

LQCD-ext Project Closeout

Bill Boroski
LQCD Contractor Project Manager

DOE FY2015 Annual Progress Review
Brookhaven National Laboratory
May 21-22, 2015

Outline

- ▶ Project scope, organization, and budget
- ▶ Performance measures and metrics
- ▶ FY14 performance and financial results
- ▶ Total project performance and financial results
- ▶ Summary

LQCD-ext Project Scope

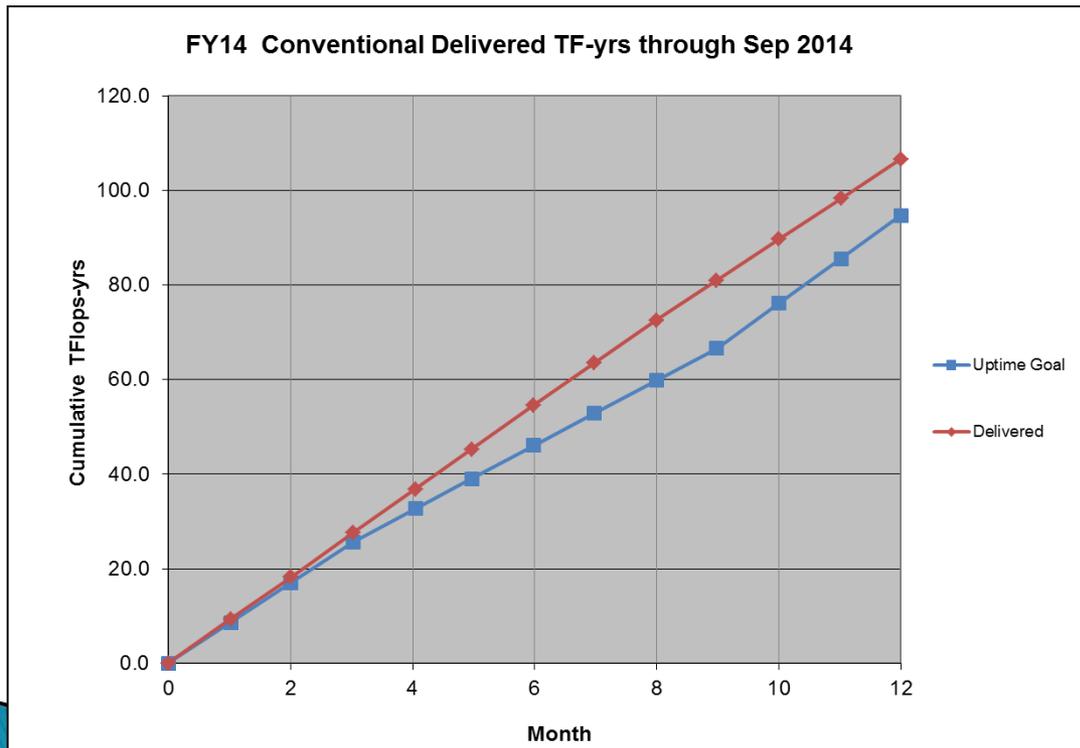
- ▶ Acquire and operate dedicated hardware at BNL, JLab, and FNAL for the study of QCD during the period FY2010–2014.
- ▶ Scope includes acquisition, deployment, and operation of computing facilities; software development is out of scope.
- ▶ At project completion, we were executing against the original baseline plan, with a few exceptions
 - Several machines were operated beyond planned 4-year lifetimes
 - QCDOC at BNL was operated through August 2011 (5.25 years)
 - 7n at JLab was operated through mid-May 2012 (5 years)
 - Kaon at FNAL was operated through June 2013 (7 years)
 - J-Psi at FNAL was operated through mid-May 2014 (5 years)
 - FY11 and FY12 procurements included a mix of conventional Infiniband cluster nodes and GPU-accelerated nodes.
 - A modest level of salary and M&S support was provided for the operation of prototype BG/Q at BNL, in exchange for 20 TF (peak) compute capacity (10% of one rack).
 - Beginning in FY13, operations support was provided for the compute hardware at JLab that was acquired under the LQCD-ARRA project.

Performance Measures & Metrics

- ▶ Performance goals and milestones for LQCD-ext are documented in the Project Execution Plan (Appendices C & D).
 - ▶ Ensures that the performance goals and milestones remain under formal change control and are readily available to the project team and stakeholders.
 - ▶ These are similar to the goals and milestones that had previously been explicitly defined in the baseline OMB Exhibit 300 document.
- ▶ 14 Level-1 project milestones
 - ▶ External reviews of future procurement plans
 - ▶ Incremental procurements/TFlops-deployed
 - ▶ Aggregate TFlops-yrs delivered
- ▶ 36 cost and schedule performance metrics
 - Planned costs and schedule completion dates
- ▶ 58 performance indicators
 - ▶ Additional computing resources brought on-line
 - ▶ System performance (i.e., % of time system available for work)
 - ▶ Process improvements (i.e., % of tickets resolved within 2 business days)
 - ▶ Customer satisfaction (measured through user surveys)
- ▶ Progress against these goals is tracked and reported periodically to the LQCD-ext Federal Project Director and Project Monitor.

FY14 Milestone Performance – (TFlops–yrs delivered) Conventional Hardware

- ▶ Computing delivered in FY14 from conventional compute hardware is shown.
- ▶ Uptime goal is 8000 hours per year (91.3%); average uptime for the metafacility = 97.5%
- ▶ The unmodified goal for FY14 was 90 TFlops–yrs.
- ▶ The project achieved 106.7 TFlops–yrs (119% of goal).



Conventional computing resources include:

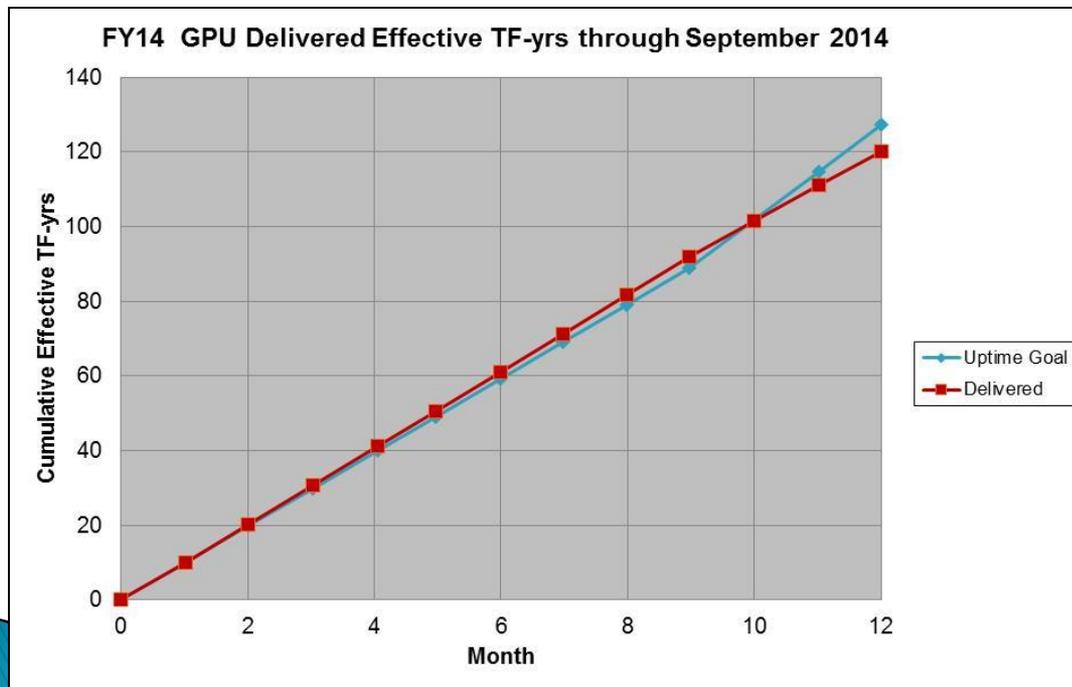
- FNAL Infiniband clusters
- JLab Infiniband clusters
- BNL BG/Q half-rack
- BNL BG/Q DD2 prototype rack

(8 machines total)

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 due to CR funding delays).

FY14 Milestone Performance – (TFlops–yrs delivered) Accelerated Hardware

- ▶ Computing delivered in FY14 from accelerated compute hardware is shown.
- ▶ Uptime goal is 8000 hours per year (91.3%); average uptime for the metafacility = 95.3%
- ▶ Conversion from GPU–hrs to effective TF–yrs is 140 GF/GPU, based on allocation–weighted performance of GPU projects running from July 2012 to December 2012.
- ▶ The goal for FY14 was 127.2 effective TF–yrs
- ▶ The project achieved 120.2 effective TF–yrs (95% of goal).



Computing resources include accelerated clusters operating at FNAL and JLab – FNAL Dsg and Pi0g and JLab 9g, 10g, 11g, and 12k clusters.
(6 machines total)

The inflection point in the Uptime Goal curve indicates the planned date for the FY14 GPU cluster deployment (which was delayed until September).

Funding delays due to the CR negatively impacted our deployment schedule, which in turn impacted our performance delivery.

FY14 Performance Summary – Milestones

Comparison of Actuals to Approved Baseline

Milestone #	Description	Actual Results	% of Plan	Planned Cost (\$K)	Actual Cost (\$K)	Planned Completion	Actual Completion
30	Architecture planning for FY14 procurement reviewed by external DOE committee	Plan reviewed & accepted	100%	59	0	06/30/13	05/15/13
34	Procurement & deployment of 57 TF (sustained) system	46.7 TF ⁽¹⁾ <i>IB Cluster</i> 13.1 TF <i>GPU Cluster</i> 33.6 TF (equivalent)	82%	2,688	2,044 ⁽²⁾	06/30/14	10/10/14 (IB cluster) 10/20/14 (GPU cluster)
35	90 TF–yrs aggregate computing delivered	106.7 TF–yrs ⁽³⁾	119%	1,512	1,826 ⁽⁴⁾	09/30/14	09/30/14
36	Security controls testing and contingency plan review complete at BNL, FNAL, & JLab	Completed as planned	100%	0	0	08/31/14	08/31/14

(1) The baseline capacity deployment target was missed for several reasons, which are described on the following slide.

(2) Includes costs for the conventional and accelerated cluster deployments at FNAL. Costs are below plan because budget revisions post-baseline had shifted funds from procurements to operations to meet operational and storage needs. In addition, the planned cost for procurement/deployment includes management reserve funds, which were not used for these procurements.

(3) Compute capacity delivered by conventional Infiniband clusters.

(4) Includes salary costs for operations, storage hardware, and other misc. operating expenses (travel, spares, repairs, tape, etc.)

FY14 Performance Summary – Milestones (2)

Comparison of Actuals to Approved Baseline (*with FY15 Pi0 expansion*)

- The baseline deployment target was missed for several reasons:
 - (1) The FY14 hardware budget was reduced from the \$2.29M baseline to \$1.8M to defer funds to cover budget uncertainty;
 - (2) memory on conventional nodes was expanded to 128GB;
 - (3) warranties were extended to 5 yrs;
 - (4) slowdown in Moore’s law for conventional nodes; and
 - (5) over-optimistic projection from Dsg for GPU performance/price.
- With the Pi0 expansion in FY15, we added 6 TF at a cost of \$451K, which brings us closer to the baseline performance target and cost.
- The effective TF rating for the GPU does not capture the effects of much larger memory on the K40s (12 GB vs. 6 GB on 12K’s K20s and 3 GB on Dsg’s M2050s. This allows larger problems to fit on a single node, and large multi-node problems required fewer nodes than on 12K which leads to better scaling.

Milestone #	Description	Actual Results	% of Plan	Planned Cost (\$K)	Actual Cost (\$K)	Planned Completion	Actual Completion
34	Procurement & deployment of 57 TF (sustained) system	46.7 TF ⁽¹⁾ <u>+ 6.0 TF</u> 52.7 TF	82% 92%	2,688	2,044 ⁽²⁾ <u>+ 451</u> 2,495	06/30/14	10/10/14 <i>(IB cluster)</i>
		<i>IB Cluster</i> 13.1 TF					10/20/14 <i>(GPU cluster)</i>
		<i>GPU Cluster</i> 33.6 TF <i>(equivalent)</i>					
		<i>IB Cluster Expansion</i> 6.0 TF					04/15/15 <i>(IB cluster expansion)</i>

**Blue text identifies items associated with the FY15 expansion procurement.*

FY14 Performance Summary – KPIs

Performance against other Key Performance Indicators (KPIs)

Measurement Indicators	Target	Actual Results
Customer Satisfaction Rating	≥92%	97%
% of tickets closed within two business days	≥95%	93% (277/298)
% of average machine uptime at the metafacility	≥95%	Conventional: 98.1% (weighted ave.) Accelerated: 95.4% (weighted ave.)
Weekly vulnerability scans	Scans performed at least weekly at each host institution	Mixture of daily and weekly scans performed at all sites.

- ▶ All KPI metrics were met or exceeded in FY14, with the exception of helpdesk ticket response time.
- ▶ Across Fermilab and JLab, we resolved a total of 277 out of 298 tickets within 2 business days. The remaining tickets were more complicated and required more time to diagnose and resolve.

FY14 Performance Summary

- ▶ We exceeded our compute delivery performance goal for our conventional Infiniband clusters.
- ▶ We missed our new cluster target deployment date of Jun 30 due to funding delays that delayed the start of procurement activities.
- ▶ Our compute capacity deployment goal was not met due to several factors:
 1. The FY14 hardware budget was reduced from the \$2.29M baseline to \$1.8M to defer funds to cover budget uncertainty;
 2. memory on conventional nodes was expanded to 128GB;
 3. warranties were extended to 5 yrs;
 4. slowdown in Moore's law for conventional nodes; and
 5. over-optimistic projection from Dsg for GPU performance/price.
- ▶ All KPI metrics were met with the exception of helpdesk ticket response time – although we came up two percent short on our goal, we still managed to resolve 277/298 tickets within two business days.

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: Summary of Tflop/s Deployed

Conventional Resources

	Site	Name	# Nodes	Capacity (TF)
FY10/FY11	FNAL	Ds	412	21.0
FY12	JLab	12s	280	12.8
FY13	BNL	BG/Q / DD2	614	26.3
FY13	FNAL	Bc	224	12.7
FY14	FNAL	Pi0	214	13.1
Total			1744	85.9

Total Capacity Deployed (TF)

Conventional	85.9
Accelerated	90.2
Total Deployed	176.1
Baseline Deployment Goal	145.0
% of Goal Achieved	121%

Accelerated Resources

	Site	Name	# Nodes	# GPUs	Capacity (TF)
FY12	FNAL	Dsg	76	152	21.3
	JLab	12k	42	168	35.3
FY13	None	---	0	0	0
FY14	FNAL	Pi0g	32	128	33.6
Total			150	448	90.2

LQCD-ext: Summary of System Uptimes

Year	Uptime Goal (%)	Actual (conventional systems) (%)	Actual (accelerated systems) (%)	Overall Average (%)	% of Goal
FY10	91.3	97.4	---	97.4	107
FY11	91.3	96.6	---	96.6	106
FY12	91.3	95.9	95.4	95.7	105
FY13	91.3	97.6	95.4	96.3	105
FY14	91.3	98.1	95.4	96.4	106

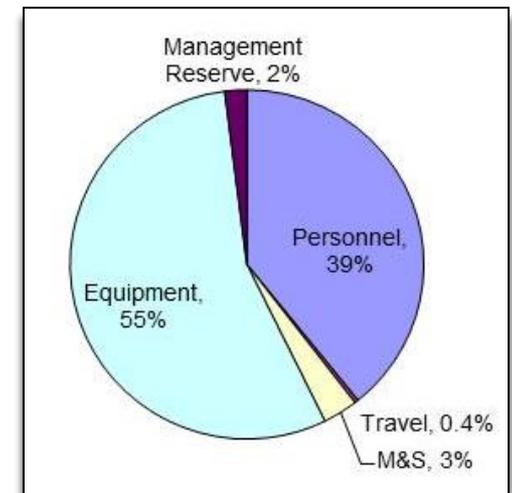
LQCD-ext Project Budget

- ▶ Approved Baseline Budget = \$18.15 million
 - Jointly funded by DOE Offices of High Energy and Nuclear Physics

Approved Funding Profile (in \$K)

Expenditure Type	FY10	FY11	FY12	FY13	FY14	Total
Personnel	1,139	1,306	1,456	1,340	1,644	6,885
Travel	13	11	12	12	12	60
M&S	104	84	84	84	84	440
Equipment	1,684	1,779	1,974	2,589	2,379	10,405
Management Reserve	60	69	75	75	81	360
Total	3,000	3,250	3,600	4,100	4,200	18,150

Planned Budget Distribution



LQCD-Ext Project Budget

- ▶ Changes in the budget, relative to the baseline.
 - Personnel Budget Changes
 - Updated salary cost basis for FY13-14, to account for changes in salaries and overheads
 - Modified staffing model in FY13, and again in FY14, based on operating experience
 - Increased staffing support to operate BG/Q and ARRA facilities in FY13-14
 - Increased staffing support at JLab in FY14 to improve support for accelerated clusters.
 - Reduced level of staffing support for project management in FY14.
 - M&S Budget Changes
 - Added funds to cover IBM maintenance contract for BNL BG/Q
 - Storage Hardware Budget Changes
 - Increased storage allocation to accommodate growing storage needs (5% -> 8%)
 - Compute Hardware Budget Changes
 - Reduced HW budget to accommodate staffing support for BG/Q and ARRA in FY13-14
 - Reduced HW budget to accommodate increased storage needs

LQCD-ext Cost Performance

- ▶ Comparison of final costs to baseline budget (\$K)

	Operating Funds	Equipment Funds	Total
LQCD Carryover	222	18	240
5-yr Baseline Budget	8,985	9,165	18,150
Fund Type Conversion	188	(188)	0
Total Actual Costs	<u>(8,611)</u>	<u>(8,974)</u>	<u>(17,585)</u>
Funds Remaining	784	21	805
% Spent	92%	99.8%	96%

- Throughout the LQCD-ext project, we managed the budget to maximize scientific output, and in accordance with the procedures and processes defined in the approved Project Execution Plan.
- Regarding the *Funds Remaining*:
 - Decision was made to defer a portion of FY14 funds to FY15 to cover one month of salary support at each host site (contingency against another CR).
 - Decision was made to hold back a portion of FY14 equipment funds due to uncertainty in LQCD-ext II CD-2/3 approval (to cover operations if LQCD-ext II approval was delayed or not granted).
- Once LQCD-ext II was granted CD-2/3 approval (Oct 1, 2014), a spending plan for the remaining funds was developed and presented to our DOE Federal Project Director and Project Monitor for review and concurrence. We are executing that plan through the LQCD-ext II project.

LQCD-ext Project Summary (FY10–14)

- ▶ The LQCD-ext Computing Project officially concluded on September 30, 2014.
- ▶ Successfully deployed and operated computing facilities at BNL, FNAL and JLab over the period FY10–FY14 (*Oct 1, 2009 through Sep 30, 2014*)
- Average uptime across the metafacility over the 5-year project: 97%
- Met or exceeded nearly all of our performance metrics and KPIs. Biggest impediments to success were out of our control (CRs, funding delays, slow down of Moore's Law, unexpected deviations in price/performance trends, etc.). We remained agile and worked hard to minimize the impact of these impediments.
- Final Project Cost
 - Project Budget: \$18.15M
 - *\$9.17M for equipment*
 - *\$8.99M for personnel, materials & supplies (e.g. storage hardware)*
 - Final Cost: \$17.59M (*96% of budget*)
 - *\$8.97M for equipment*
 - *\$8.61M for personnel, materials & supplies (e.g. storage hardware)*
 - ~\$805K in deferred and unspent funds have been carried forward to the Extension Project (LQCD-ext II) and are being expended in accordance with an approved spending plan.

Questions?