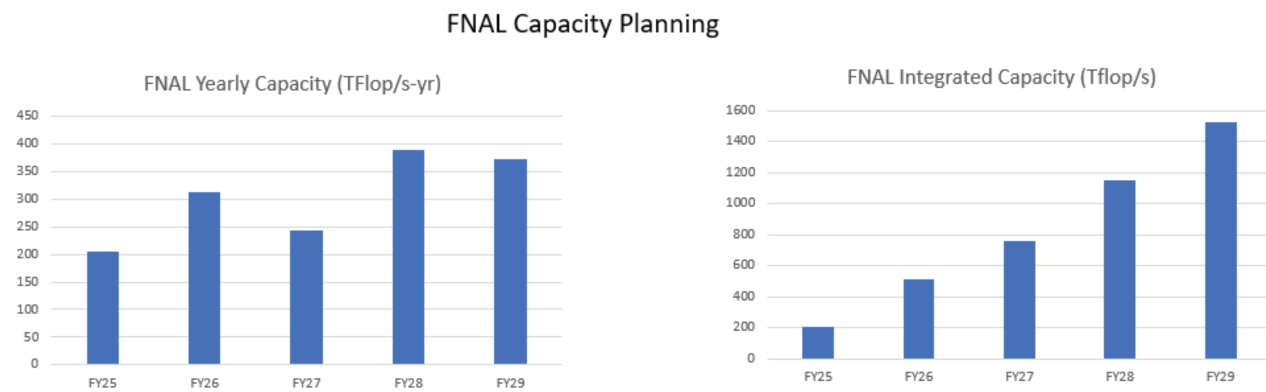

Figure 6: FNAL Yearly and Integrated Planned Capacity

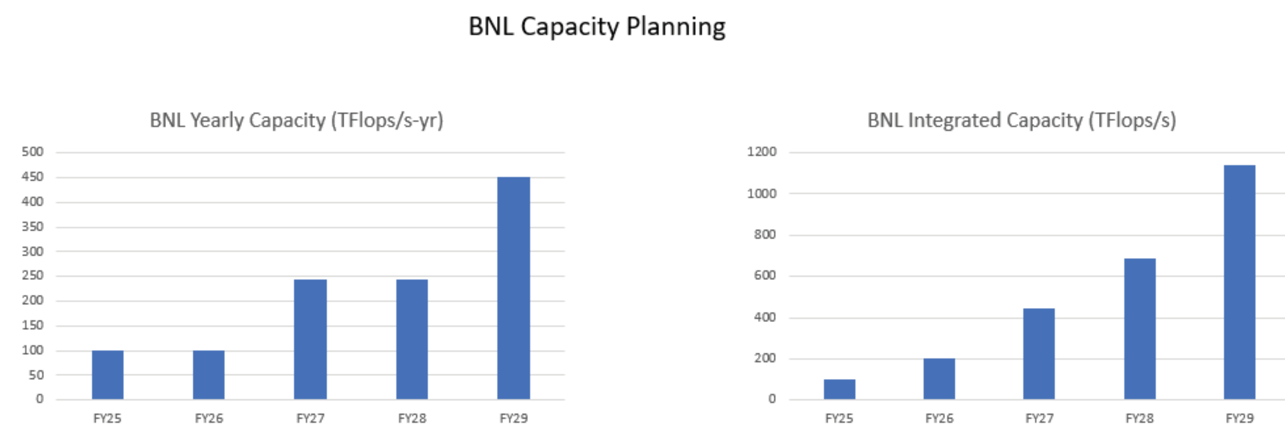

Figure 7: BNL Yearly and Integrated Planned Capacity

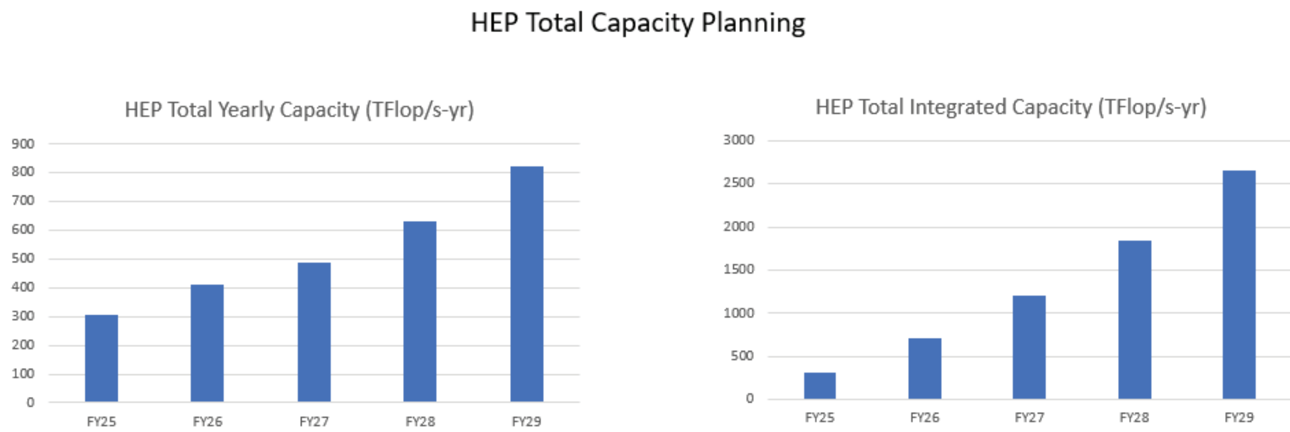

Figure 8: Combined (BNL+FNAL) Yearly and Integrated Planned Capacity

In each year, the host laboratories procuring new clusters will continue to build to cost and LQCD will commit to procuring computing cycles in accordance with the approved budget. It is estimated that approximately $1.68M will be available yearly to the relevant institution making that year's hardware procurement.

5.5 Cost and Schedule Management Controls

Overall performance at the two host laboratories is managed under the terms of the performance-based management contract with the DOE. Under these terms, laboratories are expected to integrate contract work scope, budget, and schedule to achieve realistic, executable performance plans. The table in Appendix C lists all facility performance metrics for the entire LQCD program. The metrics in these tables are associated with a $15 million Program budget. Should the Program budget change because of changes in available funding, performance metrics will be revised accordingly.

Following existing financial and operational procedures and processes at FNAL and BNL, the program has implemented methods of collecting and analyzing program performance data. The LQCD Program Office is responsible for the overall management of the Program and for implementing controls to ensure that cost, schedule, and technical performance metrics are met.

Memoranda of Understanding (MOU) are executed between the Program and the participating laboratories that detail work scope, level of funding, and the in-kind support provided to the Program by the host laboratories.

The LQCD Program Office has implemented a performance-based management system in which cost, and effort data are collected from both laboratories and analyzed monthly. Site Managers are responsible for tracking cost and schedule elements, and for reporting these to the CPM monthly. The CPM prepares and reviews monthly cost and schedule performance data against schedule, cost, and technical goals, and reports the result to the CPM. Every month the CPM reports overall cost, schedule, and technical performance to the Federal Program Director.

Technical performance is monitored throughout the Program to ensure conformance to approved functional requirements. Design reviews and performance testing of the completed systems are used to ensure that equipment and systems meet functional requirements.

On an annual basis, the DOE Office of High Energy Physics organizes an external review of Program performance. The review typically covers aspects of scientific, technical, cost, and schedule performance against goals. Results are recorded in a written report; all recommendations are carefully considered and implemented as appropriate. The CPM is responsible for preparing a document summarizing the Program's response to each recommendation.

6 PROGRAM MANAGEMENT

6.1 Security Management

The institutional clusters at the host laboratories are contained within computing enclaves of the host laboratory. Each computing enclave is protected according to the procedures implemented by the corresponding laboratory. The LQCD Program Office maintains copies of the Certification and Accreditation documents for each laboratory.

Performance is monitored by the DOE site office at each laboratory, in accordance with the requirements specified in the contracts between the DOE and the respective contracting agencies (Brookhaven Science Associates (BSA) for BNL and Fermi Research Alliance (FRA) for FNAL). These contracts include requirements for compliance with pertinent government (NIST 800-53) and DOE Computer Security policies (e.g., DOE O 205.1 Department of Energy Cyber Security Management Program). At each laboratory, contractor security procedures are monitored, verified, and validated by numerous external entities including: 1) DOE-OCIO, 2) DOE Office of Performance Management and Oversight Assessment, 3) the DOE-IG, and 4) external reviews.

6.2 Privacy Management

None of the computing systems being used by LQCD contain, process, or transmit personally identifiable information. These systems are not privacy systems of record.

6.3 Risk Management

Risk management is an ongoing activity accomplished by continuously identifying, analyzing, mitigating, and monitoring risks that arise during program execution. Risk is a measure of the potential of failing to achieve overall program objectives within the defined scope, cost, schedule, and technical constraints. The purpose of risk analysis is not solely to avoid risks, but to understand the risks associated with the program and devise strategies for managing them.

The final responsibility for risk management rests with the CPM, in consultation with the USQCD Executive Committee and LQCD Site Managers. However, effective risk management is a multi-step process that requires the continuous involvement of all program team members. The LQCD team plans for and tracks operational, technical, and financial risks as defined in the LQCD Risk Management Plan. The Risk Management Plan is reviewed and updated whenever changing conditions warrant a review and revision of the risk register. The Risk Management Plan is also reviewed on a periodic basis to review the status of identified risks and to consider the potential existence of new risks. During these reviews, the risk register is updated by adding and/or closing risks, and initiating and revising risk mitigations, as needed.

A full discussion of potential risks and mitigation strategies is contained in LQCD Risk Management Plan. Performance risks associated with computing and storage systems are estimated to be low due to the use of commercial off-the-shelf components.

The distributed nature of LQCD computing partially mitigates the risk of natural disasters. Additionally, the Program employs a disaster recovery strategy for valuable data by storing data files redundantly at two different locations. Although the equipment at each facility is not insured against disasters, standard disaster recovery protections are provided by each laboratory.

6.4 Quality Assurance

Quality is defined as the "fitness of an item or design for its intended use" and Quality Assurance (QA) as "the set of actions taken to avoid known hazards to quality and to detect and correct poor results." Program personnel follow quality control procedures established at the two host laboratories. In addition, the Program has put into place various methodologies to monitor and improve quality, as described in the following document: *Quality Assurance Plan for the LQCD Computing Program*. As new systems are brought online by the host laboratories, a series of tests are conducted to verify quality at the system level. Once a cluster is released to "user-friendly mode" usage, LQCD codes are run and performance measured to verify operational readiness, before full-production use commences.

6.5 Program Oversight

On a monthly basis, the LQCD Program Office prepares and presents a program status report to the Federal Program Director. The purpose of the monthly meeting is to inform the Federal Program Director of cost, schedule, and technical performance, and discuss other issues related to program execution.

To determine the health of the program and to provide guidance on progress, an annual DOE Office of High Energy Physics progress review is held, generally in May. During this review, past performance and future are presented and reviewed. Review results are presented in written form and transmitted to the Contractor Program Manager via the DOE Office of High Energy Physics. The CPM is responsible for responding to all review recommendations.

7 ENVIRONMENT, SAFETY AND HEALTH

The LQCD program is a collaborative effort among the two DOE-sponsored laboratories with stringent environment, safety, and health (ES&H) policies and programs. LQCD integrates ES&H into all phases of the program (planning, acquisition, operations, and maintenance) using appropriate procedures defined by the participating laboratories. All individuals supported by Program funds follow procedures specific to the laboratory at which they work.

The LQCD program follows the five core functions associated with integrated safety management:

1. Define work and identify the potential hazards
2. Analyze potential hazards and design the equipment or activities to appropriately mitigate or eliminate those hazards.
3. Establish controls for hazards that cannot be eliminated through design features
4. Perform work in accordance with the procedures

5. Review the effectiveness of the hazard analyses and controls and provide feedback for improvement.

Line management at each laboratory retains supervisory authority of their personnel and responsibility for the safety of work performed. Line management keeps the CPM informed about their laboratory's management and ES&H organization structures. Any safety concerns by personnel assigned to the LQCD program are to be communicated to the line management where the concern occurs and if appropriate, the employee's home laboratory or university.

Site Managers at each laboratory work with safety officers at their laboratory to ensure that any hazards found are documented according to plans and procedures of the laboratory and mitigated appropriately. Information pertaining to these hazards is documented as needed using appropriate safety documentation guidelines for the laboratory. Also, laboratory personnel receive specific training required to perform their job in a safe and proper manner.

There is no direct construction activity under the direction and control of this Program. Any facility upgrades or improvements involving construction activities will be managed by the host laboratory. The LQCD program will comply with all necessary rules, regulations, policies, and procedures related to working in or around construction areas. Any required NEPA reviews related to facility upgrades associated with the LQCD computing facilities will be coordinated and/or conducted by the host laboratory.

Appendix A: Integrated Program Team

LQCD Federal Program Director (HEP)	John Kogut (chair)
Contractor Program Manager (CPM)	Alan Prosser
BNL Site Manager	Zhihua Dong
BNL Site Architect	Costin Caramarcu
BNL Scientific Consultant	Chulwoo Jung
FNAL Site Manager	Lydia Brynmoor
FNAL Site Architect	Tim Skirvin
USQCD Executive Committee Chair	Robert Edwards
USQCD Executive Committee Deputy	Thomas Blum
USQCD Scientific Program Committee Chair	Peter Petreczky
USQCD Scientific Program Committee Deputy	(open)

Note: As of October 1st, 2021

- Robert Edwards is the Spokesperson and Chair of EC
- Thomas Blum is the Deputy Spokesperson and Deputy Chair of EC
- Andreas Kronfeld has returned to an EC member and become the FNAL Scientific Consultant within the IPT

Appendix B: Committees and Members

USQCD Executive Committee

Robert Edwards (chair), Tom Blum (deputy), Norman Christ, Carleton E. DeTar, William Detmold, Anna Hasenfratz, Huey-win Lin, Swagato Mukherjee, Kostas Orginos, Peter Petreczky (SPC Chair)

Scientific Program Committee

Peter Petreczky (chair), James Simone, (SPC deputy chair), Martha Constantinou, George Fleming, Christopher Kelly, Stefan Meinel, Sergey Syritsyn

Appendix C: Computing Facility Key Performance Indicators (KPIs)

ID	Fiscal Year	Measurement Category	Measurement Indicator	Target	Actual Results	Rating
1	2025	Scientific Program Support	TF-Yrs. delivered towards the completion of the Scientific Program – *Combined Resources*	93TF-Yrs	Available in Q4 FY26	
2	2025	Responsiveness	% Of tickets responded to within three business days	≥95%	Available in Q1 FY26	
3	2025	Security and Privacy	Frequency of vulnerability scans performed at each site on nodes visible from the Internet	Vulnerability scans performed at least weekly at each host site (minimum of 52 scans per year per site)	Available in Q1 FY26	
4	2025	Reliability and Availability	% Of average machine availability across all LQCD computing sites	≥95%	Available in Q1 FY26	
5	2025	Quality of Service Delivery	Customer satisfaction rating (Customer's rate satisfaction with the service provided on a scale of 1 to 5)	≥92%	Available in Q1 FY26	
6	2026	Scientific Program Support	TF-Yrs. delivered towards the completion of the Scientific Program – *Combined Resources*	77 TF-Yrs.	Available in Q1 FY27	
7	2026	Responsiveness	% Of tickets responded to within three business days	≥95%	Available in Q1 FY27	
8	2026	Security and Privacy	Frequency of vulnerability scans performed at each site on nodes visible from the Internet	Vulnerability scans performed at least weekly at each host site (minimum of 52 scans per year per site)	Available in Q1 FY27	
9	2026	Reliability and Availability	% Of average machine availability across all LQCD computing sites	≥95%	Available in Q1 FY27	
10	2026	Quality of Service Delivery	Customer satisfaction rating (Customer's rate satisfaction with the service provided on a scale of 1 to 5)	≥92%	Available in Q1 FY27	
11	2027	Scientific Program Support	TF-Yrs. delivered towards the completion of the Scientific Program – *Combined Resources*	77 TF-Yrs.	Available in Q1 FY28	
12	2027	Responsiveness	% Of tickets responded to within three business days	≥95%	Available in Q1 FY28	
13	2027	Security and Privacy	Frequency of vulnerability scans performed at each site on nodes visible from the Internet	Vulnerability scans performed at least weekly at each host site (minimum of 52 scans per year per site	Available in Q1 FY28	
14	2027	Reliability and Availability	% Of average machine availability across all LQCD computing sites	≥95%	Available in Q1 FY28	
15	2027	Quality of Service Delivery	Customer satisfaction rating (Customer's rate satisfaction with the service provided on a scale of 1 to 5)	≥92%	Available in Q1 FY28	
16	2028	Scientific Program Support	TF-Yrs. delivered towards the completion of the Scientific Program – *Combined Resources*	92 TF-Yrs.	Available in Q1 FY29	
17	2028	Responsiveness	% Of tickets responded to within three business days	≥95%	Available in Q1 FY29	
18	2028	Security and Privacy	Frequency of vulnerability scans performed at each site on nodes visible from the Internet	Vulnerability scans performed at least weekly at each host site (minimum of 52 scans per year per site)	Available in Q1 FY29	
19	2028	Reliability and Availability	% Of average machine availability across all LQCD computing sites	≥95%	Available in Q1 FY29	
20	2028	Quality of Service Delivery	Customer satisfaction rating (Customer's rate satisfaction with the service provided on a scale of 1 to 5)	≥92%	Available in Q1 FY29	
21	2029	Scientific Program Support	TF-Yrs. delivered towards the completion of the Scientific Program – *Combined Resources*	98 TF-Yrs.	Available in Q1 FY30	
22	2029	Responsiveness	% Of tickets responded to within three business days	≥95%	Available in Q1 FY30	
23	2029	Security and Privacy	Frequency of vulnerability scans performed at each site on nodes visible from the Internet	Vulnerability scans performed at least weekly at	Available in Q1 FY30	

ID	Fiscal Year	Measurement Category	Measurement Indicator	Target	Actual Results	Rating
				each host site (minimum of 52 scans per year per site)		
24	2029	Reliability and Availability	% Of average machine availability across all LQCD computing sites	≥95%	Available in Q1 FY30	
25	2029	Quality of Service Delivery	Customer satisfaction rating (Customer's rate satisfaction with the service provided on a scale of 1 to 5)	≥92%	Available in Q1 FY30	

Appendix D: Controlled Documents

The set of documents submitted to DOE are designated as controlled Program documents. These documents are tracked using DocDB, the Document Database Control system managed by the Fermilab Core Computing Division. The LQCD document control area is password protected and only accessible by the IPT. Access requests should be made to the CPM.

The following are considered controlled documents, with formal version control and signature approval.

1. Program Execution Plan
2. Risk Management Plan
3. Quality Assurance Program
4. Acquisition Strategy
5. Annual Acquisition Plans
6. Certification and Accreditation Document
7. Cyber Security Plan (formerly called the Security Vulnerability Assessment Report)

In addition to controlled documents, the following documents are also stored in DocDB under limited access.

1. Memoranda of Understanding
2. DOE Annual Review Reports

Appendix E. Historical Background

The development and operation of a large-scale computing facility dedicated to the study of quantum chromodynamics (QCD) plays an important role in expanding our understanding of the fundamental forces of nature and the basic building blocks of matter.

Since 2000, members of the United States lattice gauge theory community have worked together to plan the computational infrastructure needed for the study of QCD. In February 2003, the lattice QCD computational infrastructure effort was reviewed by a panel of physicists and computer scientists chaired by Frank Wilczek. One of its conclusions was: "The scientific merit of the suggested Program is very clearly outstanding." Since then, the High Energy Physics Advisory Panel (HEPAP) and the Nuclear Science Advisory Committee (NSAC) have both recommended that DOE funds should be allocated for dedicated computer hardware for lattice QCD simulations because of the importance of the calculations to their respective fields. Thus, the scientific need for this Program has been validated by leading experts in high energy and nuclear physics.

The LQCD-research Program continues to meet the planning, budgeting, and reporting criteria for an OMB Exhibit 53 IT investment; therefore, this classification remains intact.

In fall 2017, the LQCD Program began transitioning from a dedicated compute cluster model to a new cooperating model in which program funds were used to purchase computing cycles from institutional clusters (ICs) operating at BNL and FNAL. LQCD began purchasing compute cycles from BNL in January 2017 and Fermilab in FY18.

Beginning in FY18 and concluding in FY19, LQCD partnered with FNAL on the design and implementation of a new institutional cluster. Building on an acquisition strategy and annual planning process that was used by LQCD for many years, a joint committee was formed to understand computing needs, create viable options, and supply a recommendation to LQCD and FNAL management regarding a preferred solution for the FY19 Fermilab institutional cluster procurement.

The Acquisition Planning Committee was asked to provide input into the FY18 computing hardware planning process. Specific activities included:

1. Gather and review computing needs of the LQCD, CMS, neutrino Program user groups
2. Understand the capabilities of the existing hardware portfolio available to LQCD
3. Assess the vendor landscape for viable architecture options
4. Prepare an Alternatives Analysis of viable options
5. Present a recommendation, with technical design and cost estimates to the LQCD and FNAL computing leadership on the most cost-effective hardware solution

A strong alignment of the LQCD hardware portfolio with anticipated computing needs for the USQCD scientific Program were provided by completing the above tasks. This also assisted with the alignment of the new FNAL Institutional Cluster with anticipated future computing needs for the scientific Program.

Based on recommendations and data the team agreed that the preferred solution for the new FNAL institutional cluster was to deploy and commission a conventional cluster of 89 nodes and a GPU accelerated cluster of 2 hosts (4 GPUs per host total 8 GPUs capable of delivering respectively at least 39 TF and 11 effective TF, 50 TFlops total, with at least a memory capacity of 12TB. The committed hardware funds were used in exchange for compute cycles on the new Fermilab Institutional Cluster machine.

8 Bibliography

[1] A. Kronfeld, "Lattice QCD and Particle Physics," 2022. [Online]. Available: https://doi.org/10.48550/arXiv.2207.07641.

[2] P. Boyle, "Lattice QCD and the Computational Frontier," 2022. [Online]. Available: https://doi.org/10.48550/arXiv.2204.00039.