DOE FY2015 Annual Progress Review

Response to Technical Questions

Bill Boroski, Rob Kennedy, Amitoj Singh, and Chip Watson

for the

LQCD-ext II Project Team
and USQCD Collaboration

May 22, 2015

DOE FY15 Annual Progress Review
Brookhaven National Laboratory
May 21-22, 2015

1) LQCD-ext II Milestones

Level-1 and Level-2 project milestones are defined and recorded in the Project Execution Plan. A high-level definition of the milestones is provided in Section 6.3, with specific target levels and target completion dates defined in Appendices C and D.

The following elements are taken from the LQCD-ext II PEP, v1.1:

> 6.3 ***Project Milestones***
>
> Table 2 shows the Level 1 project milestones that are tracked by the DOE Federal Project Director and Project Monitor. These milestones are also defined and tracked in the project WBS. The target levels for new computing capacity deployed and aggregate computing delivered are defined in Appendix D - Computing Facility Performance Metrics.

1) LQCD-ext II Milestones (cont.)

These tables are taken from Section 6.3 of the PEP, v1.1.

Table 2. Level-1 Milestones

No.	Level 1 Milestone	Fiscal Year
1	Computer architecture planning for the FY16 procurement complete & reviewed	Q3 2015
2	Procurement and deployment of zero teraflops (sustained – *Conventional Resources*) in FY15 (no deployment in FY15 is planned, but this placeholder will account for any change in budget profile)	Q3 2015
3	Target level of aggregate *Conventional Resources* computing delivered in FY15	Q4 2015
4	Target level of aggregate *GPU-accelerated Resource* computing delivered in FY15	Q4 2015
5	Computer architecture planning for the FY17 procurement complete & reviewed	Q3 2016
6	Procurement and deployment of *Conventional Resources* in FY16	Q4 2016
7	Procurement and deployment of *Accelerated Resources* in FY16	Q4 2016
8	Target level of aggregate *Conventional Resource* computing delivered in FY16	Q4 2016
9	Target level of aggregate *GPU-accelerated Resource* computing delivered in FY16	Q4 2016
10	Computer architecture planning for the FY18 procurement complete & reviewed	Q3 2017
11	Procurement and deployment of *Conventional Resources* in FY17	Q3 2017
12	Procurement and deployment of *Accelerated Resources* in FY17	Q3 2017
13	Target level of aggregate *Conventional Resource* computing delivered in FY17	Q4 2017
14	Target level of aggregate *GPU-accelerated Resource* computing delivered in FY17	Q4 2017
15	Computer architecture planning for the FY19 procurement complete & reviewed	Q3 2018
16	Procurement and deployment of *Conventional Resources* in FY18	Q4 2018
17	Procurement and deployment of *Accelerated Resources* in FY18	Q4 2018
18	Target level of aggregate *Conventional Resource* computing delivered in FY18	Q4 2018
19	Target level of aggregate *GPU-accelerated Resource* computing delivered in FY18	Q4 2018
20	Procurement and deployment of *Conventional Resources* in FY19	Q3 2019
21	Procurement and deployment of *Accelerated Resources* in FY19	Q3 2019
22	Target level of aggregate *Conventional Resource* computing delivered in FY19	Q4 2019
23	Target level of aggregate *GPU-accelerated Resource* computing delivered in FY19	Q4 2019

Table 3. Level-2 Milestones

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

These Level-2 milestones are included in the PEP as an example of the milestones associated with annual procurement activities. Actual Level-2 milestones, with specific target dates, are formally defined and tracked in the project schedule, which is maintained in an MS Project file maintained by the Project Office.

1) LQCD-ext II Milestones (cont.)

Appendix C defines the planned costs and schedules for the key performance metrics that are tracked and reported monthly to the DOE Federal Project Director (OHEP) and Federal Project Monitor (ONP).

Appendix C. Cost and Schedule Performance Metrics

Description of Activity	DME, SS, MR	Total Cost		Current Baseline (07/10/2014)			
		Planned Cost ($M)	Actual Cost ($M)	Planned Start Date	Actual Start Date	Planned Completion Date	Actual Completion Date
FY15 SS - Aggregate sustained computing delivered to USQCD community. 88 TFlops-yrs (*Conventional Resources*) 92 Eff. TFlops-yrs (*Accelerated Resources*)	SS	$1.954		10/01/2014		09/30/2015	
FY16 DME Procurement and deployment of new sustained computing capacity. 10 Tflops (*Conventional Resources*) 39 Eff. Tflops (*Accelerated Resources*)	DME (FY16 DME + FY15 MR)	$1.116 ($1.070 + $0.046)		10/01/2015		08/30/2016	
FY16 SS - Aggregate sustained computing delivered to USQCD community. 68 TFlops-yrs (*Conventional Resources*) 67 TFlops-yrs (*Accelerated Resources*)	SS	$1.843		10/01/2015		09/30/2016	
FY17 DME Procurement and deployment of new sustained computing capacity. 14 Tflops (*Conventional Resources*) 52 Eff. Tflops (*Accelerated Resources*)	DME (FY17 DME + FY16 MR)	$1.313 ($1.226 + $0.087)		10/01/2016		06/30/2017	
FY17 SS - Aggregate sustained computing delivered to USQCD community. 70 TFlops-yrs (*Conventional Resources*) 95 Eff. TFlops-yrs (*Accelerated Resources*)	SS	$1.713		10/01/2016		09/30/2017	
FY18 DME Procurement and deployment of new sustained computing capacity. 28 Tflops (*Conventional Resources*) 106 Eff. Tflops (*Accelerated Resources*)	DME (FY18 DME + FY17 MR)	$1.459 ($1.398 + $0.061)		10/01/2017		08/30/2018	
FY18 SS - Aggregate sustained computing delivered to USQCD community. 85 TFlops-yrs (*Conventional Resources*) 145 Eff. TFlops-yrs (*Accelerated Resources*)	SS	$1.516		10/01/2017		09/30/2018	
FY19 DME Procurement and deployment of new sustained computing capacity. 35 Tflops (*Conventional Resources*) 135 Eff. Tflops (*Accelerated Resources*)	DME (FY19 DME + FY18 MR)	$1.723 ($1.637 + $0.086)		10/01/2018		06/30/2019	
FY19 SS - Aggregate sustained computing delivered to USQCD community. 80 TFlops-yrs (*Conventional Resources*) 290 Eff. TFlops-yrs (*Accelerated Resources*)	SS	$1.299		10/01/2018		09/30/2019	
FY19 Management Reserve	MR	$0.064		10/01/2018		09/30/2019	
Total		**$14.000**		**10/1/2015**		**09/30/2019**	

Legend
DME = Development/Modernization/Enhancement; SS = Steady-State Operations; MR = Management Reserve

Notes:
1) Following project policy, unspent management reserve from one year is rolled into the hardware procurement budget for the following year. The DME planned costs in this table are based on the assumption that management reserve will not be used and will thus be available to augment the hardware budget.
2) Planned steady-state (SS) costs include Operations & Maintenance; and Project Management.

1) LQCD-ext II Milestones (cont.)

Appendix D defines the key performance indicators that are tracked and reported annually to the DOE Federal Project Director (OHEP) and Federal Project Monitor (ONP). Results are also reported at the DOE Annual Progress Review.

Whereas Appendix C tracks cost and schedule performance, Appendix D tracks performance against our KPIs.

The table included on this slide is a portion of the actual table contained in the PEP and is presented as an example of the level of detail defined and tracked.

Appendix D. Key Performance Indicators (KPIs)

ID	Fiscal Year	Measurement Category	Measurement Indicator	Target	Actual Results	Rating
1	2015	Scientific Program Support	TF-Yrs delivered towards the completion of the Scientific Program – *Conventional Resources*	88 TF-Yrs	Available in Q1 FY16	
2	2015	Scientific Program Support	TF-Yrs delivered towards the completion of the Scientific Program – *Accelerated Resources*	92 TF-Yrs	Available in Q1 FY16	
3	2015	Responsiveness	% of tickets resolved within 2 business days	≥95%	Available in Q1 FY16	
4	2015	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 FY16	
5	2015	Reliability and Availability	% of average machine uptime across all LQCD computing sites	≥95%	Available in Q1 FY16	
6	2015	Quality of Service Delivery	Customer satisfaction rating (Customers rate satisfaction with the service provided on a scale of 1 to 5)	≥92%	Available in Q1 FY16	
7	2016	Effectiveness	Additional computing resources deployed by the project, expressed as an average of the HISQ and DWF algorithm performances in TFlops. – *Conventional Resources*	≥10 TF	Available in Q4 FY16	
8	2016	Effectiveness	Additional computing resources deployed by the project, expressed as an average of the HISQ and DWF algorithm performances in TFlops. – *Accelerated Resources*	≥39 TF	Available in Q4 FY16	
9	2016	Scientific Program Support	TF-Yrs delivered towards the completion of the Scientific Program – *Conventional Resources*	68 TF-Yrs	Available in Q4 FY16	
10	2016	Scientific Program Support	TF-Yrs delivered towards the completion of the Scientific Program – *Accelerated Resources*	67 TF-Yrs	Available in Q4 FY16	
11	2016	Responsiveness	% of tickets resolved within 2 business days	≥95%	Available in Q1 FY17	
12	2016	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 FY17	
13	2016	Reliability and Availability	% of average machine uptime across all LQCD computing sites	≥95%	Available in Q1 FY17	
14	2016	Quality of Service Delivery	Customer satisfaction rating (Customers rate satisfaction with the service provided on a scale of 1 to 5)	≥92%	Available in Q1 FY17	

5) Walk us through the Risk Register and how it is maintained.

The Risk Register is a workbook maintained by the ACPM according to the process defined in the Risk Management Plan. The Risk Register in PDF format is posted on the review web site. *The Risk Register spreadsheet has now been added to the review web site too.*

On the first of the month, the ACPM checks whether a risk's Next Review Date is coming up in the month. If any are, a list of the risks coming due is sent to the IPT with a target date for the actual review (usually held as part of the biweekly Site Managers Mtg). The IPT conducts preliminary discussion by email, with the formal review in a teleconference. Notes are maintained for each Risk Review and posted, documenting changes and intent. Risk Register changes are then drafted and sent to the IPT to assure the final wording and settings reflect consensus. The "Next Review Date" is then assigned based on the earlier of:

(a) yearly/quarterly/monthly review frequency based on risk priority, or

(b) Agreed-upon date more appropriate for risk review (e.g. after a funding decision is due).

(walk through the Risk Register, example of risk review notes, Risk Management Plan)

2) Do you have a Succession Plan, especially for your technical staff?

We at least partly address this in Risk ID 37:

- Description: Changes in staff can have adverse effects on the project.
- Mitigation: The project maintains staff depth in key roles: Project Manager, BNL Site Manager, FNAL Site Manager, and TJNAF Site Manager. For these roles, an active deputy exists who can fill the role if and when necessary. This should keep the impact of any one key staff member Low, assuming we lose only one key staff member within a period of 6 months.

More specifically, for our technical staff:

Site	Technical Staff Member	"Operations" Back-up	"Site Architect" Back-up
BNL	Frank Quarant	Nick D'Imperio	Bob Mawhinney
FNAL	Don Holmgren	Amitoj Singh	Amitoj Singh
TJNAF	Chip Watson	Sandy Philpott	Robert Edwards

3) Considering past and present architectures -- how were technical decisions made, including the process that led to the BG/Q? How was the process executed?

Architecture Decision Process:

1. Research

   The project actively tracks the computing market to understand trend lines of existing and similar technologies, as well as watching for new alternatives, such as when GPUs became viable for USQCD. This research activity includes understanding the behavior of current and near term LQCD algorithms, such as when multi-grid became available as a production algorithm.

   As part of this tracking of the market, we maintain good relationships with silicon/component companies (Intel, AMD, Mellanox, Nvidia) and a representative sampling of OEMs. LQCD is one of many reference applications that both Intel and Nvidia are using in the design of future chips.

   Tracking includes attending SuperComputing, the GPU Technology Conference (Nvidia), and the Intel HPC Forums.

3) Considering past and present architectures -- how were technical decisions made, including the process that led to the BG/Q? How was the process executed? (cont.)

2. Forward looking design

   Each year the project produces an alternatives analysis document that estimates the cost and benefits of multiple known options as well as a "do nothing" option for completeness. This contributes to the process of site selection and feeds into the RFI process in the fall. As part of building up a procurement RFP, details are refined to constrain the range of responses that vendors will submit, with details on network bandwidth, memory speed, etc. to help guide the vendors in their optimizations.

   In FY13, the alternatives analysis design led to the decision to procure a half rack of BG/Q at BNL. (See AA document in the "Other Project Documents" on the review website.)

3. Adaptive optimization

   Finally, the project frequently takes advantage of just-in-time adaptive design, such as when FNAL doubled the memory per node to help support G-2 deflation and other applications, and when a decision was taken to spend the deferred funds on expanding the conventional Infiniband cluster.

4) What is the plan for BNL after the planned retirement of the BG/Q half-rack in FY18?

The current baseline plan for LQCD-ext II calls for new hardware deployments at JLab in FY16-17 and FNAL in FY18-19. This plan was created during the DOE Critical Decision process to demonstrate what the project could deliver within the defined 5-year funding profile. We are baselined against compute capacity and delivery goals, as well as project costs and schedules.

We have a documented procurement strategy and a well-defined process for determining the optimum acquisition and deployment plan for every acquisition we undertake. This process is executed on an annual basis and has been exercised successfully many times.

BNL has recently expressed interest in hosting a project-funded Intel Knights Landing based machine in FY17. This interest will be factored into our annual acquisition and deployment decision process.

4) What is the plan for BNL after the planned retirement of the BG/Q half-rack in FY18. (cont.)

This consideration of all reasonable options is not unlike a similar situation that occurred in LQCD-ext. The initial baseline plan called for deployments at FNAL and JLab in alternating years and the retirement of QCDOC in FY10. However, we operated QCDOC through FY11 and in each year all options for new deployments were considered in our annual alternatives analysis process. In FY13, the best option involved the acquisition and deployment of the BG/Q half-rack at BNL, and a small cluster at FNAL.

Going forward, if it is determined through our alternatives analysis process that deploying a cluster at BNL is the most cost-effective solution in a given year, then we will work with BNL to develop a deployment and operations plan that ensures that we have adequate staffing and infrastructure support.

It may also be determined that operating the BG/Q beyond FY18 is cost-effective, in which case we would modify our plans to include staffing and maintenance support beyond FY18.

If it is determined that operating a cluster at BNL and/or extending the life of the BG/Q is not cost-effective, then we will consider other options for maintaining BNL's involvement in the project.

4) What is the plan for BNL after the planned retirement of the BG/Q half-rack in FY18. (cont.)

As a reference, the following excerpt is taken from the document "LQCD-ext II Acquisition Strategy, v 1.3", which is posted on the review website.

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.

6) User Survey – Why are some sites rating lower than others, in particular for Documentation? What is the plan to address this?

Individual issues are addressed in the 2014 User Survey Report if they fell below 80%:

FY14 Computing Facilities	All Sites	BNL	FNAL	JLab
Overall Satisfaction	97%	94%	100%	93%
Documentation	88%	70%	91%	89%
User Support	96%	85%	100%	95%
Responsiveness	96%	93%	100%	90%
Reliability	96%	97%	100%	86%
Ease of Access	91%	86%	96%	82%
Other Tools	97%	91%	100%	94%

BNL Documentation

Past Action Plan from 2013 User Survey:

- BNL Site Staff: Setup a web page with links to relevant BG/Q documentation to assist those new to the BG/Q technology in getting the information they need.
  - 2014: This was done.

Future Action Plan for 2015 User Survey:

- BNL Site Staff: The documentation web pages went down due to the retirement of old hardware. We are working with Bob M to host this documentation at Columbia instead.

Users said:

- Question 7 (Documentation): "The BNL documentation is out of date."

JLab Ease of Access

This area has had a "yellow" rating in recent years. However the issue we believe is due to the "2 hops" required to access JLab LQCD clusters, since the head nodes are not on the public network. This approach is required by JLab Cybersecurity Plan, but is still annoying to users.

We have no action plan to address this as long as it remains above 80%.

12) In the context of Infiniband clusters, GPU clusters, and the BG/Q, should the 8.5% spent on warranties be used to purchase additional hardware?

A typical warranty contract at Fermilab includes the following:

- Four-hour response during business hours is required.
- The vendor is our single point of contact for problems found in any component of the delivered system and accepts complete responsibility for its resolution.
- Vendor is fully responsible to bring back into service defective systems within five business days of the report of the problem.
- Fermilab reserves the right to require replacement of the entire server if it has experienced five hardware service calls within a six month time period.

A typical warranty contract at JLab includes the following:

- Maintain a pool of spare parts to replace failed components.
- The vendor is our single point of contact for cross-shipping replacement parts.
- Vendor is responsible for replenishing the spare parts pool during the warranty period.

At BNL it is not possible to procure replacement parts for the USQCD BG/Q half-rack without a service contract with IBM; therefore, foregoing the maintenance contract is not a viable option.

12) In the context of Infiniband clusters, GPU clusters, and the BG/Q, should the 8.5% spent on warranties be used to purchase additional hardware? (cont.)

All of our hardware acquisitions include a 3-year manufacturers warranty on all hardware components, for which the precise cost is unknown as the warranty cost is included in the line item costs provided in the vendor quotes. For the most recent FY14/15 hardware acquisition we upgraded to a 4th and 5th year warranty at a cost of $75,600, a 3% cost of the total acquisition of $2.25M.

Without any warranty, a significant amount of time and effort would be spent in diagnosing problems and ordering replacement parts from various hardware manufacturers. Also, there is the increased risk in a vendor delivering a low quality product.

From our operational experience since FY05, our portfolio of machines have useful lifetimes beyond the 3-year warranty period making the exercising of the option to purchase an extended 4^{th} and 5^{th} year warranty a reasonable choice.