

# LQCD-ext II Project Management & Performance

Bill Boroski  
LQCD Contractor Project Manager

Robert D. Kennedy  
Associate 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
- ▶ FY15 year-to-date performance results
- ▶ FY15 year-to-date financial results
- ▶ User survey results
- ▶ Summary

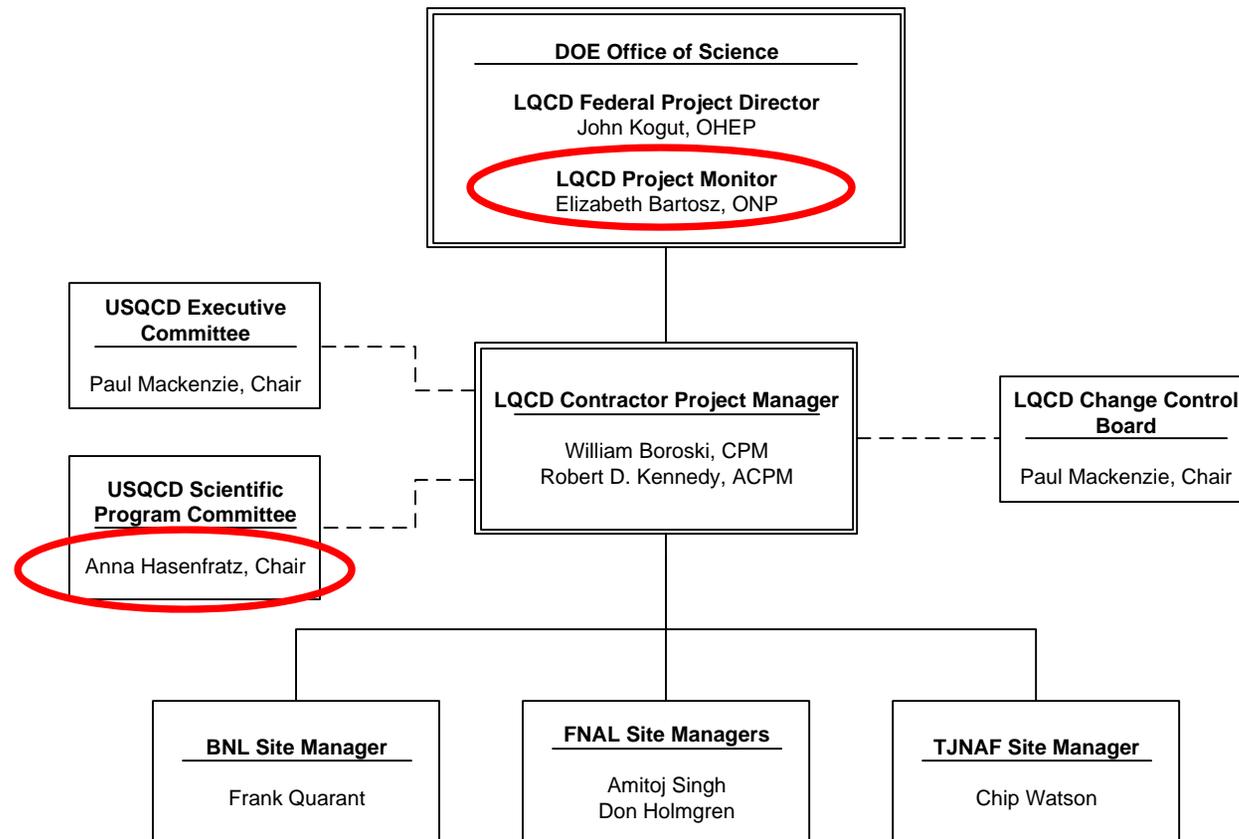
# LQCD-Ext II Project Scope

- ▶ Acquire and operate dedicated hardware at BNL, FNAL, and JLab for the study of QCD during the period FY2015–2019.
- ▶ Scope includes acquisition, deployment, and operation of computing facilities; software development is out of scope.
- ▶ Currently executing against baseline plan, with one exception
  - Baseline plan did not include a new system deployment in FY15.
  - Executed FY15 procurement of 100 node expansion for Pi0 cluster at Fermilab using deferred FY14 funds.

# 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
- ▶ **Risk Management**
  - Risks are regularly reviewed as part of our active Risk Management program. All risks are reviewed at a frequency commensurate with their risk priority (e.g., High=monthly; Medium=quarterly, etc.). In addition, each risk is assigned a “Next Review Date” to ensure adequate risk monitoring.

# Management & Oversight



- ▶ Organizational changes between LQCD-ext and LQCD-ext II:
  - Elizabeth Bartosz replaced Kawtar Hafidi as NP Project Monitor
  - Anna Hasenfratz replaced Robert Edwards as SPC Chair

# LQCD-ext II Project Budget

- ▶ Approved Baseline Budget = \$14 million
  - Jointly funded by DOE Offices of High Energy and Nuclear Physics

## Approved Funding Profile (in \$K)

<b>Expenditure Type</b>	<b>FY15</b>	<b>FY16</b>	<b>FY17</b>	<b>FY18</b>	<b>FY19</b>	<b>Total</b>
Personnel	1,654	1,766	1,525	1,634	1,328	7,908
Travel	17	17	17	17	17	84
M&S	283	283	283	102	102	1,053
Compute/Storage Hardware	-	847	1,114	1,161	1,489	4,611
Management Reserve	46	87	61	86	64	344
<b>Total</b>	<b>2,000</b>	<b>3,000</b>	<b>3,000</b>	<b>3,000</b>	<b>3,000</b>	<b>14,000</b>
Planning Budget Guidance	2,000	3,000	3,000	3,000	3,000	14,000

# LQCD-ext II Project Budget

- ▶ Approved Baseline Budget = \$14 million
  - Jointly funded by DOE Offices of High Energy and Nuclear Physics

## Approved Funding Profile (in \$K)

Expenditure Type	FY15	FY16	FY17	FY18	FY19	Total
Personnel	1,654	1,766	1,525	1,634	1,328	7,908
Travel	17	17	17	17	17	84
M&S	283	283	283	102	102	1,053
Compute/Storage Hardware	-	847	1,114	1,161	1,489	4,611
Management Reserve	46	87	61	86	64	344
<b>Total</b>	<b>2,000</b>	<b>3,000</b>	<b>3,000</b>	<b>3,000</b>	<b>3,000</b>	<b>14,000</b>
Planning Budget Guidance	2,000	3,000	3,000	3,000	3,000	14,000

## Hardware Budget Breakdown (in \$K)

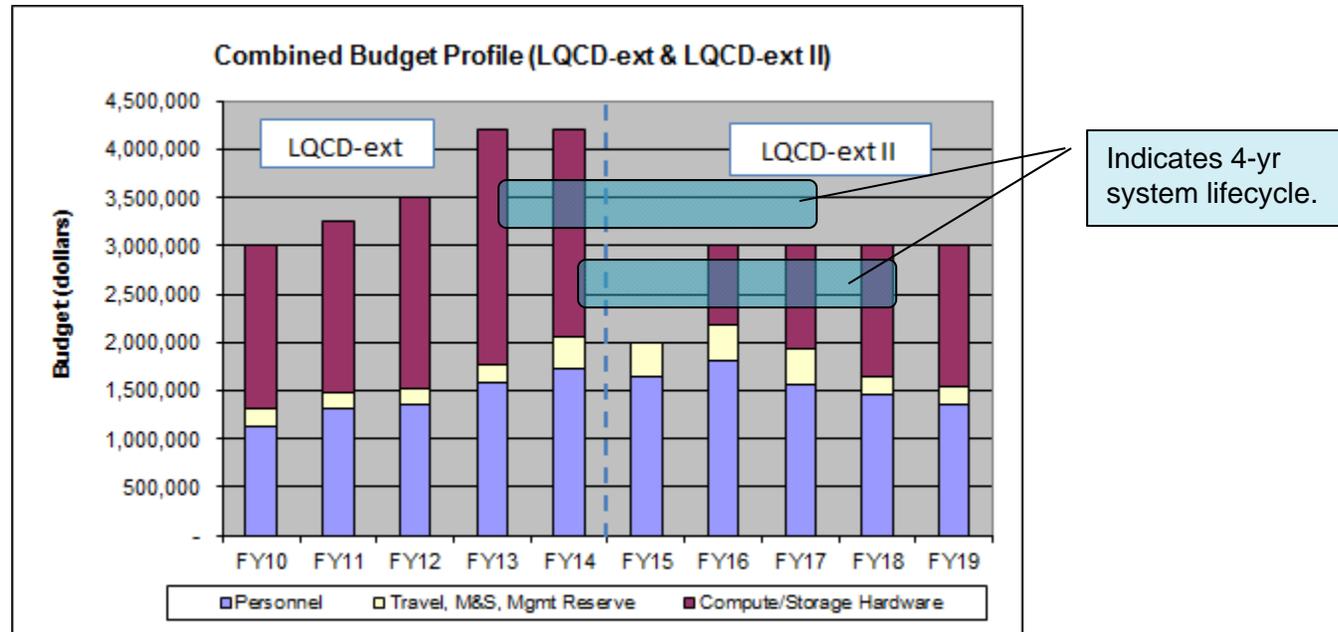
Fiscal Year	Compute Hardware	Storage Hardware	Total
FY15	-	-	-
FY16	779	68	847
FY17	1,024	89	1,114
FY18	1,068	93	1,161
FY19	1,370	119	1,489
<b>Total</b>	<b>4,242</b>	<b>369</b>	<b>4,611</b>

Equipment budget is used to procure compute and storage hardware

Storage budget set at 8% of total hardware budget to meet collaboration needs.

# Project Budget Comparison

*Approved Budget Baseline: \$14 million*



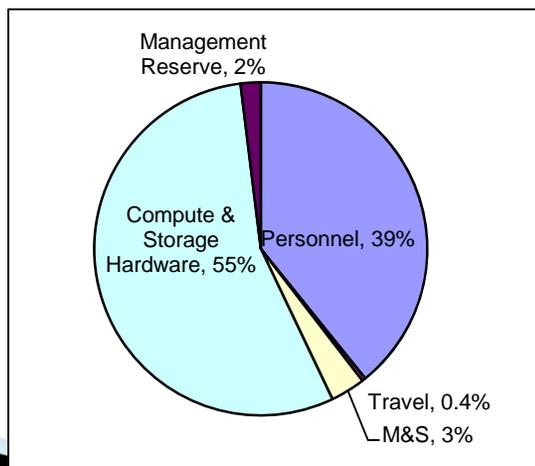
- ▶ The approved \$14 million budget represents a significant reduction in funding from prior levels, which had been back-loaded in the funding profile for the previous project (LQCD-ext).
- ▶ Personnel cost requirements are based on a refined staffing model. Level of operations support is based on number 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 Budget Comparison

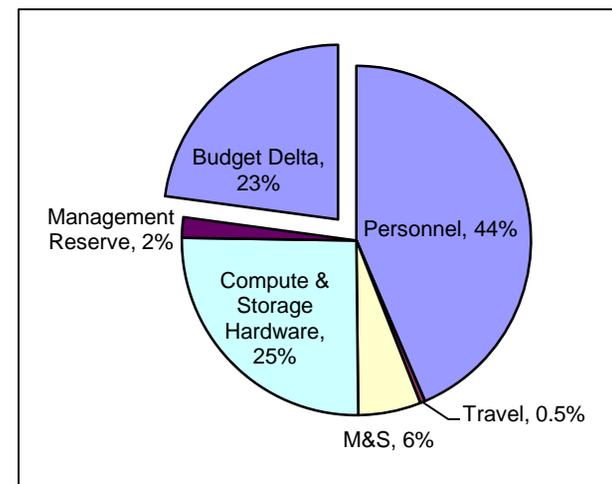
## Approved Funding Profiles (in \$K)

Expenditure Type	LQCD-ext Total	LQCD-ext II Total	Change
Personnel	6,885	7,908	1,023
Travel	60	84	24
M&S	440	1,053	613
Compute/Storage Hardware	10,405	4,611	(5,794 )
Management Reserve	360	344	(16 )
<b>Total</b>	<b>18,150</b>	<b>14,000</b>	<b>(4,150 )</b>

LQCD-ext



LQCD-ext II  
(Normalized to \$18.15 M Total)



# Performance Measures & Metrics

# 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 deferred 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.
- We are considering planning and executing procurements across fiscal year boundaries at JLab in FY16–17 and FNAL at FY18–19 in order to reduce planning and overhead costs. However, as part of our annual acquisition planning process, we will determine the procurement strategy that best matches the collaboration needs and optimizes the use of our resources to maximize scientific output.

# Performance Measures & Metrics

- ▶ Performance goals and milestones for LQCD-ext II 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.
  
- ▶ 23 Level-1 project milestones (for LQCD-ext II)
  - External reviews of future procurement plans
  - Incremental procurements/TFlops-deployed
  - Aggregate TFlops-yrs delivered
  
- ▶ 10 cost and schedule performance metrics
  - Planned costs and schedule completion dates
  
- ▶ 37 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 II Federal Project Director and Project Monitor.

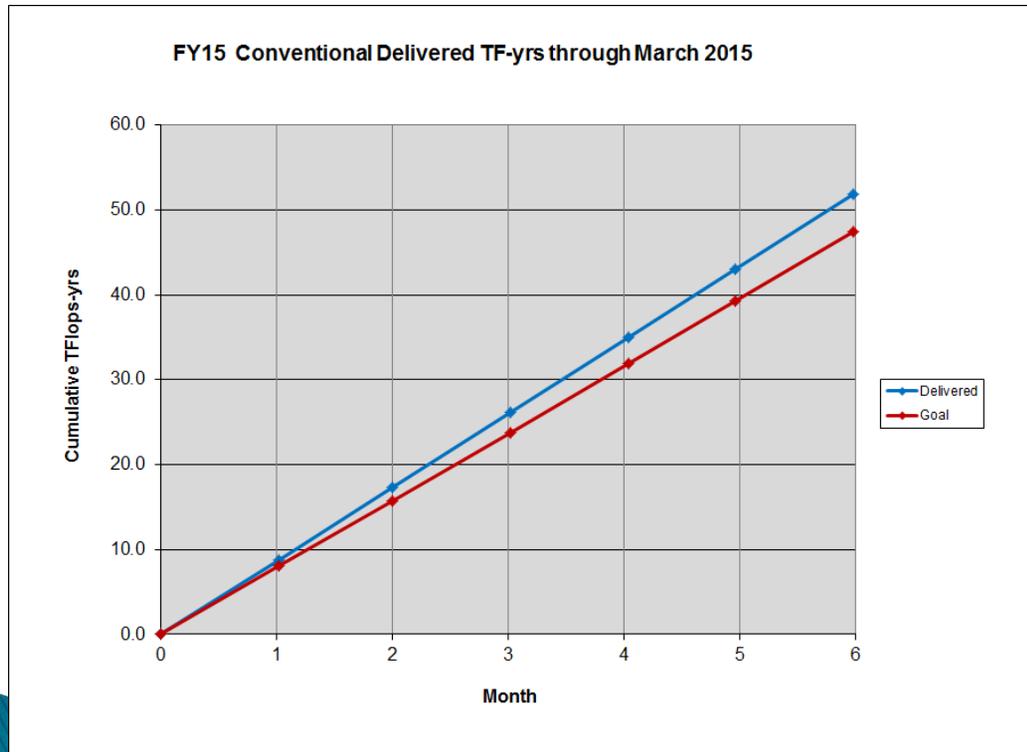
# Compute Hardware Performance Data

- ▶ Performance and utilization data are available online for LQCD-ext II resources at FNAL and JLab
  - 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, power usage.
  - 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

# FY15 Performance To Date

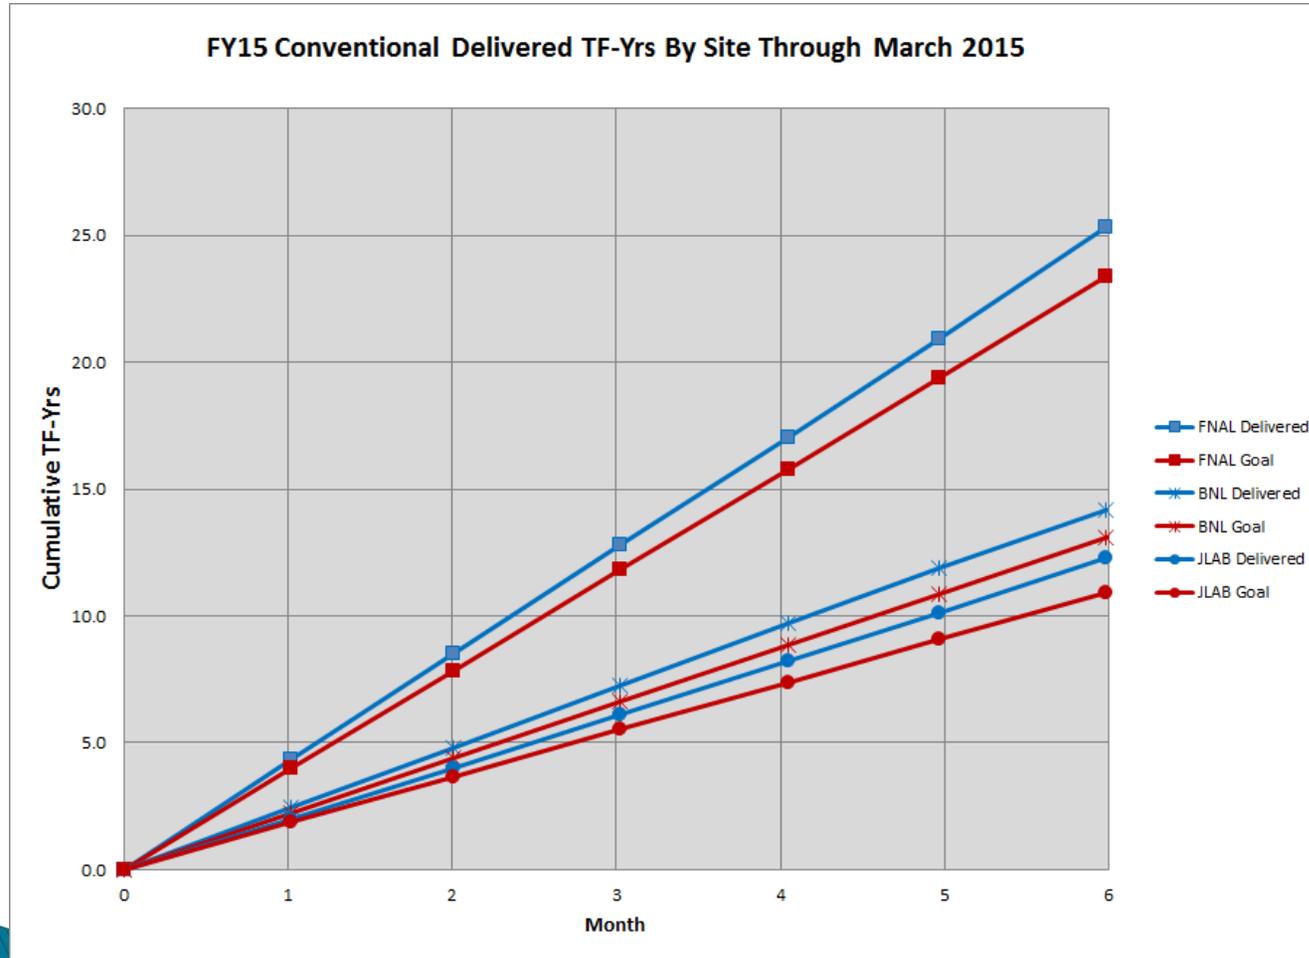
# FY15 Milestone Performance – (TFlops–yrs delivered) Conventional Hardware

- ▶ Data for FY15 conventional systems thru March 2015 are shown.
- ▶ The uptime goal is 8000 hours per year (91.3%), which equates to a conventional hardware goal for FY15 of 95.1 TFlops–yrs.
- ▶ Goal through March = 47.4 TFlops–yrs; Actual = 51.8 TFlops–yrs (109% of goal)



Computing resources included are the FNAL and JLab Infiniband clusters, the BNL BG/Q DD2 rack, and 10% of the BNL BG/Q DD2 prototype rack.

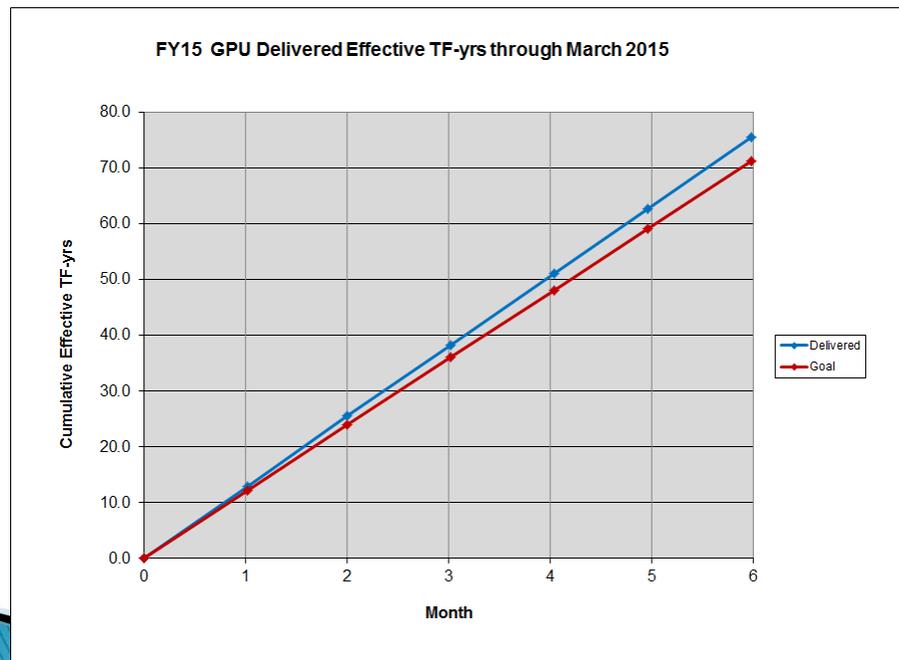
# Conventional Hardware Delivered Performance by Site



All sites  
exceeding  
delivery goals  
due to high  
system  
uptimes

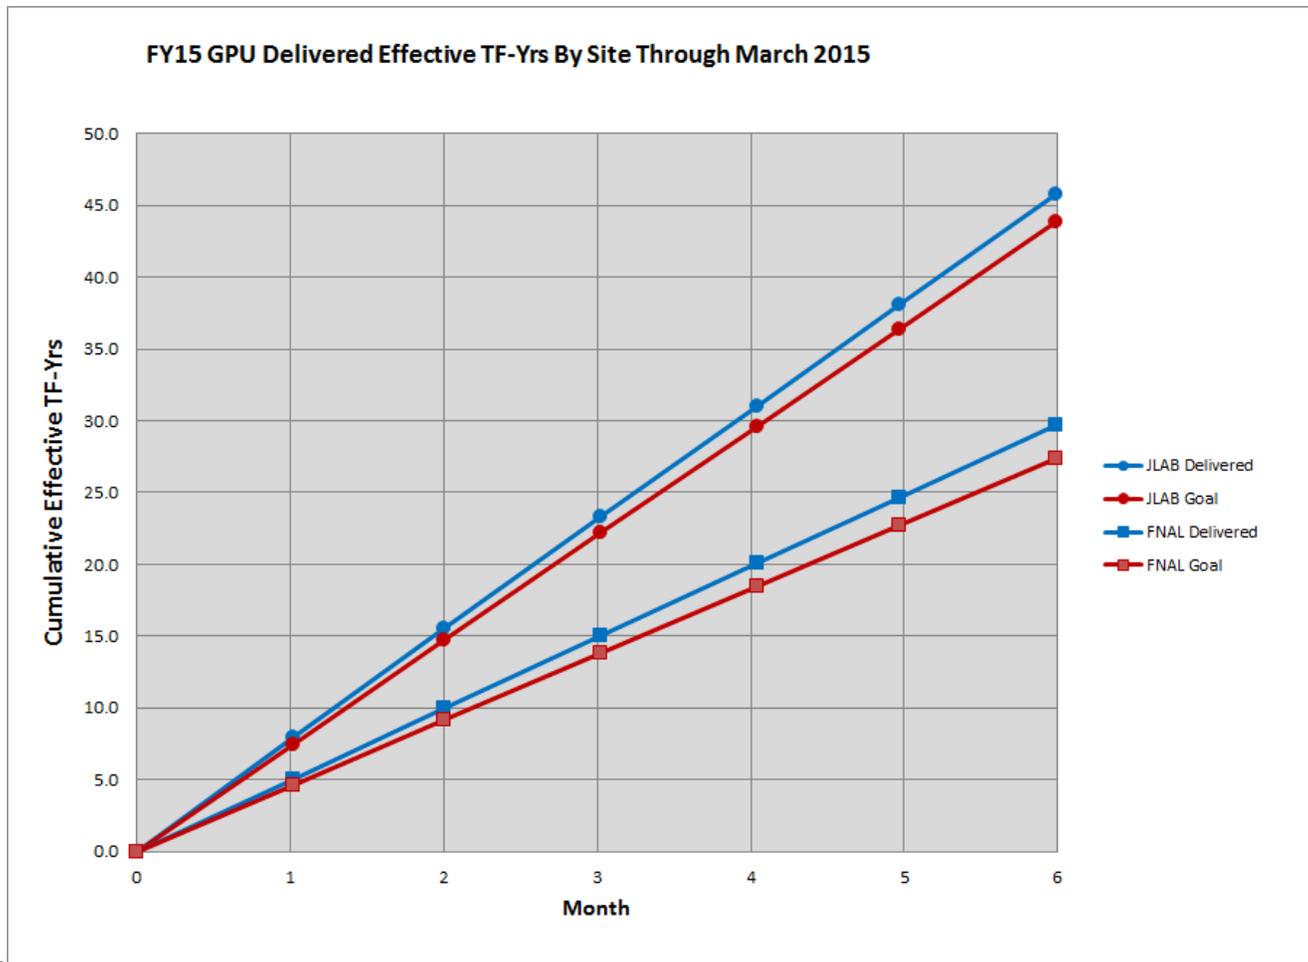
# FY15 Milestone Performance – (TFlops–yrs delivered) Accelerated Hardware

- ▶ Data for accelerated hardware systems thru March 2015 are shown.
- ▶ The uptime goal is 8000 hours per year (91.3%), which equates to an FY15 accelerated hardware goal of 142.8 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 from July 2012 through December 2012.
- ▶ Goal through March = 71.2 effective TF–yrs; actual = 75.5 effective TF–yrs (106% of goal)



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

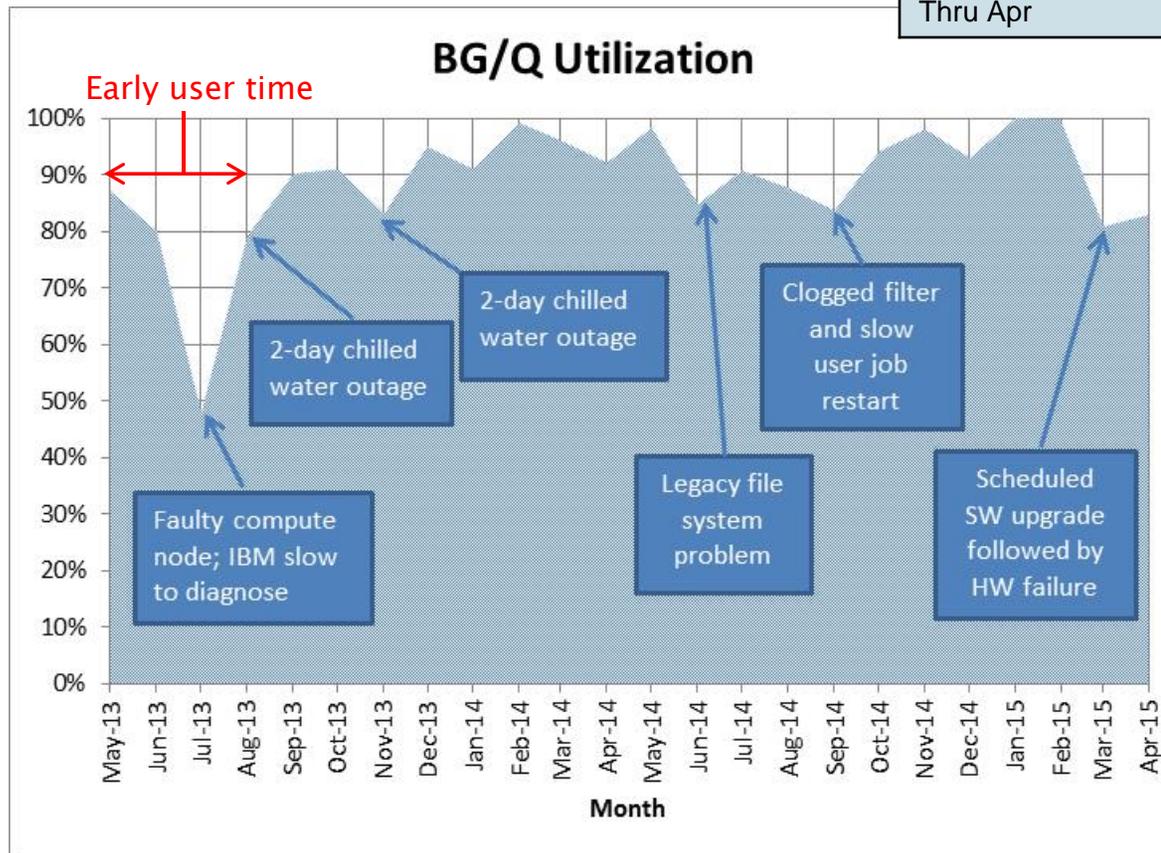
# Accelerated Hardware Delivered Performance by Site



All sites  
exceeding  
delivery goals  
due to high  
system  
uptimes

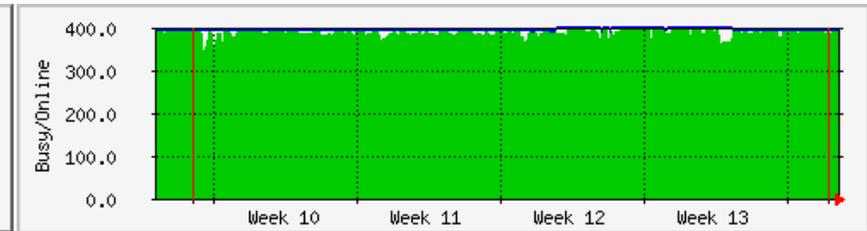
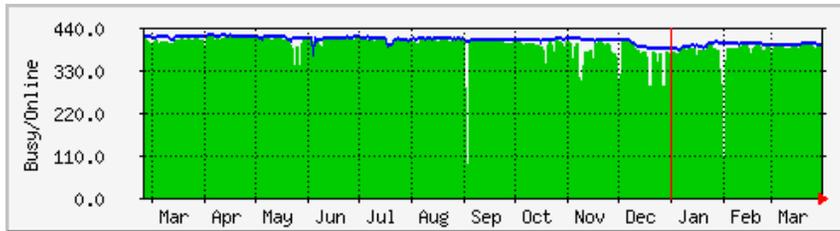
# BG/Q Utilization – BNL

	BG/Q Average Utilization
FY13	85%
FY14	91%
FY15 Thru Apr	93%

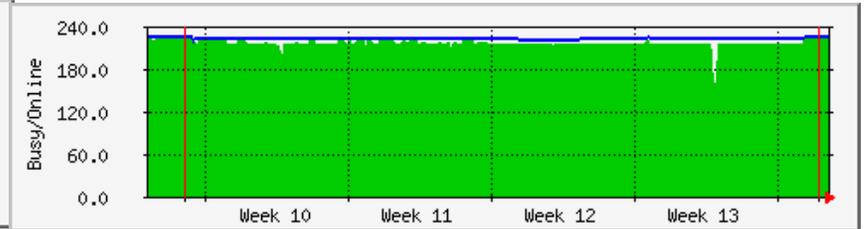
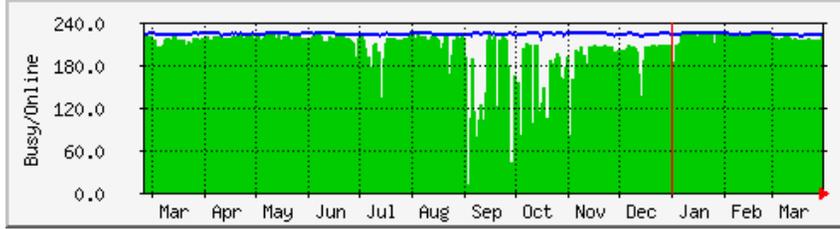


# FY15 Cluster Utilization – FNAL

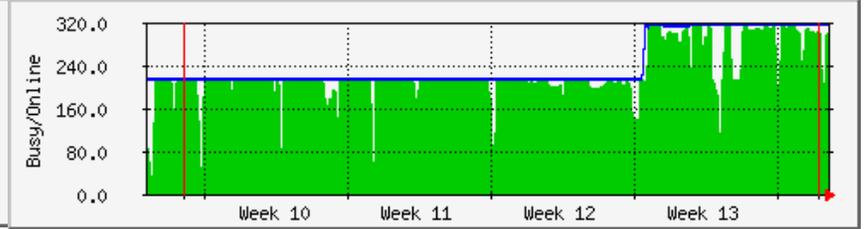
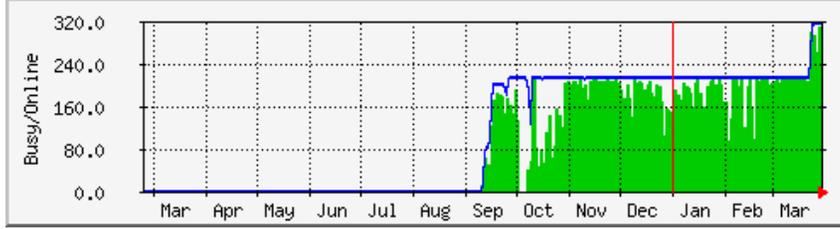
Ds



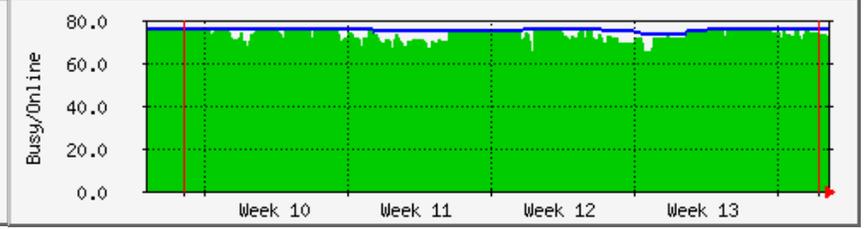
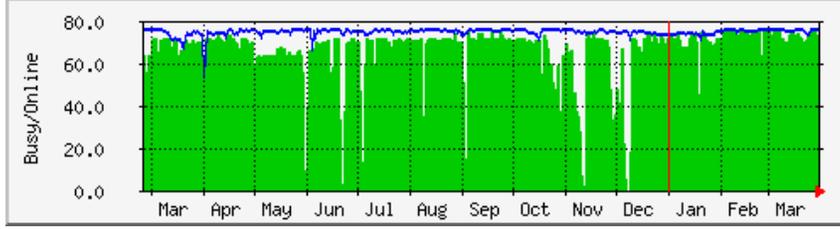
Bc



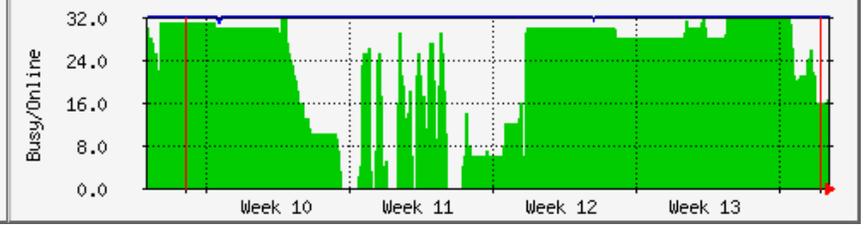
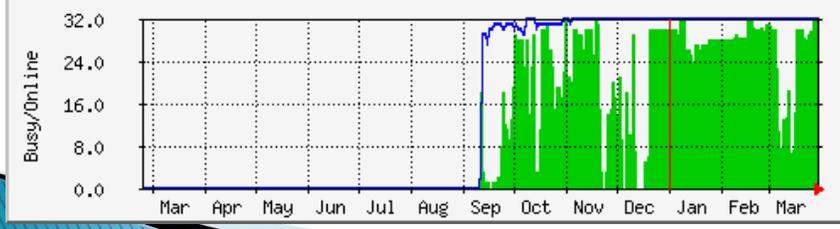
Pi0



Dsg

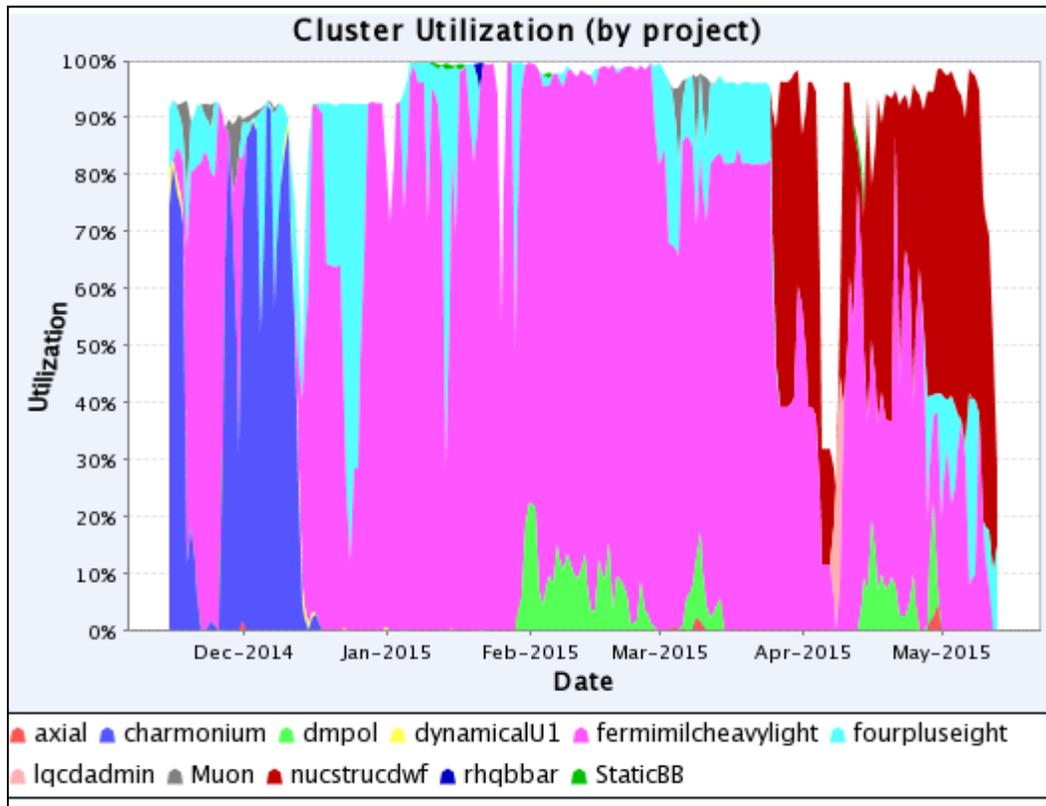


Pi0g



# Cluster Utilization – FNAL

Fermilab Bc Conventional Cluster Utilization



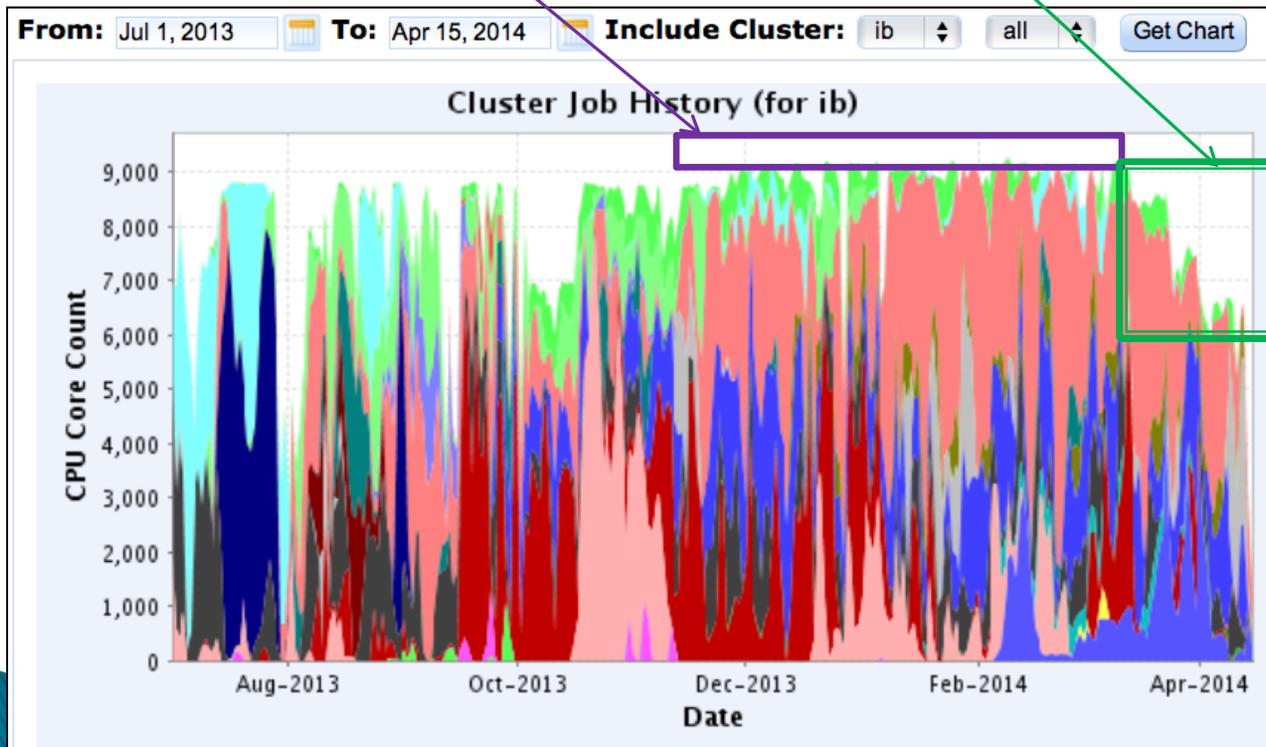
	Conventional Clusters Weighted Utilization	Accelerated Clusters Weighted Utilization
FY13	91%	85%
FY14	94%	81%
FY15 (thru Apr)	90%	78%

Note that on the FNAL GPU clusters, 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

1 million core-hr  
"loan" from physics

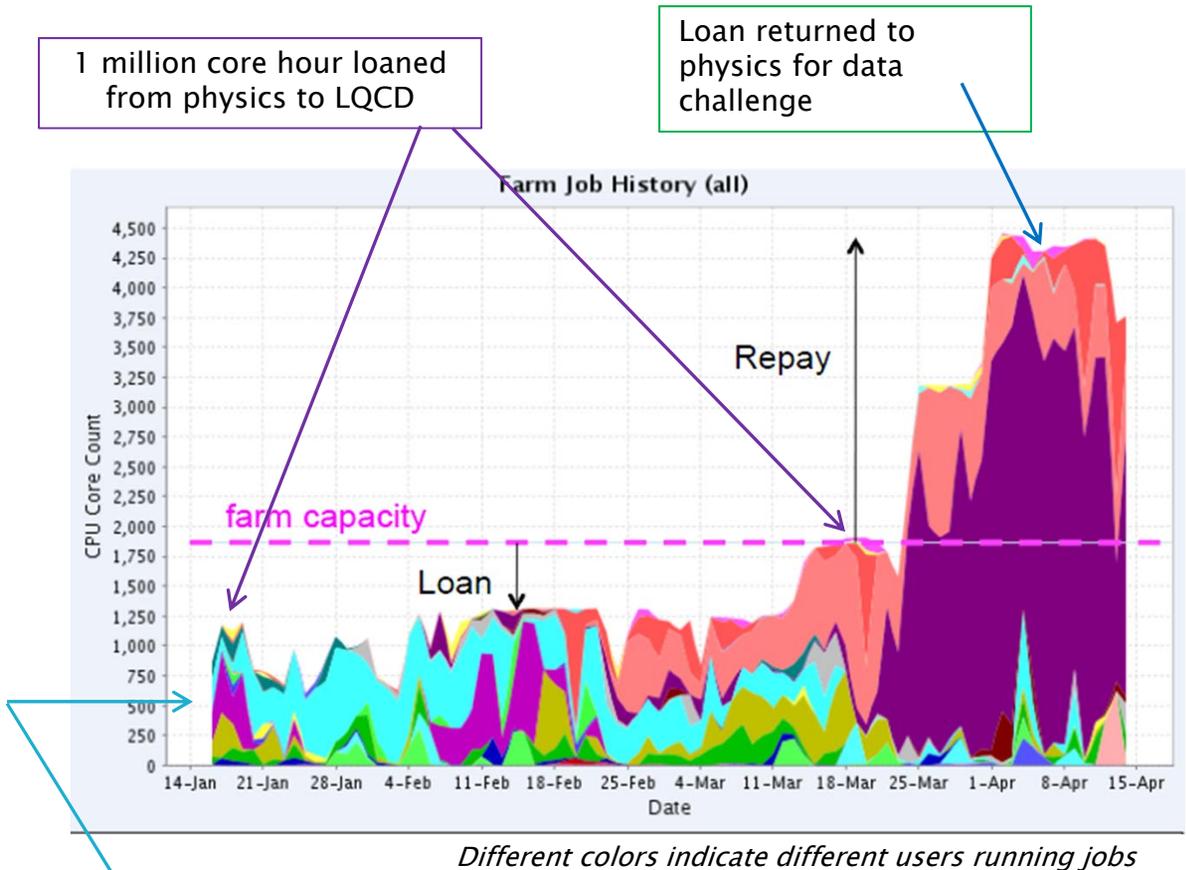
"Loaned core-hrs" returned to  
physics for Glue-X / CLAS-2  
data challenges



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



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

# FY15 Project Cost Summary

Status through **March 2015**; fiscal year complete: **50.0%**

Fund Type	FY14 Carry-over	FY15 Budget	Total FY15 Funds Available	FY15 Actual Costs	FY15 Obligations	% Spent & Obligated
Equipment	\$ 0 K	\$ 0 K	\$ 0 K	\$ 0 K	\$ 0 K	0%
Operating	\$ 809 K	\$ 1,954 K	\$ 2,763 K	\$ 1,408 K	\$ 0 K	51%
Sub-total	\$ 809 K	\$ 1,954 K	\$ 2,763 K	\$ 1,408 K	\$ 0 K	51%
Mgmt Reserve	\$ 0 K	\$ 46 K	\$ 46 K	---	---	0%
<b>TOTAL</b>	\$ 809 K	\$ 2,000 K	\$ 2,809 K	\$ 1,408 K	\$ 0 K	51%

## Cost Performance Analysis

- ▶ Spend rate across the three laboratories is on track, in reasonable agreement with operations budget plan when FY15 OPS costs are compared to FY15 budget:  $(\$1,408\text{ K} - \$451\text{ K}) / \$1,954\text{ K} = 49\% \text{ FY15 spent}$
- ▶ There has been no draw on the management reserve.
- ▶ FNAL Pi0 Expansion purchase obligations are now realized costs in March (\$451 K, from FY14).
- ▶ Forecast for FY15 Carryover = FY15 MR + FY14 unallocated deferred funds = \$46 K + \$127 K = \$173 K.

# FY14 User Survey Results

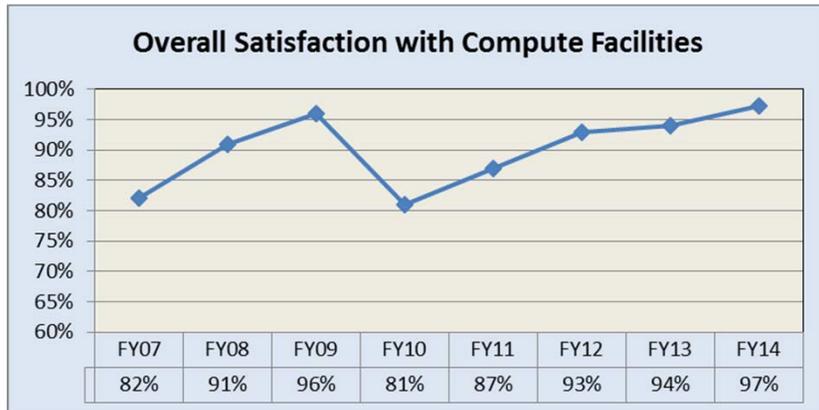
Rob Kennedy

LQCD-ext II Associate Contractor Project Manager

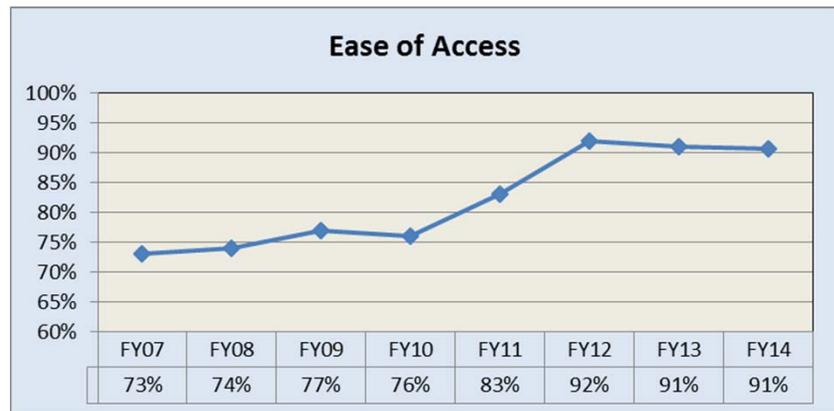
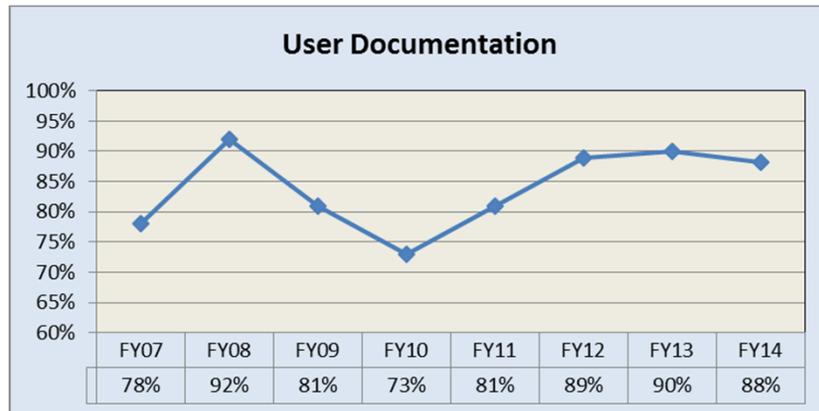
# FY14 Survey Results

- ▶ 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 upward spike in FY13

# Compute Facility Satisfaction Trends

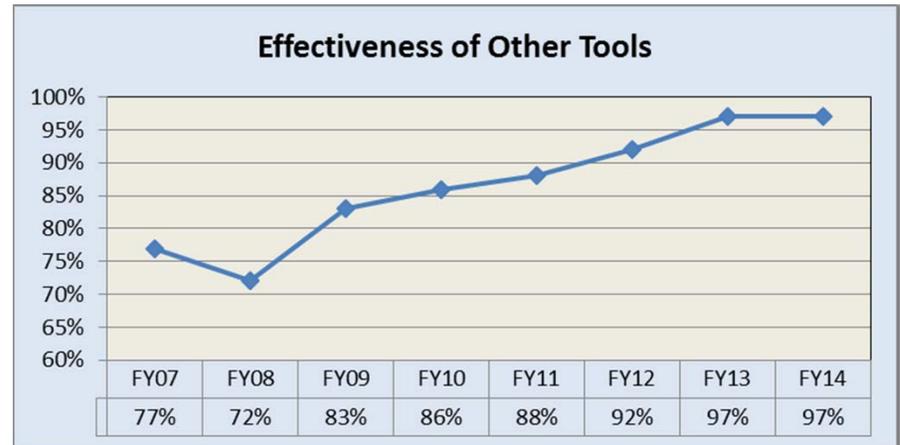
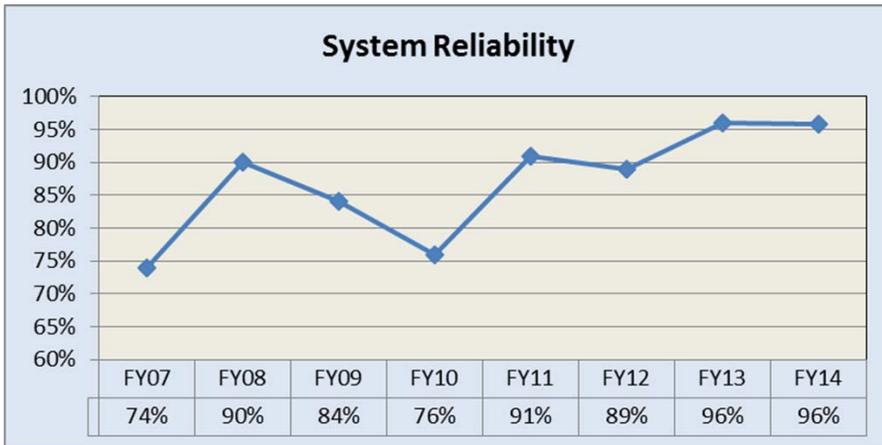
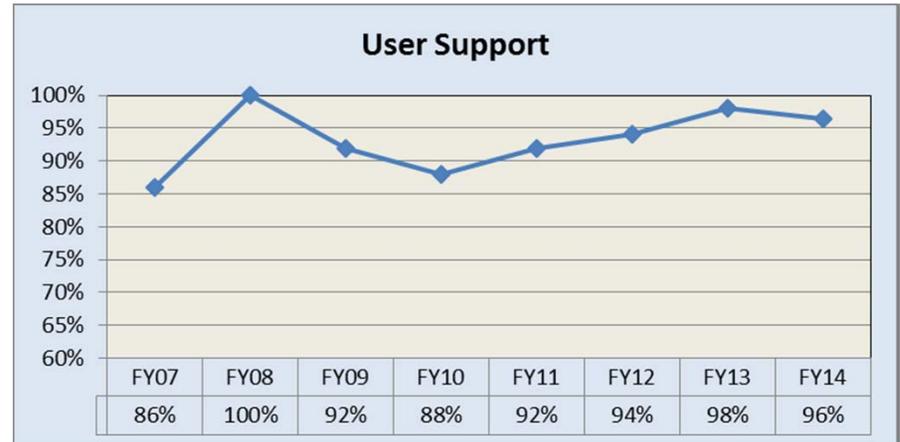
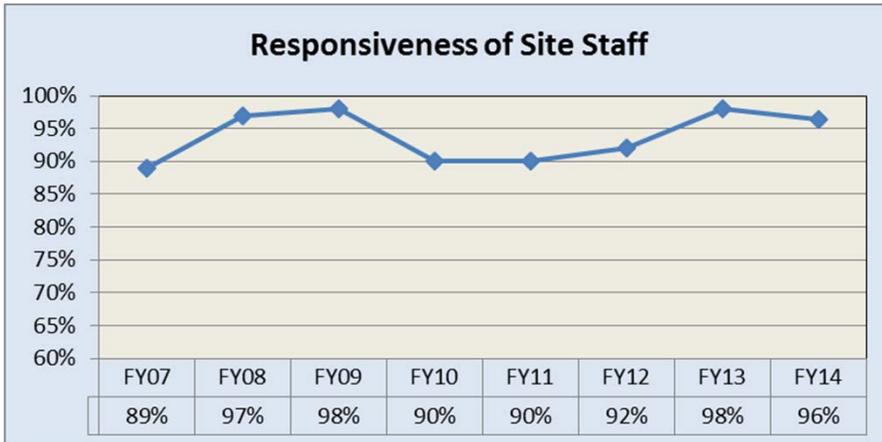


FY14 Computing Facilities	All Sites	BNL	FNAL	JLab
Overall Satisfaction	97%	94%	100%	93%
Documentation	88%	70%	91%	89%
User Support	96%	85%	100%	95%
Responsiveness	96%	93%	100%	90%
Reliability	96%	97%	100%	86%
Ease of Access	91%	86%	96%	82%
Other Tools	97%	91%	100%	94%



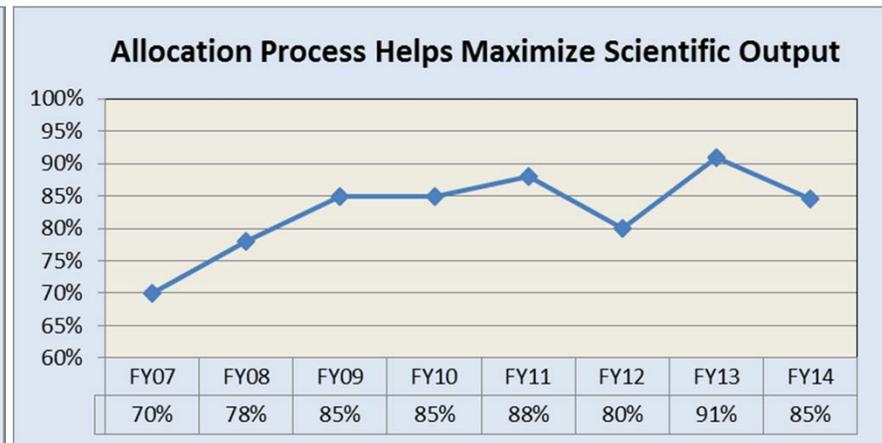
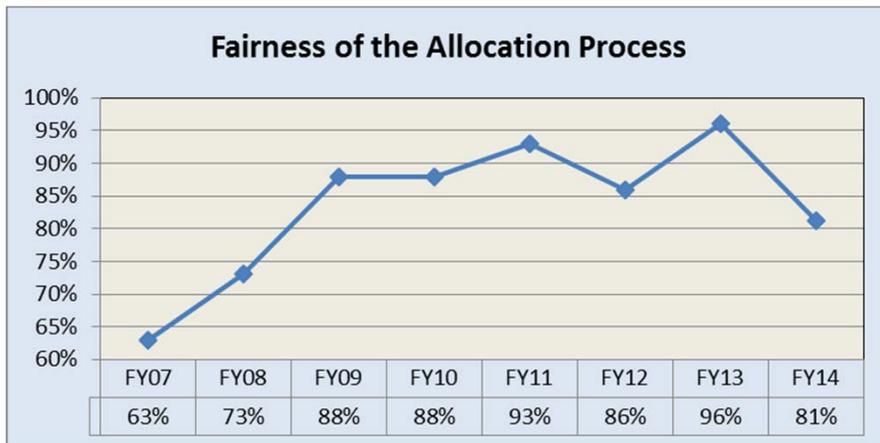
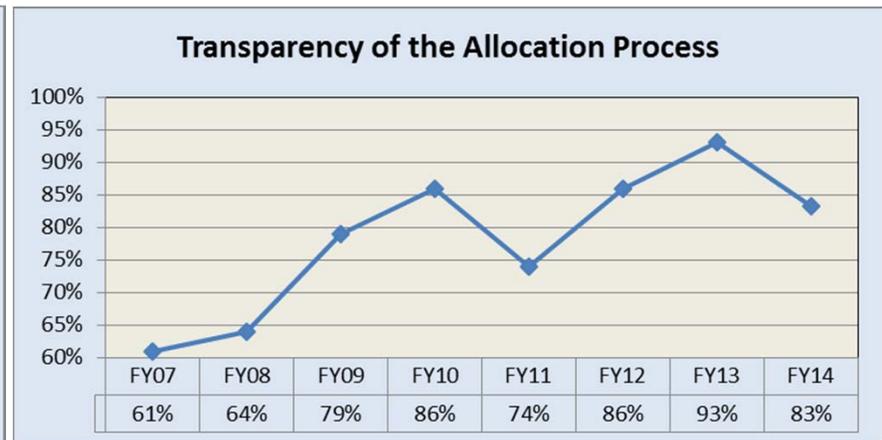
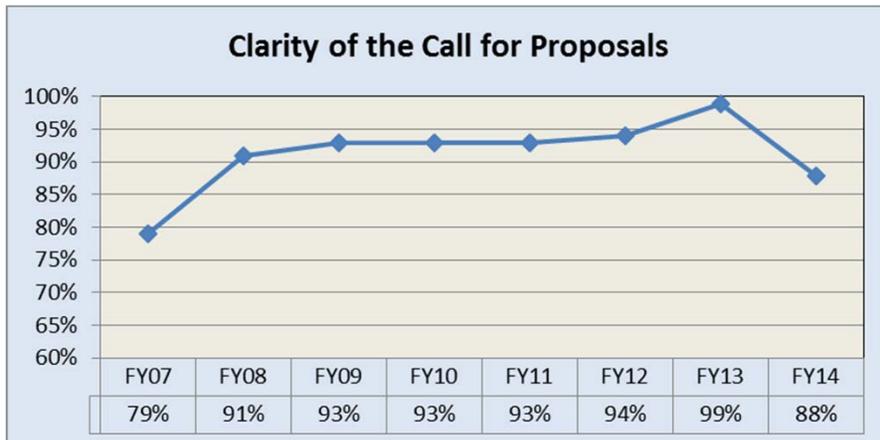
- ▶ The overall satisfaction rating has been trending upward over the previous four years.
- ▶ FY14 rating of 97% exceeds our target goal of 92%, and is similar to recent past.
- ▶ **JLab's Overall Satisfaction rating** of 93% in FY14 continues high level regained in FY13.
- ▶ **BNL's rating for User 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



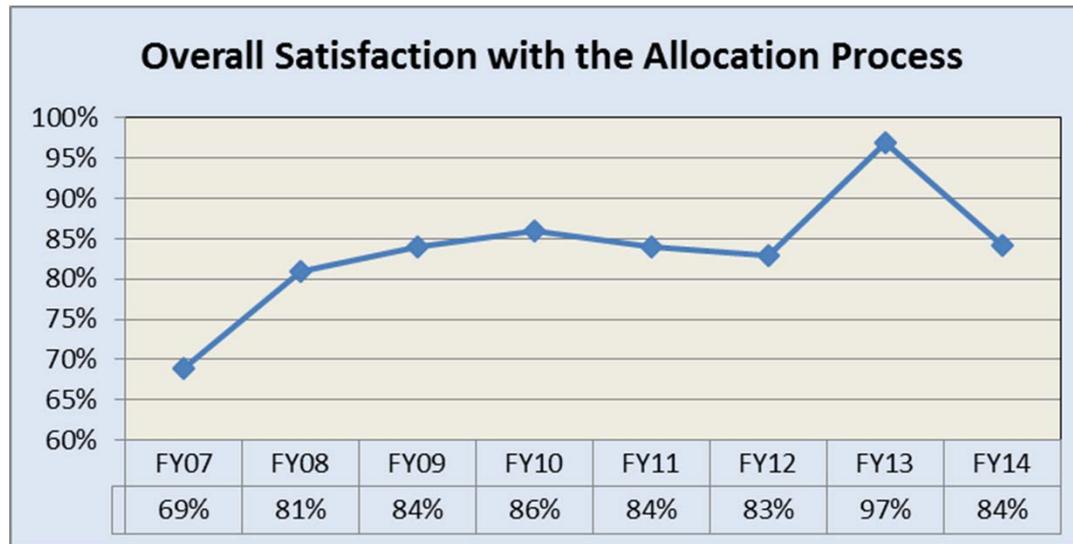
- ▶ Responsiveness of Site Staff, User Support maintain high satisfaction ratings.
- ▶ System Reliability, Online Tools also maintain high satisfaction ratings.

# Allocation Process Satisfaction Trends



- ▶ 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.
- ▶ 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.
- ▶ 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.

# 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.
- Action Plan Highlights:
  - Continue to improve logistics, communication for User Survey process
  - User documentation improved, remains an opportunity for improvement

## ▶ Satisfaction with Allocation Process

- The overall satisfaction rating of 84% is about the same as pre-FY13.
- Action Plan Highlights:
  - Project, Exec Committee, SPC to discuss cost-effective means to reduce confusion expressed by some users due to the wide range of systems in the allocation process... while continuing to focus on goal of optimizing science output per \$ rather than abstract machine performance per \$.

# Summary

- ▶ Compute facilities are running well and we are successfully executing against our plans.
- ▶ We are on target to meet or exceed our FY15 performance goals.
  - 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 will be following a similar approach going forward.