

Report from the Project Manager

Bill Boroski
LQCD-ext II Contractor Project Manager
boroski@fnal.gov

Rob Kennedy
LQCD-ext II Assoc. Contractor Project Manager
kennedy@fnal.gov

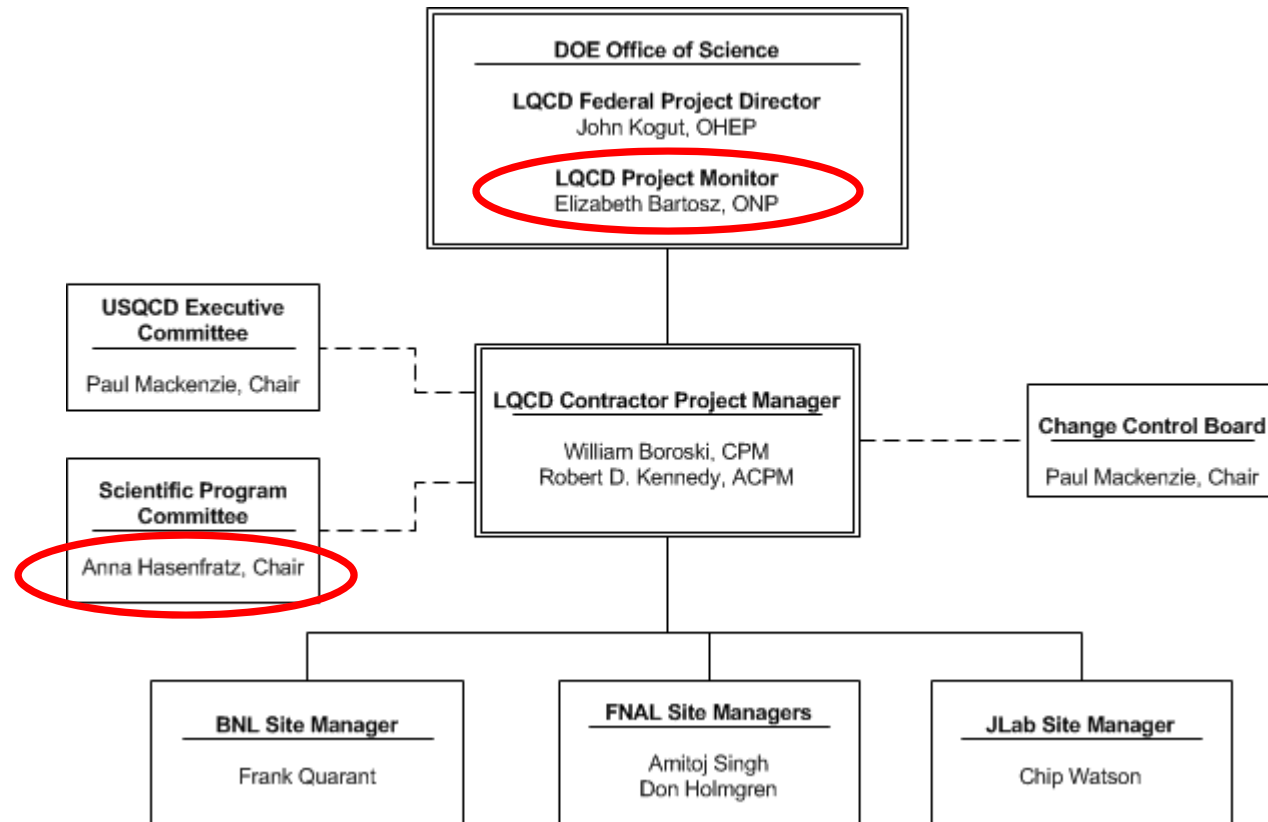
USQCD All-Hands Meeting
Fermi National Accelerator Laboratory
May 1–2, 2014

Outline

A year in transition

- ▶ Organizational changes
- ▶ LQCD-ext wrap-up
- ▶ LQCD-ext II start-up
- ▶ User survey results
- ▶ Summary


Management & Oversight



Organizational changes:

- Elizabeth Bartosz replaced Kawtar Hafidi as NP Project Monitor
- Anna Hasenfratz replaced Robert Edwards as SPC Chair


FY14 LQCD-ext Project Performance



FY14 data for conventional resources are shown.
The uptime goal is 8000 hours per year (91.3%).

Performance goal is based on an average of the sustained performance of domain wall fermion (DWF) and highly improved staggered quark (HISQ) algorithms

The inflection points in the Uptime Goal curve correspond to the retirement of the FNAL J/Psi cluster in January (which actually occurred in mid-May) and the deployment of the new cluster in July (delayed until September).



FY14 data for GPU-accelerated clusters is shown.
The uptime goal is 8000 hours per year (91.3%).

Conversion from GPU-hrs to effective TF-ys is 140 GF/GPU, based on allocation-weighted performance of GPU projects running from July 1, 2012 through Dec 2012. Resources included are the FNAL Dsg and JLab 9g, 10g, 11g, and 12k clusters.

The inflection point in the Uptime Goal curve is the FY14 GPU cluster (which was delayed until September).

LQCD-ext Project Wrap-up

- ▶ The LQCD-ext project officially concluded on Sep 30, 2014.
- ▶ Computing Resources Deployed and Delivered

	# Machines	# Nodes	# GPUs	Delivered Computing Capacity of New Deployments (Tflop/s)	Delivered Performance (Tflop/s-yr)
Conventional Resources	6	1744	---	85.9	270
Accelerated Resources	3	150	448	90.2	258

These values reflect only the resources purchased using LQCD-ext project funds. Additional resources purchased with LQCD-ARRA funds during this period were also available for collaboration use.

- ▶ Cost Performance
 - We completed the LQCD-ext project with a modest level of unspent funds. Because the follow-on project is an extension of the current project, we were able to carry these funds forward into FY15. Part of these funds will be held and used to cover one month of FY16 operations should a federal budget Continuing Resolution occur again. The remaining funds were used to purchase the FY15 Pi0 expansion.

LQCD-ext II

(FY15 – FY19)


LQCD-ext II Approval

- ▶ LQCD-ext II approved Oct 1, 2014.
 - An extension of the previous project – not a new start
 - Approved project duration: Oct 1, 2014 through Sep 30, 2019
 - Approved total funding level: \$14 million

- ▶ The project was formally reviewed and approved following the formal DOE Order 413.3B Critical Decision (CD) approval process.
 - **CD-0: Approval of Mission Need**
 - Granted Sep 9, 2013
 - **Science Review**
 - Review held Nov 18, 2013; report issued Jan 29, 2014; recommended proceeding to CD-1
 - **CD-1: Approve Alternative Selection and Cost Range**
 - Review held Feb 25, 2014; CD-1 approval granted May 1, 2014
 - **CD-2: Approve Performance Baseline**
 - **CD-3: Approve Start of Execution**
 - Combined CD-2/3 review held Jul 10, 2014; approval granted Oct 1, 2014

Project Budget Comparison

Approved Budget Baseline: \$14 million



- ▶ The approved \$14 million budget represents a significant reduction in funding from prior levels, which had been back-loaded in the funding profile for the previous project (LQCD-ext).
- ▶ Personnel cost requirements are based on a refined staffing model. Level of operations support is based on number of nodes and GPUs in production during each year.
- ▶ Reduced funding level directly affects the amount of compute capacity we will be able to deliver to the science program.

LQCD-ext II Project Performance Goals

	Target Goals				
	FY15	FY16	FY17	FY18	FY19
Planned computing capacity of new deployments (Tflop/s)	0	49	66	134	172
Planned delivered performance (Tflop/s-yr)	180	135*	165	230	370

(DWF + HISQ averages used). Integrated performance figures use an 8000-hr year.

* The dip in performance is due to the retirement of aging clusters.

Hardware acquisition strategy will be similar to prior years

- FY15 budget does not provide sufficient funding for new hardware purchase. Fortunately, we were able to use carry-over funds from LQCD-ext to expand Pi0.
- For FY16–19, collaboration needs will be assessed annually and hardware procurement decisions will be made to make the best use of available funds.
- Preliminary plan calls for combined procurements at JLab in FY16–17 and FNAL at FY18–19. However, we may do two-year combined buys or individual buys, whichever best matches the collaboration needs and the hardware in the marketplace.

User Survey Results


Rob Kennedy

LQCD-ext II Associate Contractor Project Manager



FY14 Survey Results

- ▶ The FY14 User Survey measured user satisfaction during the 7 month period from March 2014 through September 2014 inclusive, dovetailing the previous User Survey that covered early part of FY14.
- ▶ The survey consisted of 29 questions designed to measure satisfaction with the compute facilities and the resource allocation process.
- ▶ The survey was distributed to 177 individuals
 - Responses were received from 61 individuals
 - By comparison, 66 individuals responded to the FY13 survey
 - 24 of top 48 most Active Users responded: 50% response rate
 - 20 of 27 PI's responded: 74% response rate
- ▶ FY14 overall satisfaction rating with Compute Facilities = 97%
 - Exceeds our KPI goal of 92%
- ▶ FY14 overall satisfaction with Resource Allocation Process = 84%
 - Similar to recent past years, except for spike in FY14 (97%)

Compute Facility Satisfaction Trends







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%



- ▶ The overall satisfaction rating has been trending upward over the previous four years.
- ▶ FY14 rating of 97% exceeds our target goal of 92%, and is similar to recent past.
- ▶ JLab's Overall Satisfaction rating of 93% in FY14 continues high level regained in FY13.
- ▶ BNL's rating for User Documentation was still below par, but a little higher than FY13.
- ▶ Ease of Access and User Documentation ratings were about the same as recent past.


Compute Facility Satisfaction Trends




- ▶ Responsiveness of Site Staff and User Support maintain high satisfaction ratings.
- ▶ System Reliability and Online Tools also continue to maintain high satisfaction ratings.

Allocation Process Satisfaction Trends


Clarity of the Call for Proposals




Transparency of the Allocation Process



Apparent Fairness of the Allocation Process




Belief that Allocation Process Helps Maximize Scientific Output



- ▶ Satisfaction rating trends for Allocation Process survey areas fell back to levels seen before FY13, a little lower perhaps in some cases.

Allocation Process Satisfaction Trends



- ▶ The overall satisfaction rating for the Allocation Process was 84% in FY14.
- ▶ We speculate that the largest single increase in resources that occurred in 2013 was mostly responsible for the high satisfaction rate that year, although improvements in the communications with users and the establishment of the Scientific Advisory Board also occurred in 2013.
- ▶ Some user comments suggest also that the wide range of systems handled in the allocations process can be confusing or seem to be handled inconsistently.

FY14 User Survey Summary

► Satisfaction with Compute Facilities

- The overall satisfaction rating of 97% exceeds our target goal of 92%.
- Very good satisfaction ratings overall for all individual sites.
- User documentation has improved, but there remains an opportunity for improvement at one of our sites.

► Satisfaction with Allocation Process

- The overall satisfaction rating of 84% is about the same as pre-FY13.
- We speculate that the largest single increase in resources that occurred in 2013 was mostly responsible for the high satisfaction rate that year, although improvements in the communications with users and the establishment of the Scientific Advisory Board also occurred in 2013.
- Some user comments suggest also that the wide range of systems handled in the allocations process can be confusing or seem to be handled inconsistently.

Summary

- ▶ We successfully met or exceeded all but two of key performance goals in FY14. We did not meet our target deployment dates, which caused us to fall short in terms of Accelerated Cluster computing delivered.
 - We missed deployment milestones due to Continuing Resolution and other factors.
- ▶ Although we exceeded our overall user satisfaction goal, user survey results indicate areas for potential improvement. Thank you for your honest feedback.
- ▶ The LQCD-ext II project was formally approved for 5 more years of acquisitions and operations.
 - A smaller budget means that we'll need to continue working hard to minimize operations costs, in order to channel more funds to hardware purchases.
- ▶ We continue to have a strong project team and we remain dedicated to meeting the needs of the USQCD collaboration.