

**Acquisition Strategy
for the
SC Lattice QCD Computing Project Extension II
(LQCD-Ext II)**

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

Operated at
Brookhaven National Laboratory
Fermi National Accelerator Laboratory
Thomas Jefferson National Accelerator Facility

for the
U.S. Department of Energy
Office of Science
Offices of High Energy and Nuclear Physics

Version 1.3

April 20, 2015

PREPARED BY:
Don Holmgren, FNAL

CONCURRENCE:



April 24, 2015

William N. Boroski
LQCD Contractor Project Manager

Date

**Acquisition Strategy
Change Log**

Revision No.	Description/ Pages Affected	Effective Date
Revision 0.0	Adopt from SC LQCD-ext to create initial version for CD1	February 18, 2014
Revision 1.0	Update with information from FY14 cluster procurement	June 19, 2014
Revision 1.1	Expanded discussion of Xeon Phi processors	June 27, 2014
Revision 1.2	Performance goals reflect budget revision	July 2, 2014
Revision 1.3	Updates to reflect actual 2014 buys, Intel model names	April 20, 2015

Table of Contents

Introduction.....	1
Previous LQCD Computing Projects.....	2
Overview of LQCD-Ext II Project Deployments	2
Design Considerations and Strategies for LQCD-Ext II Deployments	4
Compute Nodes.....	4
High Performance Network	5
Service Networks	6
Network Plan	6
File I/O	6
Procurement Strategy.....	7
Strategy for the FY2014 LQCD-Ext Deployment.....	8

Introduction

The Lattice QCD Computing Extension Project (LQCD-Ext II) develops and operates new and existing systems in each year from FY2015 through FY2019. These computing systems are deployed at Fermilab (FNAL), Jefferson Lab (JLab), and Brookhaven (BNL). Table 1 shows the planned total computing capacity of the new deployments, as well as the planned delivered (integrated) performance.

	FY 2015	FY 2016	FY 2017	FY 2018	FY 2019
Planned (FY15- FY19) computing capacity of new deployments, Tflop/s	0	49	66	134	172
Planned (FY15-FY19) delivered performance (JLab + FNAL + BNL), Tflop/s-yr	180	135	165	230	370

Table 1 – Performance of New System Deployments, and Integrated Performance (DWF+HISQ averages used). Integrated performance figures use an 8000-hour year. The capacity and delivered performance figures shown in each year sum the conventional (TFlop/s and TFlop/s-yr) and accelerated (effective TFlop/s and effective TFlop/s-yr) resources deployed and operated. All deployment figures assume that 50% of the annual hardware budget is used to purchase accelerated hardware, and 50% to purchase conventional hardware.

In all discussions of conventional cluster and supercomputer performance, unless otherwise noted, the specified figure reflects an average of the sustained performance of domain wall fermion (DWF) and highly improved staggered quark (HISQ) algorithms. For accelerated clusters, the figure is based on the USQCD rating of an NVIDIA K20 model GPU of 157 effective GFlop/s. These performance figures are defined in the document “System Description for the SC Lattice QCD Computing Extension II Project”.

All LQCD-Ext II Project hardware procurements will utilize firm, fixed-price contracts. Hardware purchases will use contracts with vendors specializing in COTS hardware. The steady-state operations of the project computing facilities are performed by the three host laboratories, each of which is a government-owned contractor-operated facility.

In each year of the project, the hardware that best accomplishes the scientific goals for LQCD calculations will be purchased. In FY2015, the project will not deploy new hardware but will operate the hardware deployed at the three labs during the prior project (SC LQCD-ext, FY2010-FY2014). This hardware consists of conventional clusters, GPU-accelerated clusters, and BlueGene/Q hardware. Starting in FY2016 and continuing through FY2019, the project will perform at most one annual purchase and deployment of new hardware. These individual deployments are anticipated to consist of a combination of conventional and accelerated cluster hardware, similar to the deployments during the SC LQCD-ext project. To save manpower and overhead (G&A) costs, the project will likely utilize combined procurements in FY2016 and FY2017, and in FY2018 and FY2019. In a combined procurement, a single purchase contract is awarded in the first year; the contract includes option pricing for expansion of the system during the second year.

In the rest of this document we first discuss the design considerations and strategies that we will use for the procurements of the LQCD-Ext II Project. We then discuss the process by which the project will determine the best hardware combination for the FY2016 acquisition. This document will be updated each year to concentrate on the upcoming year’s hardware acquisition.

Previous LQCD Computing Projects

From FY2006 through FY2009, the DOE High Energy Physics (HEP) and Nuclear Physics (NP) program offices funded the DOE Office of Science LQCD Computing Project (SC LQCD). The total project cost of \$9.2M funded the deployment and operation of four clusters at Jefferson Lab and Fermilab and the QCDOC supercomputer at Brookhaven, and the operation of several SciDAC LQCD clusters at JLab and FNAL acquired in 2003 through 2005.

From FY2010 through FY2014, the SC LQCD project was extended (SC LQCD-ext), with a total project cost of \$18.15M. This project funded the continued operation of the QCDOC at BNL and the clusters deployed by the SC LQCD project, plus the deployment and operation of four conventional clusters, three GPU-accelerated clusters, and a half-rack of BlueGene/Q hardware. In FY2013, the operation of two conventional clusters and two GPU-accelerated clusters purchased and deployed at JLab in 2009-2011, funded as part of the American Recovery and Reinvestment Act of 2009 via the NP program office, were added to the SC LQCD-ext project.

The systems deployed during SC LQCD-ext and the ARRA project that will be operated by the SC LQCD-ext II project beginning in FY2015 are as follows:

- “Ds”, based on quad-socket eight-core Opteron processors and QDR Infiniband at FNAL (2010/2011)
- “12s”, based on dual-socket eight-core Intel “Sandy Bridge” processors and QDR Infiniband at JLab (2012)
- “Bc”, based on quad-socket eight-core Opteron processors and QDR Infiniband at FNAL (2013)
- “Pi0”, based on dual-socket eight-core Intel “Ivy Bridge” processors and QDR Infiniband at FNAL (2014 with additional nodes purchased 2015)
- “10g”, based on NVIDIA GTX-480, GTX-580, GTX-690, and C2050 GPUs in Intel-processor-based hosts at JLab (2010)
- “11g”, based on NVIDIA C2050 GPUs in Intel-processor-based hosts at JLab (2011)
- “Dsg”, based on NVIDIA M2050 GPUs in Intel-processor-based hosts at FNAL (2012)
- “12k”, based on NVIDIA K20 GPUs in Intel-processor-based hosts at JLab (late 2012)
- “Pi0g”, based on NVIDIA K40 GPUs in Intel-processor-based hosts at FNAL (2014)
- BlueGene/Q half-rack, deployed at BNL (2013)

A number of the systems listed above (Ds, Dsg, 10g, 11g) will be operated during part or all of FY2015, but will be decommissioned as new hardware comes online in FY2016.

Overview of LQCD-Ext II Project Deployments

As the LQCD-Ext II project begins in FY2015, based on experience from the LQCD-ext project, the most effective hardware for the planned calculations will be a combination of commodity clusters with Infiniband interconnects and similar clusters with accelerators. The LQCD-ext commodity clusters were based on either AMD (“Opteron”) or Intel x86_64 processors, and the accelerated clusters were based on Intel x86_64 host machines containing two or four NVIDIA-brand GPU accelerators.

From 2009 forwards, a growing fraction of USQCD computing has been able to take advantage of highly optimized LQCD codes written for NVIDIA GPUs in the “CUDA” environment. The clusters deployed under the ARRA LQCD and LQCD-ext projects were based on a combination of NVIDIA “gaming” GPU cards and “Tesla” cards. The Tesla cards include features of advantage to some areas of high

performance computing, such as memory error detection and correction (“ECC memory”) and better double precision floating point performance relative to the considerably less expensive gaming cards. The Tesla cards also supported larger local memory and so could be used for larger problems. For the set of LQCD workloads that do not require ECC memory, the gaming cards provided much better price/performance, both because of lower cost but also because of higher clock rates and greater memory bandwidth compared to the Tesla cards. As the NVIDIA GPU product lines have evolved, features important for high performance computing have been restricted to the Tesla line. We anticipate in LQCD-ext II that GPU deployments will be restricted to the Tesla models; however, we will evaluate the applicability of future gaming cards during the course of the project.

Since late 2012, Xeon Phi accelerators have been available from Intel. In early 2014, LQCD production code that could exploit Phi accelerators was just becoming available, with 18M hours used on Stampede at TACC. We anticipate that future generations of the Xeon Phi may be very cost effective for LQCD production. Once Phi processors are no longer deployed as accelerators but rather as the system processor (“self-hosted”) this architecture may provide a more agile software development platform compared to the CUDA environment. In this document for the purposes of categorization, clusters based on computers with self-hosted Phi processors are considered to be commodity clusters. Depending upon the availability of software, to optimize the scientific potential of the portfolio of systems operated by the LQCD-ext II project, deployment of new commodity clusters may include systems based on conventional processors and/or systems based on “advanced” processors such as self-hosted Xeon Phi processors.

During the LQCD-ext project, in FY2013, the IBM BlueGene/Q was available for purchase. Based on the project’s alternatives analysis completed in early August of 2013 we determined that the combination of hardware that would best optimize the total portfolio of dedicated hardware consisted of a half-rack of BG/Q, to be deployed at BNL, and a conventional Infiniband cluster at Fermilab. In the LQCD-ext II project, a similar process will be used in each of FY2016-FY2019 to determine the hardware to be deployed which will best optimize the total portfolio of dedicated hardware. The input data to these alternative analyses will include scientific requirements, the anticipated mix of LQCD job types, the performance of available hardware, and the estimated lifecycle costs of the alternatives. The project will rely on information from the USQCD Executive and Scientific Program Committees to understand the scientific requirements and job demands. As in FY2013 in the LQCD-ext project, should non-cluster hardware suitable and cost effective for LQCD be available, such as the BlueGene/Q system purchased that year, such hardware will be included in the annual alternatives analysis.

All procurements will be performed by the host laboratory chosen for the particular hardware deployment. Such purchases will utilize firm, fixed-price contracts. The typical sequence for new deployments will be:

1. In consultation with the USQCD community (through the Executive Committee and Scientific Program Committee), determine anticipated usage profiles for new deployments (*e.g.*, distribution of job types and sizes, file I/O requirements)
2. Complete preliminary design
3. Issue a Request for Information (RFI) to likely vendors
4. Evaluate the RFI responses and complete a final design
5. Obtain host laboratory purchase approvals via the local requisition process
6. Issue a Request for Proposal (RFP) to likely vendors
7. Evaluate RFP responses and award purchase contract

8. Approve sample node and/or sample scalable unit (rack)
9. Test and approve vendor-integrated final system (acceptance test)
10. Operate final system in “friendly user” mode and tune the configuration
11. Release the final system to users

Design Considerations and Strategies for LQCD-Ext II Deployments

Compute Nodes

Lattice QCD codes are floating point intensive, with a high bytes-to-flops ratio (1.8 single precision, 3.6 double precision for SU(3) matrix-vector multiplies). When local lattice sizes exceed the size of cache, which is nearly always the case, high memory bandwidths are required.

At the start of LQCD-ext II, Intel “Ivy Bridge” and “Haswell” Xeon processors will provide the highest memory bandwidth of the available x86_64 options; the AMD Opteron processor product line, which had been very competitive for LQCD calculations prior to 2014, has not improved significantly in characteristics important to LQCD codes since about 2012. Based on processor roadmaps from both companies, we anticipate that in at least the first half of LQCD-ext II that the most cost effective clusters will be based on Intel processors.

Starting in 2008, various LQCD scientists implemented codes to run on NVIDIA graphics processing units (GPUs) using the “CUDA” extensions to the C and C++ languages. Typically these codes accelerated part of the overall computational work performed during LQCD configuration generation and analysis. Although very labor intensive to implement, these codes greatly accelerated those portions of LQCD computations, and GPU-accelerated clusters clearly can have greater cost efficiency than conventional Infiniband clusters for some of the calculations of interest. The “Dsg” cluster purchased at Fermilab with FY2011 funds utilizes NVIDIA Tesla M2050 GPUs. Starting in late calendar 2012, the successor to “Fermi,” “Kepler,” became available from NVIDIA, and the LQCD-ext “12k” cluster at Jefferson Lab was deployed using NVIDIA K20 “Kepler” accelerators. On LQCD codes, K20 accelerators have 1.5 times the performance of C2050 accelerators, when averaging the increase in performance of single and double precision codes. The K40 accelerators introduced in late 2013 achieve approximately 30% higher performance than K20 GPUs and have larger memories, further improving their usefulness.

NVIDIA GPUs are available in two forms: the “Tesla” series, which are intended for computations, and the “GTX” or “GeForce” series, which are intended for graphics. The more expensive Tesla GPUs have a number of numerical advantages, including error correcting memory (ECC) for the detection of memory errors and the correction of all single-bit errors, hardware double precision, and optimized GPU-to-GPU and GPU-to-Infiniband communications. The graphics GPUs tend to have less memory (1.5 GBytes, for example, compared to 6 or even 12 GBytes on recent Tesla models), but their memory bandwidth is higher due to faster clocking of the GPU processor, and their cost is much lower than the Tesla models. On Tesla models, use of ECC further reduces memory bandwidth and also decreases memory capacity. LQCD calculations have higher throughput on the graphics models because of their higher memory bandwidth. For calculations that can check on the validity of results, such as those dominated by Dirac operator inversions, the graphics cards are far more cost effective. One final advantage to Tesla models involves warranties: NVI worm9

DIA will replace Tesla GPUs that produce numerical errors, but the vendors of graphics cards in general will not. In the GPU acquisition at JLab in 2009 and 2010, errors using some of the graphics cards were occasionally observed on numerical calculations. Such errors can be easily detected and corrected on matrix inversions via an inexpensive correctness check, but cannot be tolerated for the more general calculations that make up a growing percentage of usage.

For those calculations to which GPUs can be applied, significant accelerations have been observed. Comparing the “JPsi” cluster with the “9g” and “10g” clusters, users in the fall of 2010 reported relative throughputs of between 5:1 and 15:1 when comparing a single GPU to a single “JPsi” node. This represents a significant increase in cost effectiveness.

Early work on the Intel Xeon Phi shows that they are capable of performance comparable to GPUs. Contemporary 2012 Tesla Kepler K20 and Xeon Phi 5110p (a.k.a. Knights Corner, or KNC) accelerators achieved Dirac inverter performance within a few percent of each other. Although the performance was good, these KNC accelerators were not as mature as the Tesla line, and the processors lacked useful features for general C/C++ use (e.g. the processors lacked out-of-order execution capability).

The next generation Xeon Phi part, referred to as Knights Landing, or KNL, will include out-of-order capability and will be self hosted – i.e. you will be able to construct a KNL machine entirely as a homogeneous Infiniband cluster of many-core processors, making it similar to Xeon clusters today. Code will not need to be re-written in CUDA, and optimizations to exploit KNL will be similar to the optimizations necessary to use all modern multi-core processors (x86, Power, ...).

From 2016 forward, the LQCD-Ext II expects to use a healthy fraction of the hardware budget for advanced systems, either GPU accelerated or many-core processors such as Xeon Phi. The portion to be used in any year will depend upon the scientific demand for this hardware, which in turn is related to the fraction of the LQCD calculations that can take advantage of these advanced architectures, i.e. constrained by the maturity of software.

High Performance Network

Based on the LQCD-ext clusters, Infiniband will be the likely choice for LQCD-ext II deployments. These clusters will use quad data rate (QDR) or faster Infiniband parts. From vendor roadmaps, by 2017 at least one competing architecture will be available from Intel and likely integrated into Intel Xeon and Xeon Phi processors. The project will evaluate such new architectures for performance and applicability.

During LQCD-ext, 36-port switches offered the best pricing per port for Infiniband. For the large clusters built in that project, leaf and spine designs were used, with a stack of 36-port switches employed for the spine. Because QDR HCA bandwidths exceeded the requirements for lattice QCD codes for some processors, such as the AMD Opteron “Magny-Cours” and “Abu Dhabi” processors used on the Ds and Bc clusters, oversubscribed designs were used as a cost savings. In a 2:1 design, for example, 24 computers attached to a 36-port switch, with the remaining 12 ports used to connect to the network spine. In LQCD-ext II we will continue these leaf and spine architectures and exploit oversubscription when suitable.

As node compute capacities increase, faster networking will be required. The project will deploy faster Infiniband or other network fabrics as needed. As of 2014, one of the Infiniband silicon vendors offers two models (FDR-10, FDR) of faster Infiniband silicon than QDR, and by 2015 that vendor will be selling EDR-rate components. The data rates for QDR, FDR-10, FDR, and EDR, respectively, are 32 Gbits/sec, 40 Gbits/sec, 54 Gbits/sec, and 100 Gbits/sec.

This trend towards higher bandwidths is being softened by newer more efficient communications avoiding algorithms such as multi-grid and related solvers. These algorithms help in addressing the much faster growth in processor performance than network bandwidth. At the start of this project, multi-grid on CPUs could actually outperform the older GPU accelerated algorithms on a per socket basis.

Service Networks

Although Infiniband supports TCP/IP communications, we believe that standard Ethernet will still be preferred for service needs. These needs include booting the nodes over the network (for system installation, or in the case of diskless designs, for booting and access to a root file system), IPMI access (IPMI-over-LAN) for remote hardware control and management, serial-over-LAN, and NFS access to “home” file systems for access to user home areas. All current motherboard candidates support two or more embedded gigabit or faster Ethernet ports.

In our experience, serial connections to each computer node are desirable. These connections can be used to monitor console logs, to allow login access when the Ethernet connection fails, and to allow access to BIOS screens during boot. Serial-over-LAN (standard with IPMI 2.0) will be used to provide these serial connections.

Network Plan

For LQCD-Ext II project clusters, we will replicate the network layout currently used on all of the FNAL and JLab lattice QCD clusters. In these designs all remote access to cluster nodes occurs via a “head node”, which connects to both the public network and to the private network that forms the sole connection to the computer nodes. Secure ID logon (Kerberos at FNAL, ssh at JLab) is required on the head node. “R-utility” (rsh, rlogin, rcp) or host-authenticated ssh are used to access the compute nodes.

File I/O

Particularly for analysis computing, large aggregate file I/O data rates (multiple streams to/from diverse nodes) are required. Data transfers over the high performance Infiniband network are preferred to transfers over Ethernet. Conventional TCP/IP over Infiniband relies on IPoIB (“IP over IB”, one of the protocols supported by the Open Fabrics Enterprise Distribution, or OFED, Infiniband software stack).

FNAL and JLab use *Lustre* as their primary data file systems, and NFS for user home areas and common software directories. *Lustre* provides a POSIX-compliant file system visible from all worker nodes and from the cluster head node. *Lustre* has the property that the storage volume and aggregate performance (instantaneous rate of data movement summed across all active transfers) can be scaled upwards by adding additional storage server nodes (known as *OSS* nodes, which serve *OST* disk volumes). Each new storage server node adds additional independent disk spindles and I/O processing capability to the file system. Current systems are cost optimized for capacity, and as needed future systems might be skewed towards higher bandwidth.

The LQCD-Ext II project will carefully watch developments in parallel file systems for changes that can impact the deployed systems. LQCD-Ext II will leverage work in this area performed by the large high energy physics experiments such as Atlas and CMS at the Large Hadron Collider. Relevant issues include concerns over the long term viability of *Lustre* given the recent series of ownership changes of the software (originally Sun Microsystems, then Oracle Corporation, then Whamcloud, and now Intel Corporation) and the emergence and/or maturation of competing parallel file systems such as GPFS and pNFS (the parallel version of NFSv4).

Procurement Strategy

LQCD-Ext II will procure as many as four separate lattice QCD computing systems, one in each of the final four years of the project. We consider a mixed conventional and GPU-accelerated cluster purchase to be a single procurement, as these would take place at a single host laboratory typically using a single purchase contract. If appropriate, the hardware budget from two years might be used to procure a single larger system. The guiding principal of all of these procurements is that the most cost effective hardware will be deployed, where effectiveness is judged by the quantity of science (and of course, quality of science in terms of the reliability of the numerical results) that will be produced during the lifetime of the individual lattice QCD system. In addition to commodity hardware and GPU-accelerated clusters, similar to those deployed during LQCD-ext, we will evaluate alternatives such as the IBM BlueGene family of computers, traditional supercomputers such as the Cray series, and other hardware suitable for lattice QCD calculations that may emerge.

At each of the annual project progress reviews, scheduled in or about the month of May of each fiscal year, LQCD-Ext II will present the plans for the deployment that will occur in the next fiscal year. For example, in spring of calendar year 2013, the LQCD-ext project presented the plans for the procurement that occurred in FY2014. The annual presentation of procurement plans will include the selection of hardware designs that will be considered, or the procedure that will be used to determine this selection, cost and performance estimates and their justifications, and a detailed schedule.

All procurements will utilize a multistep process:

1. Identify and characterize candidate computer and network hardware
2. Create (a) preliminary system design(s)
3. Solicit vendor feedback on the preliminary system design(s) through an RFI (Request for Information) solicitation
4. Update system design(s) based on vendor feedback and any new information that has emerged
5. Solicit vendor cost proposals for the system design(s) through an RFP (Request for Proposal) solicitation
6. Evaluate RFP responses and award purchase order(s) to the winning vendor(s), issuing a final system design as necessary
7. Accept or reject the delivered system(s) based on acceptance testing

Both the preliminary system designs and final system design may include one or two selections of hardware; for example, in a given year, traditional commodity clusters, advanced many-core clusters and GPU-accelerated clusters may be included. Throughout the five years of the project, LQCD-Ext II personnel will actively monitor the market, identifying and characterizing through benchmarking

candidate hardware for upcoming procurements. Project personnel will also interact closely with computer and network manufacturers to understand product features and schedule roadmaps.

The evaluation and selection of hardware for the preliminary system design, and the evaluation of vendor responses to the RFP, will rely on the projected performance of the anticipated lattice QCD applications that will be run on the hardware during its lifetime. The particular mixture of lattice QCD applications to be used will be determined by LQCD-Ext II Project staff in consultation with the USQCD Executive Committee and the USQCD Scientific Program Committee.

All awards will utilize firm, fixed-price contracts which include a minimum of 3 years of hardware warranty. Vendors will be encouraged to include modifications to the system designs in their RFI and RFP responses that would maximize the value of the delivered systems. Purchase awards will be based on best value evaluations that will include factors such as price/performance, quality of the vendor, quality of the proposed hardware, power consumption of the proposed hardware, impact on the facility infrastructure of the host laboratory, and usability of the delivered system.

LQCD-Ext II will procure storage for *Lustre* and NFS file systems separately from the computing systems. The amount of and performance of storage purchased will be determined in part from the requests that are required for all proposals to the Scientific Program Committee for allocations of time. The incremental storage added at each site annually will provide more than the sum of requested storage in the annual allocations proposals to accommodate a modest amount of estimation error.

Strategy for the FY2014 LQCD-Ext Deployment

As an example of the annual planning and selection strategy that will be used in LQCD-ext II, we can review the process employed for the FY2014 LQCD-ext deployment. For FY2014 the hardware candidates were an expansion of the IBM BlueGene/Q system at BNL, a conventional Infiniband cluster at Fermilab, an accelerated Infiniband cluster at Fermilab, or some mixture of the three.

The LQCD-Ext II strategy for determining the hardware for FY2014 took into account the availability of hardware, pricing, performance, and life-cycle costs. Because the project had to request the distribution of FY2014 funds among the three laboratories by mid-August 2013, a sequence of information gathering steps were used as listed in Table II below, culminating in the selection of Fermilab as the host laboratory for the FY2014 hardware, a combination of a conventional and an accelerated cluster.

Table II – FY2014 Acquisition Planning Process
(all dates 2013 unless otherwise noted)

Step	Description	Target Due Date
1	The LQCD-Ext Computing Project team (i.e., “the Project”) will provide the LQCD Executive Committee (EC) with data summarizing the distributions of job types and sizes during the prior year on the hardware operated by the Project (Infiniband and GPU-accelerated clusters). The Project will request that the EC provide the anticipated scientific program requirements for various architectures (i.e., leadership-class machines, BG/Q rack or Infiniband cluster, and GPU-accelerated cluster). Information on USQCD hardware usage will be presented to the collaboration at the 2013 All-Hands Meeting April 19-20.	Mar 26

2	The Project will prepare the FY14 Acquisition Strategy document for presentation and review at the FY2013 DOE Annual Progress Review. The Acquisition Strategy will outline the various options under consideration and the proposed process for selecting the mix of computing hardware that will be procured and deployed in FY13 using project funds.	May 9-10
3	The Project will request that the BNL site manager prepare a plan for procuring and operating a BG/Q half-rack, detailing estimating hardware, storage, deployment, and operations costs.	Jun 3
4	The EC, with input from the Scientific Program Committee (SPC), will provide the Project with the anticipated scientific program requirements for various architectures (i.e., leadership-class machines, BG/Q or Infiniband cluster, and GPU or MIC accelerated cluster). A helpful way of conveying this information would be for the EC to provide an estimate of the relative fractions of “analysis core-hours” and “cost-equivalent GPU-hours” needed to support the scientific program over the next 1 to 2 years. Ideally, the EC will provide the Project with anticipated needs on a per year basis for FY14 and FY15.	Jun 17
5	The BNL site manager will provide the Project with a preliminary plan for procuring and operating a BG/Q half-rack extension to the existing (FY13) BG/Q half-rack, including estimated costs and schedule.	Jul 1
6	The BNL site manager will provide the Project with a final plan for procuring and operating a BG/Q half-rack extension to the existing (FY13) BG/Q half-rack, including costs (hardware, storage, costed manpower for deployment and operations) and schedule.	Jul 22
7	The Project will review the technical landscape, conduct an alternatives analysis of the various options, and propose a cost-effective solution for the FY14 hardware deployment. When considering viable options, the Project will need to factor in the total cost of ownership (TCO) for each solution. In addition to hardware and deployment costs, TCO also includes on-going operations and support costs. Hardware costs will include any necessary storage acquisitions. For solutions involving Infiniband clusters and GPU-accelerated clusters, an operations cost model already exists. For a BG/Q option, the Project will need to understand the cost model for operating BG/Q hardware at BNL. Information on the cost of a BG/Q half-rack extension to the existing (FY13) BG/Q half-rack will also be needed. Results of the analysis and an overview of the proposed solution will be summarized in the Alternatives Analysis document. The Project will verify the host laboratory’s ability and willingness to provide the necessary space, power, and cooling for each alternative.	Jul 29
8	The EC will review the Alternatives Analysis document and proposed FY14 hardware solution, and will provide advice on how to proceed to the Project Manager.	Aug 12
9	The Project will analyze the advice of the Executive Committee as well as any new data that might have been obtained, and will produce the final plan for the FY14 hardware deployment. The Project Manager will advise the EC, the host laboratories, the Federal Project Director, and Project Monitor of the planned FY14 hardware acquisition.	Aug 15
10	The Project Manager will revise the project budget as necessary to accommodate the FY14 hardware solution. Depending on the alternative selected, changes may be required in the planned allocation of funds across the three host laboratories.	Aug 20
11	The Project Manager will provide the Federal Project Director with the FY14 Financial Plan, containing the requested distribution of project funds to the three host laboratories.	Aug 20 (est.)
12	The Project will develop a detailed acquisition plan, with timeline, based on the approved FY14 architecture solution.	Sep 30, 2013
13	The Project will execute the FY14 acquisition plan in a manner that meets approved performance goals and milestones.	Sep 30, 2014

