

LQCD-ext Project Management & Performance

Bill Boroski
LQCD-ext Contractor Project Manager

Robert D. Kennedy
Associate Contractor Project Manager

DOE FY2014 Annual Progress Review
Fermi National Accelerator Laboratory
May 15-16, 2014

Outline

- ▶ Updates to project scope, budget, and organization
- ▶ FY13 & FY14 performance results
- ▶ FY13 & FY14 financial results
- ▶ User survey 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.
- ▶ Currently executing against baseline plan, with a few exceptions
 - Several machines have been operated beyond planned lifetimes
 - QCDOC at BNL was operated through August 2011
 - 7n at JLab was operated through mid-May 2012
 - Kaon at FNAL was operated through June 2013
 - J-Psi at FNAL will be retired on May 19, 2014
 - FY11 and FY12 procurements included a mix of conventional Infiniband cluster nodes and GPU-accelerated nodes.
 - Providing a modest level of salary and M&S support for the operation of prototype BG/Q at BNL, in exchange for 20 TF (peak) compute capacity (10% of one rack).
 - Providing operations support for the compute hardware at JLab that was acquired under the LQCD-ARRA project.

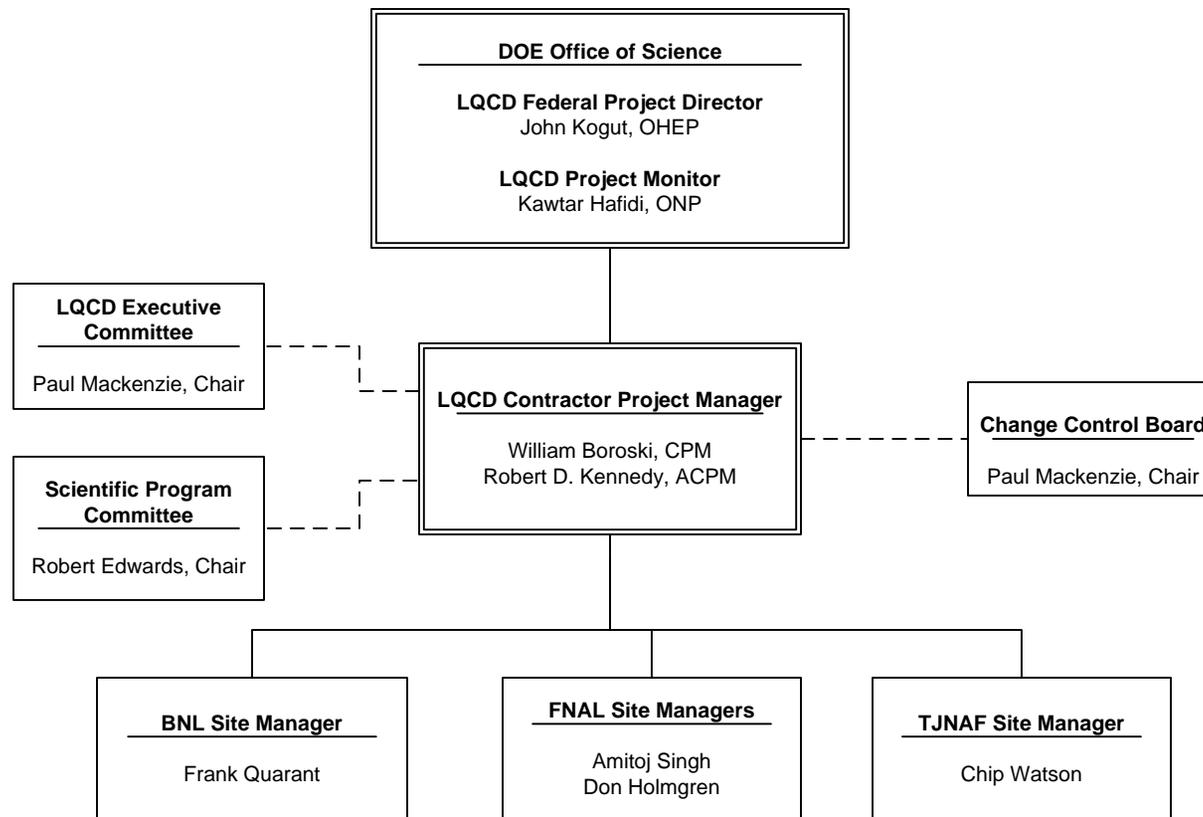
Project Execution & Work Planning

- ▶ **Project Execution Plan (PEP)**
 - Controlled document defining project need, requirements, scope, management, cost and schedule, change control, etc.

- ▶ **Work organized via WBS**
 - MS Project used to identify tasks, develop schedules, and track progress against milestones
 - Work broken down into two primary areas:
 - Steady-state operations and maintenance
 - Procurement and deployment of equipment and new systems

- ▶ **Other important project documents**
 - Risk Management Plan, Risk Register, Acquisition Strategy Documents, Annual Acquisition Plans, Quality Assurance Plan, C&A Documentation
 - All under formal version control

Management & Oversight



- ▶ No changes in organizational structure or composition since the last review.

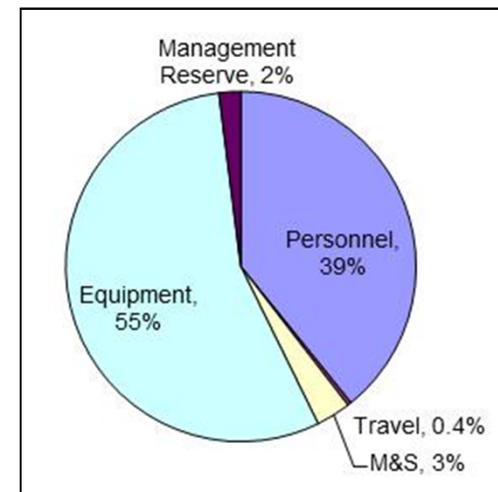
LQCD-Ext Project Budget

- ▶ 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

Budget Distribution



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



Hardware Budget Breakdown (in \$K)

Fiscal Year	Compute Hardware	Storage Hardware	Total
FY10	1,600	84	1,684
FY11	1,690	89	1,779
FY12	1,875	99	1,974
FY13	2,460	129	2,589
FY14	2,260	119	2,379
Total	9,885	520	10,405

Equipment budget is used to procure compute and storage hardware

Storage budget changes:

- Baseline storage budget was originally set at ~5% of total hardware budget.
- Currently set at 8% to meet growing storage needs.

LQCD-Ext Project Budget

- ▶ We are currently 92% through the LQCD-Ext project.
- ▶ Changes in the budget forecast, 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 Project Budget

▶ Comparison of current forecast to baseline budget (\$K)

Expenditure Type	Baseline Budget	Current Forecast	Change Relative to Baseline	% Change
Personnel	6,885	7,068	183	3%
Travel	60	74	14	23%
M&S (spares, tape, etc.)	440	681	241	55%
Compute Hardware	9,885	9,246	(639)	(6%)
Storage Hardware	520	763	243	47%
Management Reserve	360	318	(42)	(12%)
Total	18,150	18,150	---	(0%)

- “Change Relative to Baseline” shows the net effect of the changes previously described.
- We are managing the budget to maximize scientific output, and in accordance with the procedures and processes defined in the approved Project Execution Plan.

Performance Measures & Metrics

LQCD Hardware Performance Data

- ▶ Performance and utilization data are available online for LQCD-ext resources at all three sites (BNL, JLab, and FNAL)
 - BNL: <http://lqcd.bnl.gov/comp/usage/> *(was active for QCDOC thru end-of-life: Sep 2011)*
 - JLab: <http://lqcd.jlab.org/lqcd/>
 - FNAL: <http://www.usqcd.org/fnal>
- ▶ Available data include:
 - Machine usage on an hourly, daily, weekly, monthly, annual basis
 - Interactive views that allow users to select performance periods
 - System and node health monitoring
 - Node uptime, system temperature, processor temperature and fan speeds, CPU load average.
 - Job data
 - Project allocation usage, jobs running and in queue, nodes allocated to projects.
- ▶ Performance and utilization data for BG/Q is measured and analyzed monthly by the BNL site team and is available upon request

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 the same goals and milestones that had previously been explicitly defined in the baseline OMB Exhibit 300 document.

- ▶ 23 project milestones (for LQCD-ext)
 - ▶ External reviews of future procurement plans
 - ▶ Incremental procurements/TFlops-deployed
 - ▶ Aggregate TFlops-yrs delivered

- ▶ 29 performance indicators
 - ▶ Additional computing resources brought on-line
 - ▶ System performance (i.e., % of time system available for work)
 - ▶ Process improvements (i.e., % of tickets closed 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.

FY13 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%	57	0	06/30/13	05/09/13
31	Procurement & deployment of 44 TF (sustained) system	34.6 TF ⁽¹⁾ <i><u>IB Cluster</u></i> <i>12.7 TF</i> <i>and</i> <i><u>BG/Q half-rack</u></i> <i>21.9 TF</i> <i>(equivalent)</i>	79%	2,825	2,292 ⁽²⁾	06/30/13	07/10/13 <i>(IB cluster)</i> and 07/01/13 <i>(BG/Q half-rack)</i>
32	52.0 TF–yrs aggregate computing delivered	74.1 TF–yrs ⁽³⁾ (111% of goal)	111%	1,216	1,689 ⁽⁴⁾	09/30/13	09/30/13
33	Security controls testing and contingency plan review complete at BNL, FNAL, & JLab	Completed as planned	100%	0	0	08/31/13	08/31/13

(1) Total includes contributions from conventional Infiniband cluster and BG/Q half-rack. BG/Q rated at average of DWF and AQSTAD (USQCD rating). If running predominately DWF, rating is closer to 30.5 TF.

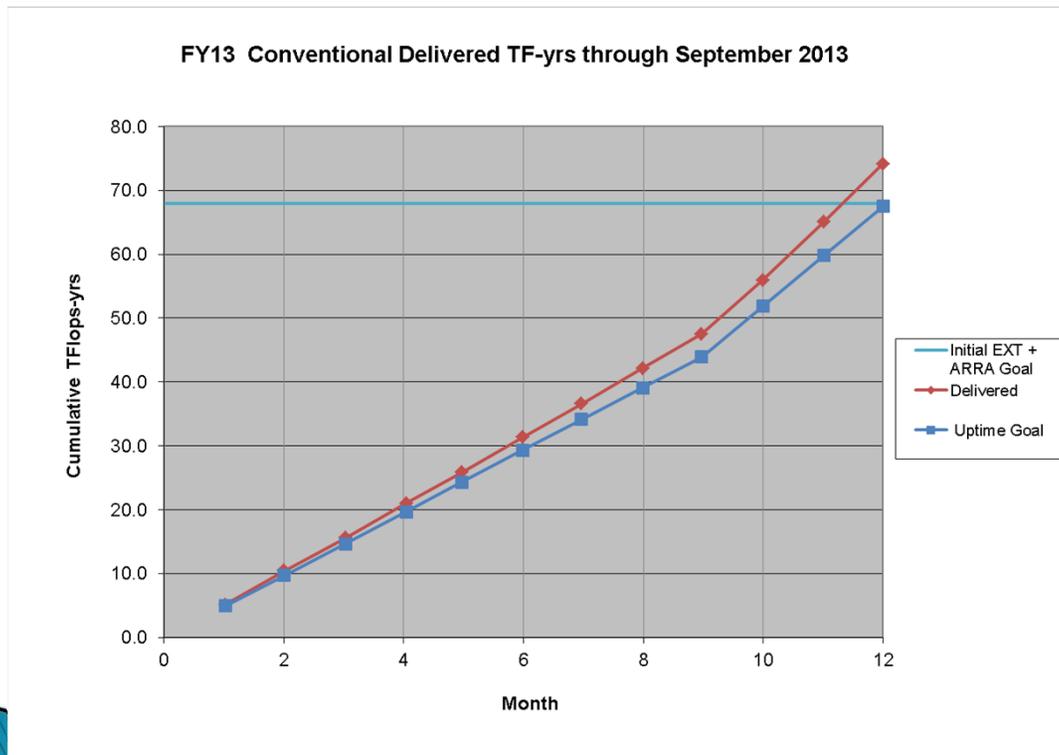
(2) Includes costs for the BG/Q deployment at BNL and the conventional cluster deployment at FNAL. Costs are below plan because effort to install the BG/Q was included in the procurement cost. In the original baseline plan, salary funds were budgeted to support a larger cluster deployment. Also, budget revisions post-baseline have shifted funds from procurements to operations to meet operational and storage needs.

(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.)

FY13 Milestone Performance – (TFlops–yrs delivered) Conventional Hardware

- ▶ Computing delivered in FY13 from conventional compute hardware is shown.
- ▶ Uptime goal is 8000 hours per year (91.3%); average uptime for the metafacility = 97%
- ▶ The unmodified goal for FY13 was 67.5 TFlops–yrs.
- ▶ The project achieved 74.1 TFlops–yrs (110% of goal).



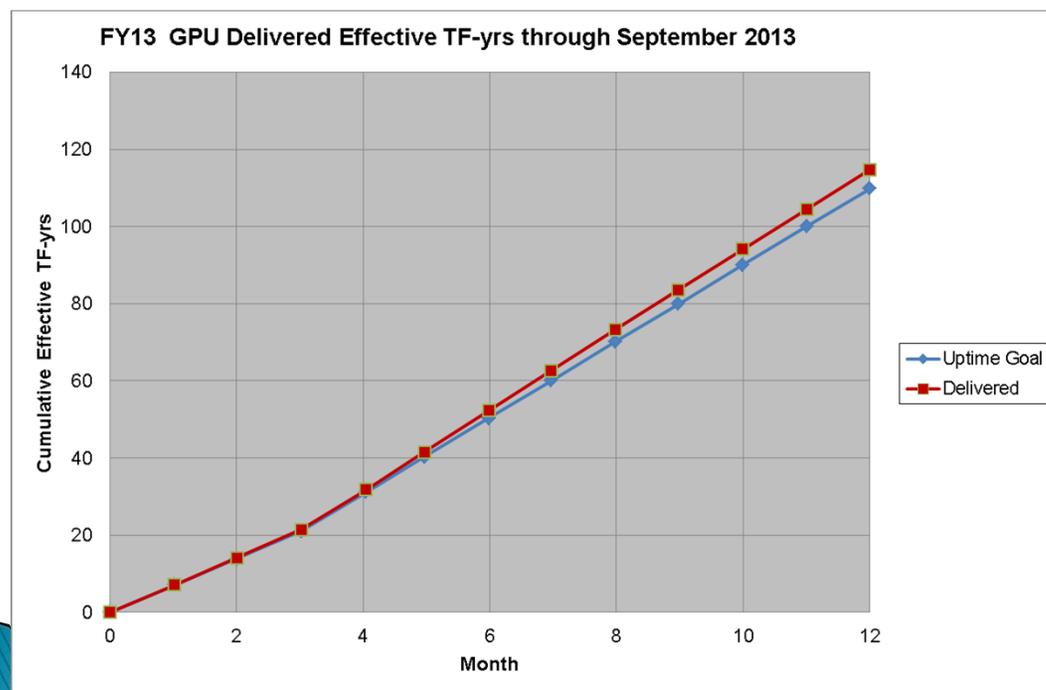
Conventional computing resources include:

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

(6 machines total)

FY13 Milestone Performance – (TFlops–yrs delivered) Accelerated Hardware

- ▶ Computing delivered in FY13 from accelerated compute hardware is shown.
- ▶ Uptime goal is 8000 hours per year (91.3%); average uptime for the metafacility = 95.6%
- ▶ Conversion from GPU–hrs to effective TF–yrs is 140 GF/GPU, based on allocation–weighted performance of GPU projects running since July 1, 2012.
- ▶ The goal for FY13 was 109.9 effective TF–yrs
- ▶ The project achieved 114.8 effective TF–yrs (105% of goal).



Computing resources include accelerated clusters operating at FNAL and JLab.

(5 machines total)

FY13 Performance Summary – KPIs

Performance against other Key Performance Indicators (KPIs)

Measurement Indicators	Target	Actual Results
Customer Satisfaction Rating	≥92%	94%
% of tickets closed within two business days	≥95%	87% FNAL: 97% JLab: 76%
% of average machine uptime at the metafacility	≥95%	Conventional: 97.6% (weighted ave.) Accelerated: 95.4% (weighted ave.)
Weekly vulnerability scans	Scans performed at least weekly at each host institution	Daily scans performed at all sites. Performance goal exceeded.

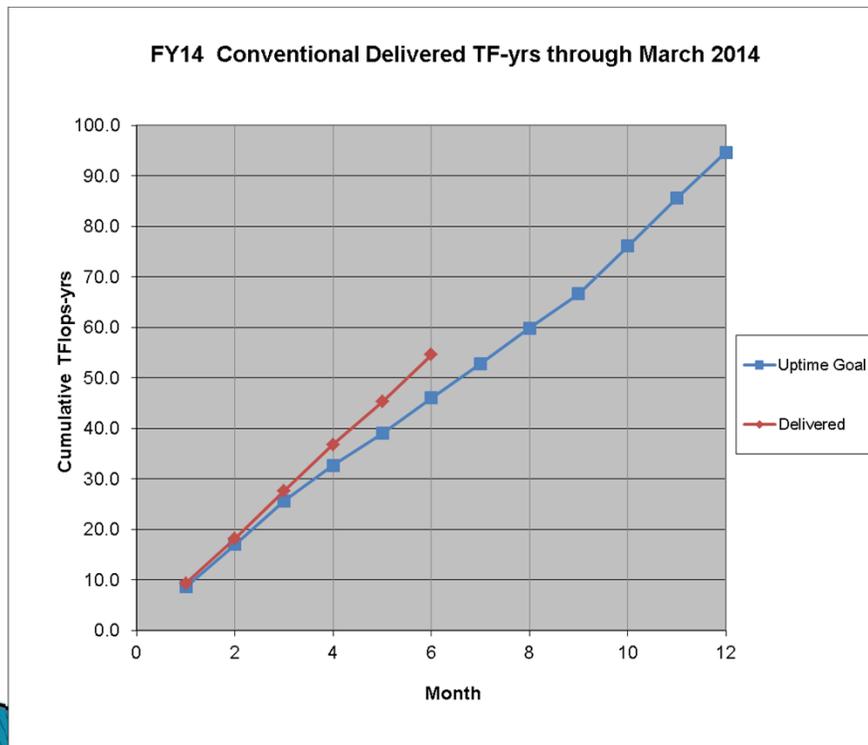
- ▶ All KPI metrics were met in FY13, with the exception of helpdesk ticket response time

FY13 Performance Summary

- ▶ All compute capacity and delivery performance goals were exceeded.
- ▶ We missed our cluster target deployment date of Jun 30 by 10 days due to funding delays and to minor complications associated with system deployment
 - Deployment of the conventional cluster at FNAL occurred 10 days later than goal, delays were due to late arrival of hardware and difficulties configuring supporting software (e.g., new Intel Infiniband hardware brought complications configuring MPI libraries) (all solved)
 - Deployment of the BG/Q at BNL occurred on schedule, with the official production release on July 1
- ▶ All KPI metrics were met with the exception of helpdesk ticket response time

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

- ▶ Data for FY14 conventional systems thru March 2014 are shown.
- ▶ The uptime goal is 8000 hours per year (91.3%), which equates to a conventional hardware goal for FY14 of 94.7 TFlops–yrs.
- ▶ Goal through March = 46.0 TFlops–yrs; Actual = 54.6 TFlops–yrs (119% of goal)



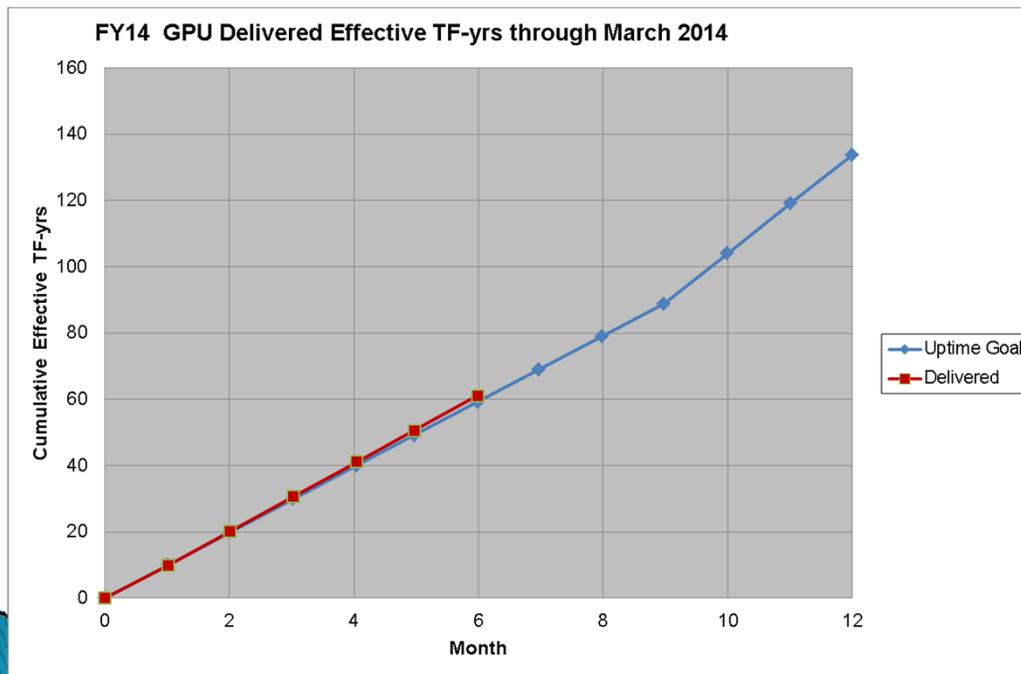
Computing resources included are the FNAL and JLab Infiniband clusters and the BNL BG/Q DD2 rack.

The inflection points in the uptime goal curve correspond to the retirement of the FNAL J/Psi cluster in January (delayed) and the deployment of the new cluster in July.

Because of funding delays related to the FY14 CRs, the procurement of the FY14 clusters has been delayed. This has resulted in higher than planned delivery capacity at FNAL early in the year, likely to be offset with shortfalls starting in July due to the deployment delay.

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

- ▶ Data for accelerated hardware systems thru March 2014 are shown.
- ▶ The uptime goal is 8000 hours per year (91.3%), which equates to an FY14 accelerated hardware goal of 133.8 Delivered Effective TFlops–yrs.
- ▶ Conversion from GPU–hrs to effective TF–yrs is 140 GF/GPU, based on allocation–weighted performance of GPU projects running since July 1, 2012.
- ▶ Goal through March = 59.2 effective TF–yrs; actual = 61.1 effective TF–yrs (103% of goal)

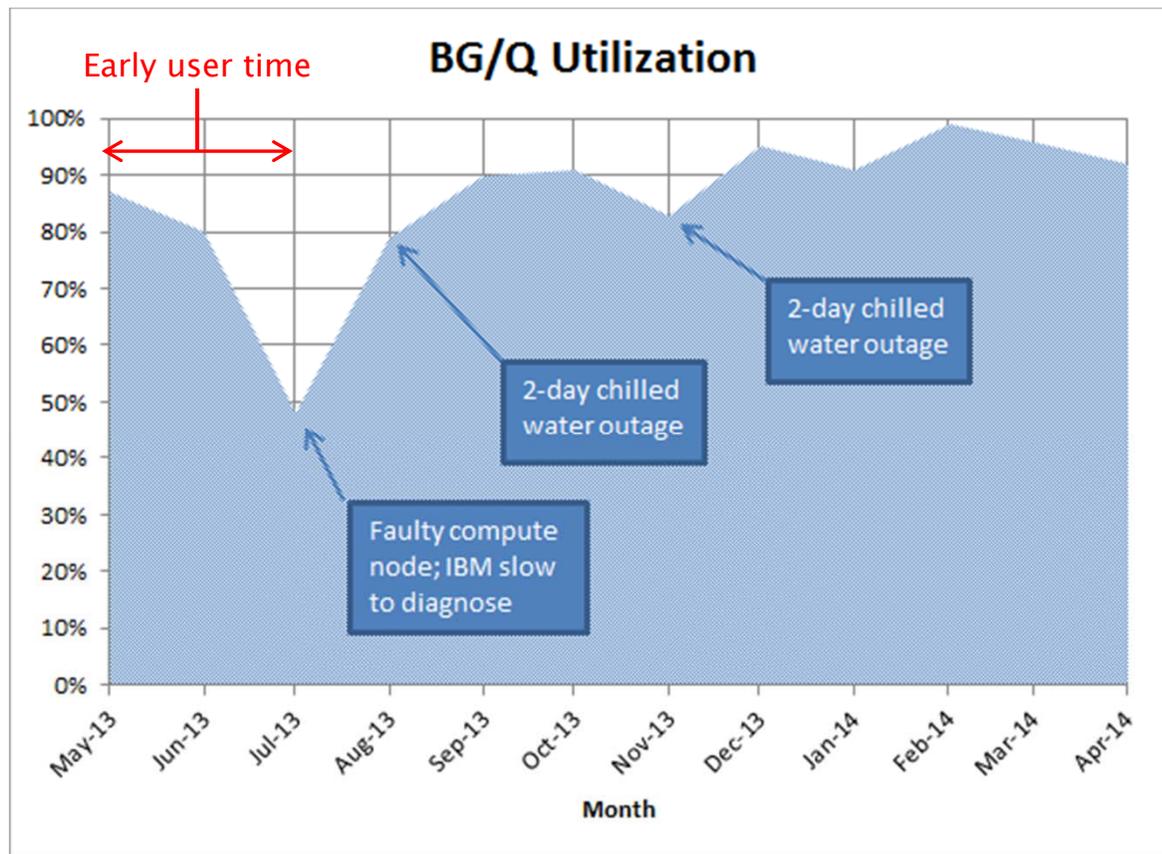


Computing resources included are the FNAL Dsg and JLab 9g, 10g, 11g, and 12k clusters.

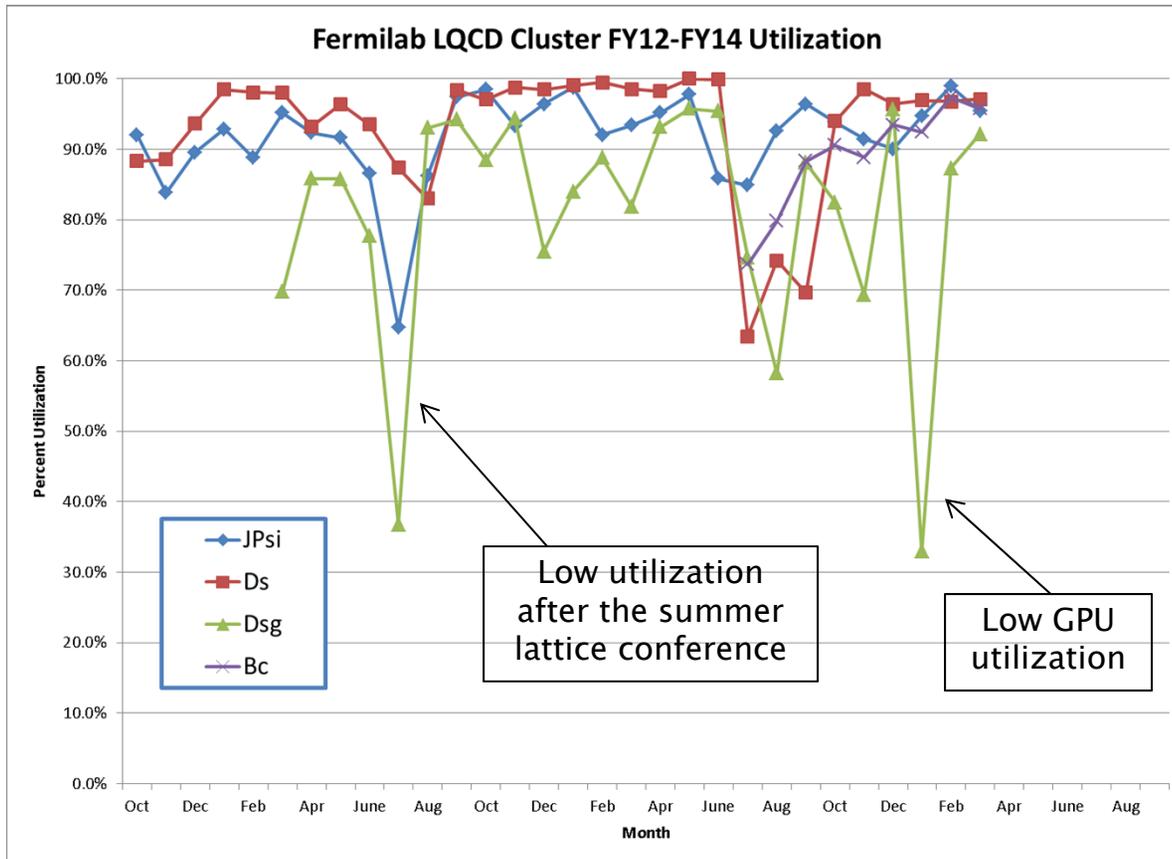
The inflection point in the curve reflects the deployment of FY14 GPU cluster.

BG/Q Utilization – BNL

	BG/Q Average Utilization
Aug '13 – Apr '14	92%



Cluster Utilization – Fermilab (conventional and accelerated clusters)



	Conventional Clusters Weighted Utilization	Accelerated Clusters Weighted Utilization
FY13	91	85%
FY14 (thru Apr)	96	79%

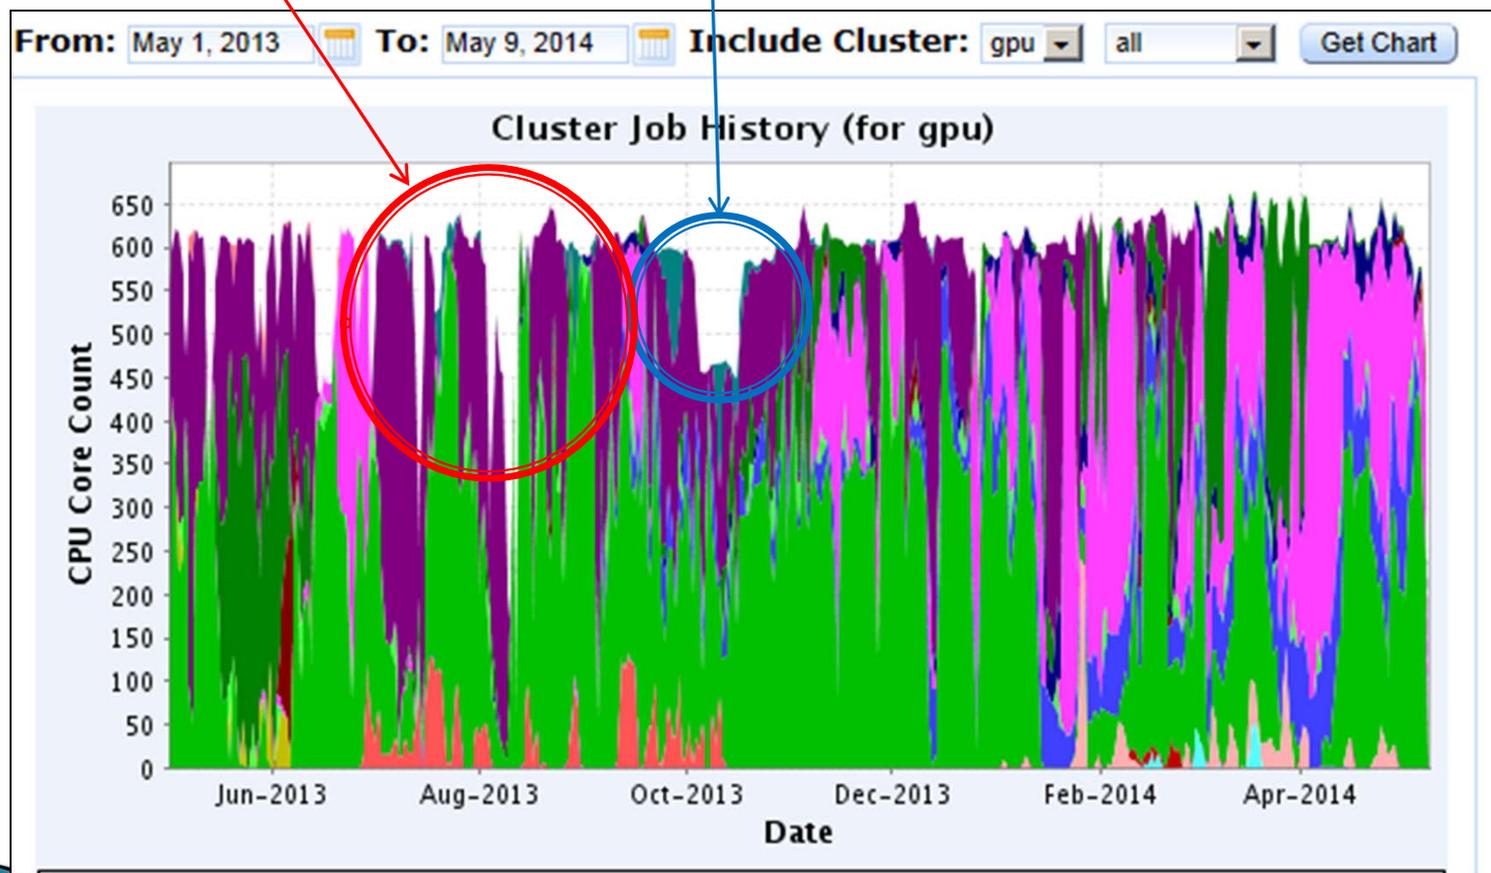
Note that on the FNAL GPU cluster, for long periods of time all jobs require 16 nodes, but since we don't have a multiple of 16 nodes in the cluster (we have 76), during those periods utilization can't exceed 84% (64/76); the projects that are allocated at FNAL tend not to have a lot of running using a small number of nodes per job.

Cluster Utilization – JLab (accelerated clusters)

Underutilization in July–Aug,
likely due to summer
conferences, vacations, etc.

Reduced operation pending
gov't shutdown

	Accelerated Clusters Weighted Utilization
FY13	90%
FY14 (thru Apr)	96%

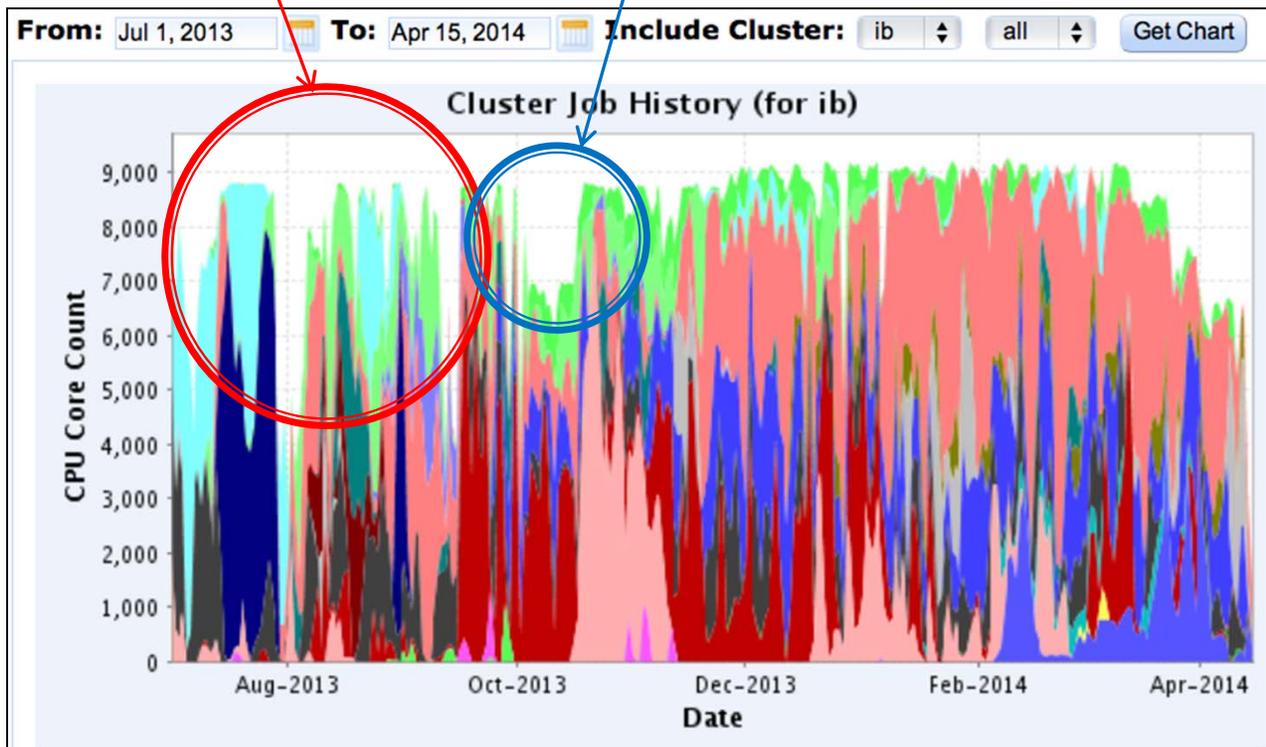


Cluster Utilization – JLab (conventional clusters)

Underutilization in July–Sep,
likely due to summer
conferences, vacations, etc.

Reduced operation pending
gov't shutdown

	Conventional Clusters Weighted Utilization
FY13	89%
FY14 (thru Apr)	96%

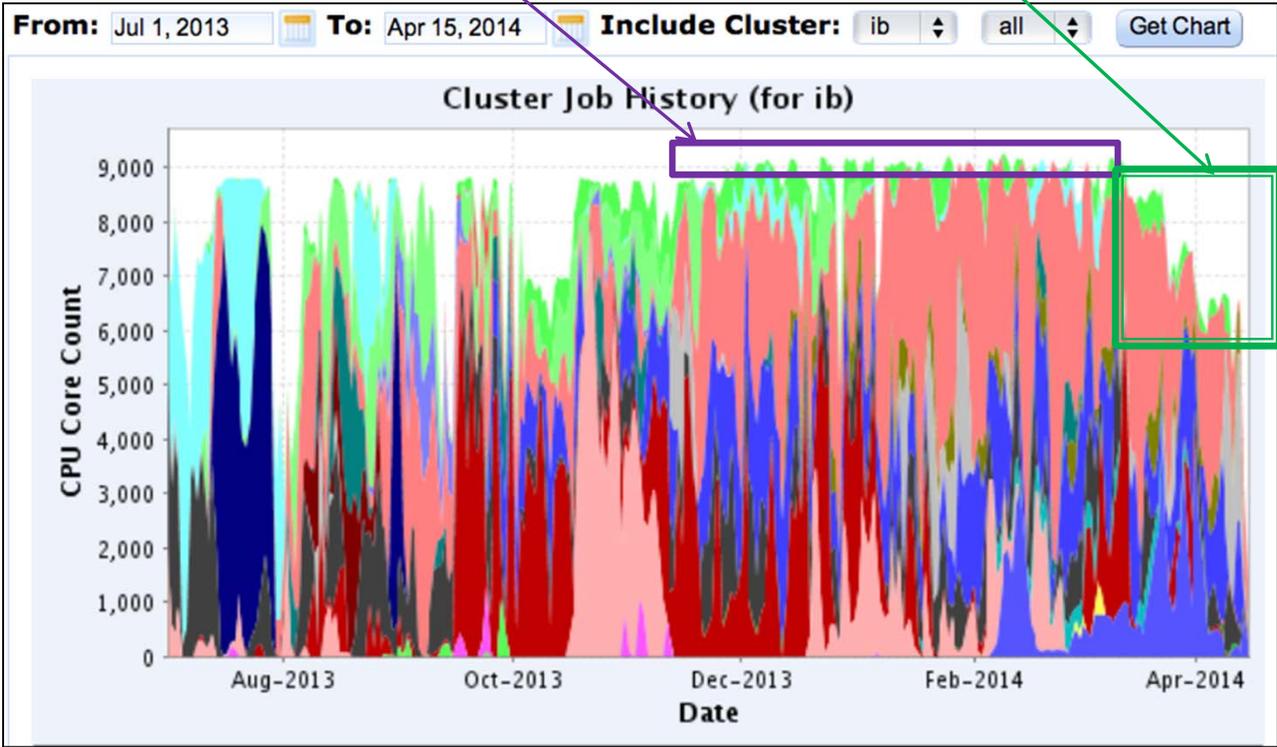


Cluster Utilization – JLab (conventional clusters)

Conventional Clusters Weighted Utilization	
FY13	89%
FY14 (thru Apr)	96%

1 million core-hr
“loan” from physics

“Loaned core-hrs” returned to
physics for Glue-X / CLAS-2
data challenges



Under-utilized cycles at JLab can be absorbed by the 12-GeV computing projects; this is a big help in late summer when each year LQCD utilization tends to sag; the cycles will be returned in the Fall when usage rebounds

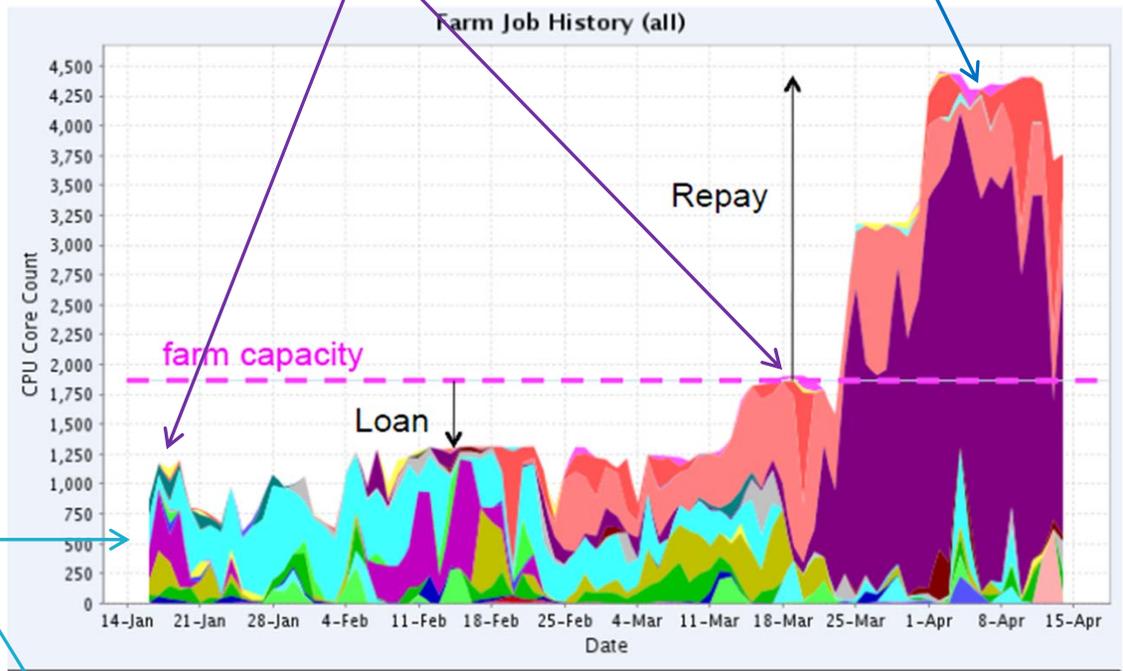
The sharing significantly benefits the large 12-GeV projects as they can also borrow significant resources for a few weeks at a time to conduct data challenges at scale in advance of full provisioning for 12-GeV.

Cluster Utilization – JLab (conventional clusters)

Conventional Clusters Weighted Utilization	
FY13	89%
FY14 (thru Apr)	96%

1 million core hour loaned from physics to LQCD

Loan returned to physics for data challenge



Clipped top area from plot on previous page

Total used by LQCD through the period of the “loan” was almost 1-million core hours.

The return flow was 4x larger (2048 cores of 12S, or about 22% of our capacity) but over a shorter period of time. However, this 22% of our capacity was sufficient to increase the Physics “farm” capacity by ~2.5x.

Win-Win
As long as the JLab farm nodes on loan are kept busy, LQCD users don’t have to worry as much about losing any of their allocations due to lost core-hrs.

Physics can do data challenges at much larger scale than their current resources allow, for this and next, which allows for delaying purchases until needed for production running.

Financial Performance

FY13 Project Cost Summary – Final

Fund Type	FY12 Carry-over	FY13 Budget	Total FY13 Funds Available	FY13 Actual Costs	FY13 Obligations	% Spent & Obligated
Equipment	\$ 811K	\$ 2,170K	\$ 2,981K	\$ 2,955K	\$ 0K	99%
Operating	\$ 253K	\$ 1,931K	\$ 2,184K	\$ 1,826K	\$ 45K	86%
Sub-total	\$ 1,064K	\$ 4,101K	\$ 5,165K	\$ 4,782K	\$ 45K	93%
Mgmt Reserve	\$ 110K	\$ 99K	\$ 209K	---	---	0%
TOTAL	\$ 1,174K	\$ 4,200K	\$ 5,374K	\$ 4,782K	\$ 45K	90%

Cost Performance Analysis

- FY13 equipment costs are associated with the BG/Q procurement at BNL and the conventional cluster procurement at FNAL. All FY13 hardware has been fully costed.
- Operating funds support steady-state and deployment SWF as well as providing funds for storage hardware, tape, spares, etc. Operating fund spend rate at BNL and FNAL was lower than anticipated.
- Open commitments include a Luster file server and GTX graphics cards for JLab.
- No management reserve funds were expended. Unspent management reserve will be applied to the FY14 hardware acquisition.

FY14 Project Cost Summary

Status through March 2014; fiscal year complete: 50%

Fund Type	FY13 Carry-over	FY14 Budget	Total FY14 Funds Available	FY14 Actual Costs	FY14 Obligations	% Spent & Obligated
Equipment	\$ 7K	\$ 1,800K	\$ 1,807K	\$ 0K	\$ 1,822K	101%
Operating	\$ 522K	\$ 2,400K	\$ 2,729K	\$ 969K	\$ 0K	36%
Sub-total	\$ 529K	\$ 4,157K	\$ 4,536K	\$ 969K	\$ 0K	21%
Mgmt Reserve	\$ 0K	\$ 43K	\$ 43K	---	---	0%
TOTAL	\$ 529K	\$ 4,200K	\$ 4,579K	\$ 969K	\$ 0K	21%

Cost Performance Analysis

- ▶ Operating expenses to date remain below where they should be due to very late arrival of funds at JLab. Base funds were used to support operations and the project will incur higher expenses later in the fiscal year to compensate.
- ▶ We have modified the FY14 budget based on operating experience. The revised financial plan has been accepted and the laboratories are fully funded per the revised plan.
- ▶ The requisition to start the RFP process for the FY14 acquisition is in process at FNAL. The req amount is show above in the “Obligations” column. The obligation amount is slightly over budget because the requisition was entered into the system before we had approval on the revised financial plan. Since we can control the amount of the actual contract award, we will not issue an award that exceeds the budget.

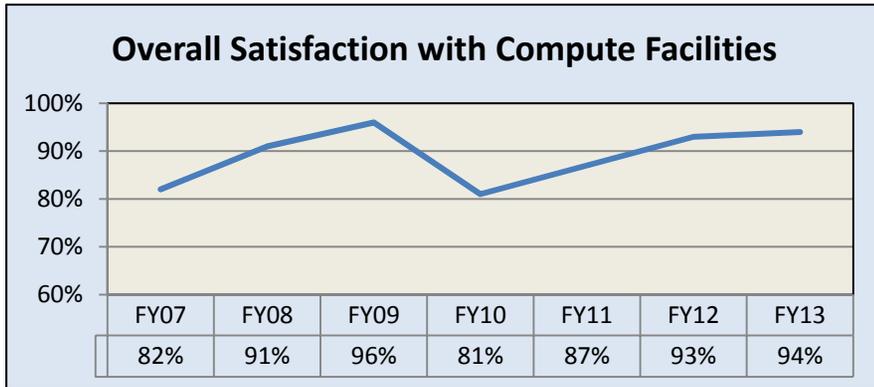
User Survey Results

Rob Kennedy

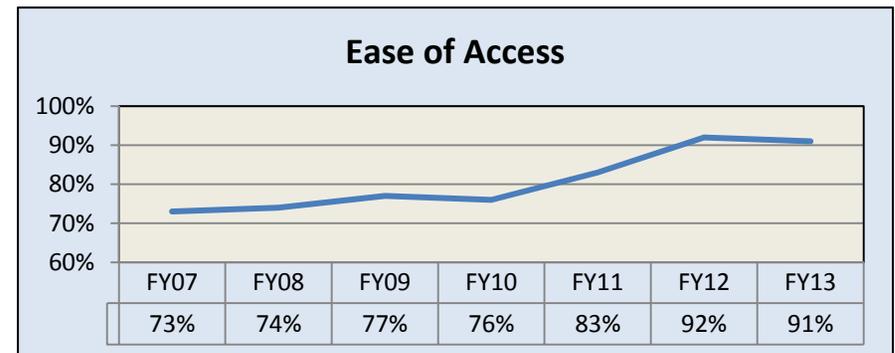
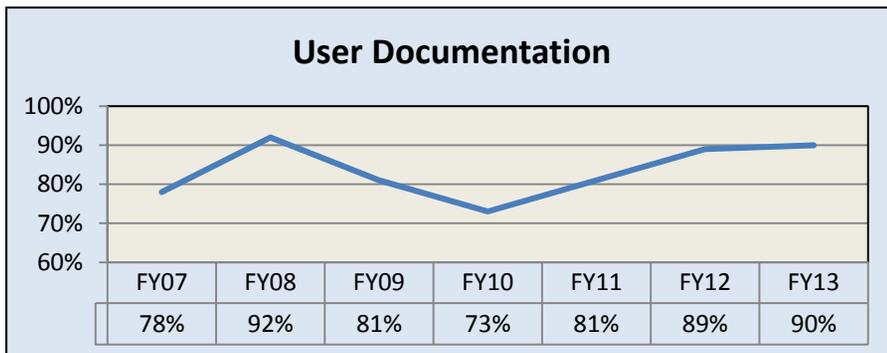
FY13 Survey Results

- ▶ The FY13 User Survey measured user satisfaction during the 17 month period from October 2012 through February 2013 inclusive.
- ▶ The survey consisted of 29 questions designed to measure satisfaction with the compute facilities and the resource allocation process.
- ▶ The survey was distributed to 158 individuals
 - Responses were received from 66 individuals
 - By comparison, 76 individuals responded to the FY12 survey
 - For FY14 – Improve response rate by improving timing and communication of survey
- ▶ FY13 overall satisfaction ratings exceeded our target goal of 92%:
 - 94% for Compute Facilities
 - 97% for Allocation Process

Compute Facility Satisfaction Trends

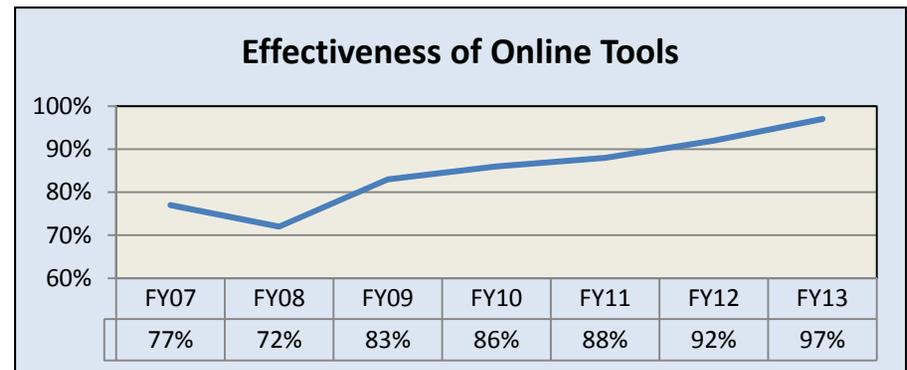
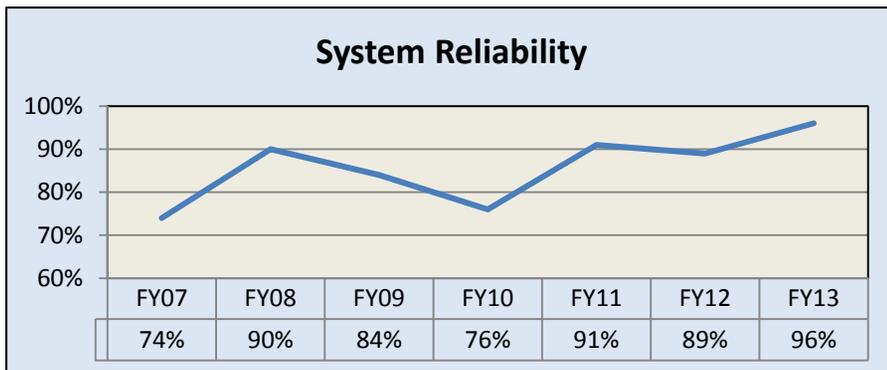
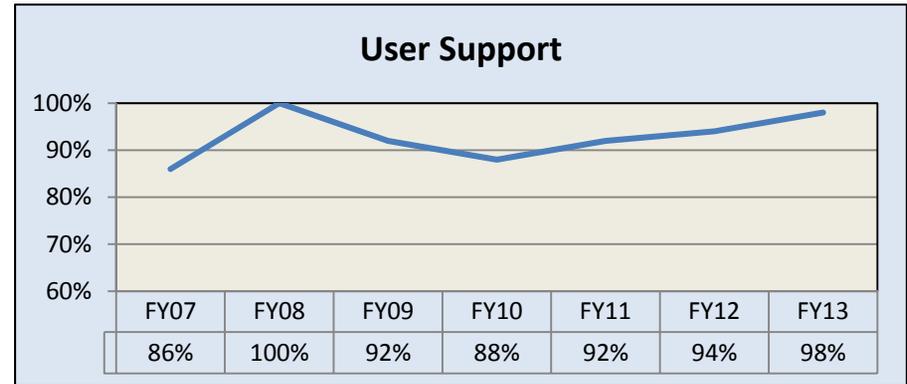
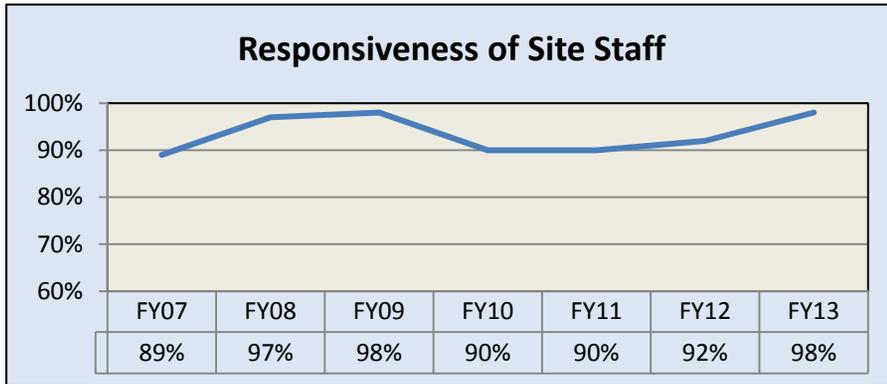


FY13 Computing Facilities	All Sites	BNL	FNAL	JLab
Overall satisfaction	94%	85%	96%	95%
Documentation	90%	64%	91%	95%
User support	98%	96%	98%	99%
Responsiveness	98%	97%	98%	99%
Reliability	96%	97%	97%	94%
Ease of access	91%	97%	94%	84%
Tools support	97%	96%	98%	95%



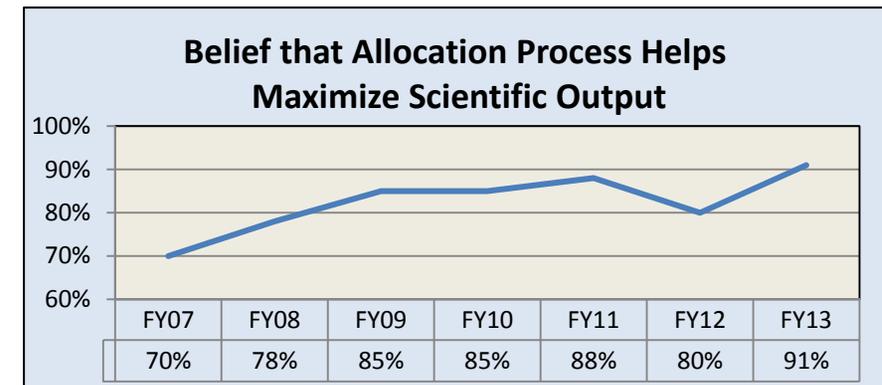
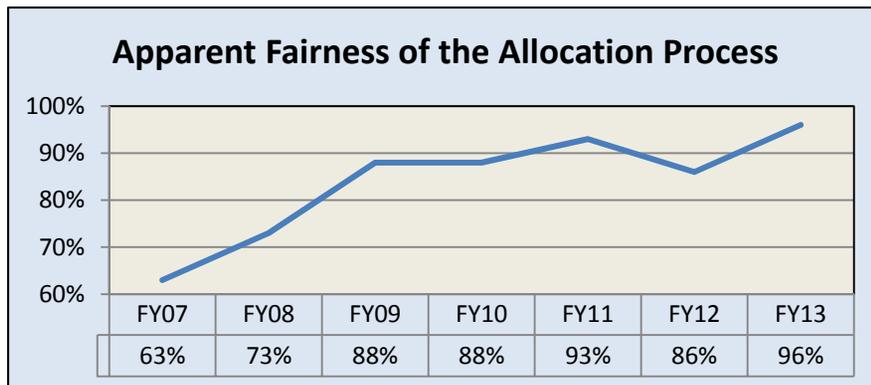
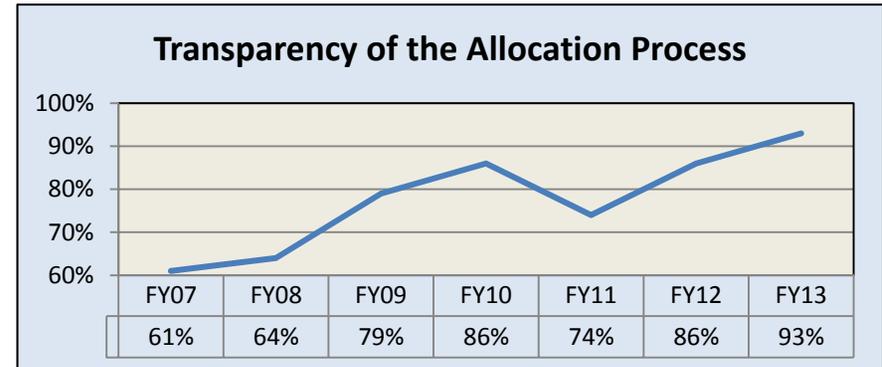
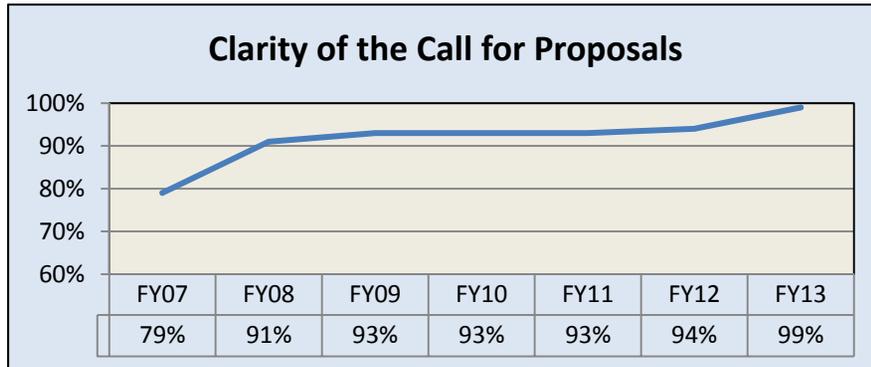
- ▶ The overall satisfaction rating has been trending upward over the past three years.
- ▶ FY13 rating of 94% exceeds our target goal of 92%, and is about the same as FY12.
- ▶ **JLab's Overall Satisfaction rating** of 95% in FY13, much higher than FY12 rating of 76%.
- ▶ **BNL's satisfaction rating for User Documentation** was below par, but very low statistics.
- ▶ Ease of Access and User Documentation ratings stayed about the same in FY13.

Compute Facility Satisfaction Trends



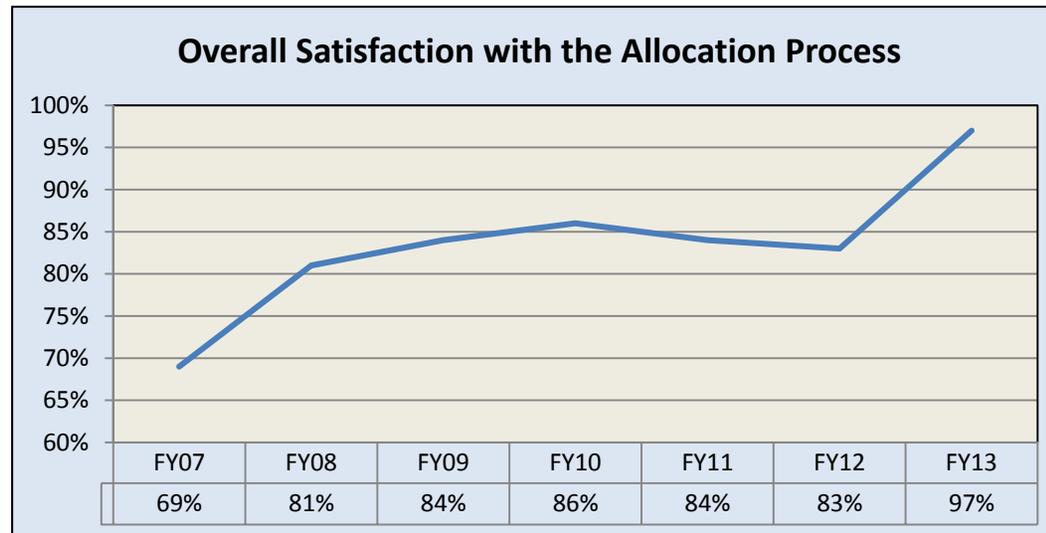
- ▶ Responsiveness of Site Staff and User Support continue their recent upward trend.
- ▶ System Reliability and Online Tools also continue their recent upward trend.

Allocation Process Satisfaction Trends



- ▶ Satisfaction rating trends for all Allocation Process survey areas are steady or improving in FY13.
- ▶ All recent dips in the satisfaction ratings have been reversed and recovered.

Allocation Process Satisfaction Trends



- ▶ The overall satisfaction rating for the Allocation Process jumped to 97% in FY13 from 83% in FY12, significant even at these statistics levels.
- ▶ User text comments do not suggest a specific cause for this jump. Project view is:
 - Scientific Program Committee is doing a better job at categorizing the proposals and communicating the breakdowns to USQCD.
 - Establishment of Scientific Advisory Board may have helped satisfaction ratings too.

FY13 User Survey Summary

▶ Satisfaction with Compute Facilities

- The overall satisfaction rating of 94% exceeds our target goal of 92%.
- User support, responsiveness, and system reliability ratings at one of our sites recovered well in FY13 from a dip in FY12.
- 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 97% exceeds our target goal of 92%.
- User satisfaction ratings exceeded prior year ratings in all categories.
- The Executive Committee and Scientific Program Committee have continued to work hard to ensure that scientific goals and the criteria for proposal evaluation are clearly understood.
- Establishment of the Scientific Advisory Board to ensure experimenters have a formal role in the allocation process may have helped ratings too.

Summary

- ▶ Compute facilities are running well and we're successfully executing against our plans.
- ▶ We successfully met or exceeded all but one of key performance goals in FY13. We did not meet our target deployment date for the cluster deployment.
 - We missed deployment milestones due to the impact of delayed funding as a result of Continuing Resolutions and minor commissioning problems.
- ▶ We are on target to meet our FY14 performance goals, with the exception of deployment milestone dates.
 - New systems are being deployed according to plan, albeit late due to funding delays.
 - Our site managers continue to do a very good job of operating their respective systems in a manner that minimizes downtime and maximizes output.
- ▶ We continue to work hard to maximize our hardware portfolio and have developed and executed an acquisition plan to optimize our procurement strategy, which allows us to make the most effective use of project resources.
 - We have successfully followed this process in past years with successful results; we are following a similar approach this year.