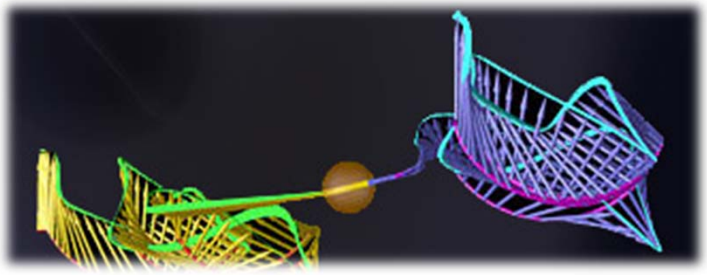

LQCD-ext II Project Management & Performance

Bill Boroski
LQCD Contractor Project Manager

Robert D. Kennedy
Associate Contractor Project Manager

DOE FY2016 Annual Progress Review
Thomas Jefferson National Accelerator Facility
June 28-29, 2016

Outline

- Project scope, organization, and budget
- Proposed changes to the baseline plan
- Performance measures and metrics
- FY15 performance and financial results
- FY16 year-to-date performance and financial 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 some exceptions
  - 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.
  - Addendum approved to Project Execution Plan to assign some Site Manager responsibilities to a newly-defined role, Site Architect.
    - Better captures actual practice, splitting/sharing operations and strategic planning responsibilities between multiple people.
    - Since no new responsibilities were added, this was a minor change approved by Federal Project Director.

Proposed Changes to Project Baseline Plan

- Currently processing two Change Requests (CRs) to formally change the project baseline plan
  - CR16-01: Add cluster-hosting at BNL to the project in a manner that fits within the approved funding profile and does not reduce the total computing delivered by the project.
  - CR16-02: Replace the current set of separate performance goals for Conventional and Accelerated computing resources with goals that do not distinguish among hardware architectures.
- You will hear more about these in later sessions.

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.

Project Change Control Process

- Defined in PEP
- Requirements based on Change Control Level.
- Trigger any threshold, then approvals are required up to and including that level.
- Communicate Level 1+ CR to USQCD: evaluate Scientific Impact
- Changes are also communicated to the Project Director and Project Monitor.

Change Control Level	Approver	Cost Threshold	Schedule Threshold	Technical Scope/ Performance Threshold
Level 4	Acquisition Executive	Any increase in Total Project Cost Or Change of > $250K in budget distribution between DME and SS O&M costs	6-month or more increase in a Level 1 milestone date	Changes to scope that affect mission need and/or performance requirements
Level 3	Federal Project Director	Change of ≥ $125K in budget distribution between DME and SS O&M costs or Movement of allocated funds between laboratories	3-month or more delay of a Level 1 milestone date	Any modification in the technical performance baseline defined in a Level-1 milestone
Level 2	Change Control Board	Change of < $125K in budget distribution between DME and SS O&M costs or Cumulative increase of ≥ $125K over baseline budget for WBS Level 2 elements	> 1-month delay of a Level 1 milestone date or > 3-month delay of a Level 2 milestone date.	> 10% decrease from baseline of either the targeted computing capability increment (Tflop/s) or integrated delivery (Tflop/s-yrs) in a single project year.
Level 1	Contractor Project Manager	Any increase of ≥$25K over baseline budget for WBS Level 2 elements	> 1-month delay of a Level 2 milestone date	Any deviation from technical deliverables that negatively affects expected performance specifications by more than 5%

Project Change Control Process

- Defined in PEP
- Requirements based on Change Control Level.
- Trigger any threshold, then approvals are required up to and including that level.
- Communicate Level 1+ CR to USQCD: evaluate Scientific Impact
- Changes are also communicated to the Project Director and Project Monitor.

CR16-01

CR16-02

Change Control Level	Approver	Cost Threshold	Schedule Threshold	Technical Scope/ Performance Threshold
Level 4	Acquisition Executive	Any increase in Total Project Cost Or Change of > $250K in budget distribution between DME and SS O&M costs	6-month or more increase in a Level 1 milestone date	Changes to scope that affect mission need and/or performance requirements
Level 3	Federal Project Director	Change of ≥ $125K in budget distribution between DME and SS O&M costs or Movement of allocated funds between laboratories	3-month or more delay of a Level 1 milestone date	Any modification in the technical performance baseline defined in a Level-1 milestone
Level 2	Change Control Board	Change of < $125K in budget distribution between DME and SS O&M costs or Cumulative increase of ≥ $125K over baseline budget for WBS Level 2 elements	> 1-month delay of a Level 1 milestone date or > 3-month delay of a Level 2 milestone date.	> 10% decrease from baseline of either the targeted computing capability increment (Tflop/s) or integrated delivery (Tflop/s-yrs) in a single project year.
Level 1	Contractor Project Manager	Any increase of ≥$25K over baseline budget for WBS Level 2 elements	> 1-month delay of a Level 2 milestone date	Any deviation from technical deliverables that negatively affects expected performance specifications by more than 5%

Management & Oversight
Changes in our LQCD-ext II Project Team

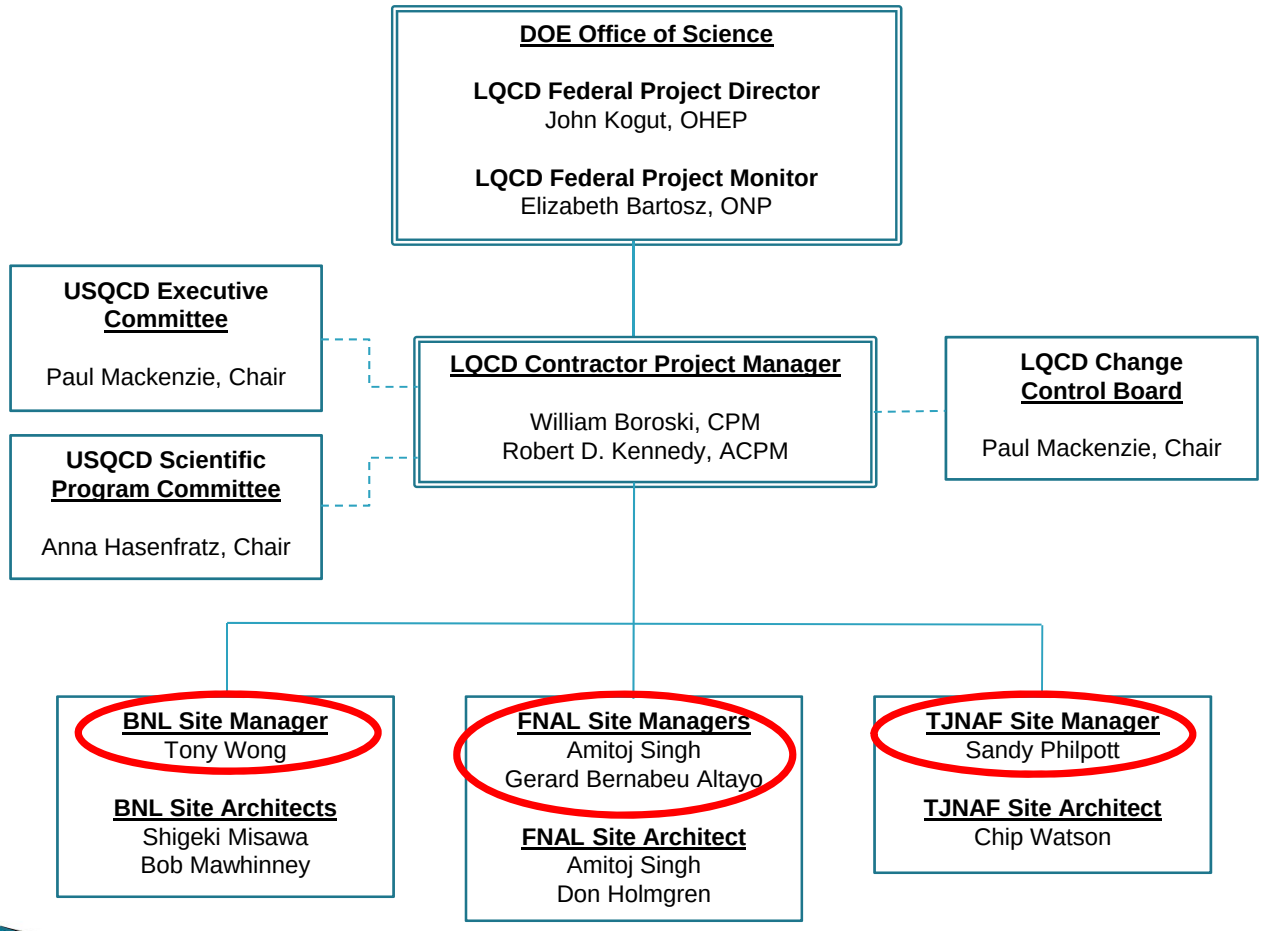

Organizational changes:

Tony Wong has replaced Frank Quarant as BNL Site Manager.

Sandy Philpott has replaced Chip Watson as JLab Site Manager.

Gerard Bernabeu Altayo has replaced Don Holmgren as co-Site Manager at FNAL

We have introduced the role of Site Architects into our org structure.

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
Total	2,000	3,000	3,000	3,000	3,000	14,000
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
Total	2,000	3,000	3,000	3,000	3,000	14,000
Planning Budget Guidance	2,000	3,000	3,000	3,000	3,000	14,000

Equipment budget is used to procure compute and storage hardware

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
Total	4,242	369	4,611

Storage budgeted at 8% of total hardware budget, based on historical experience.

Adjusted as needed at the few percent level as part of the annual planning process, in order to meet USQCD needs.

Performance Measures & Metrics

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.

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.

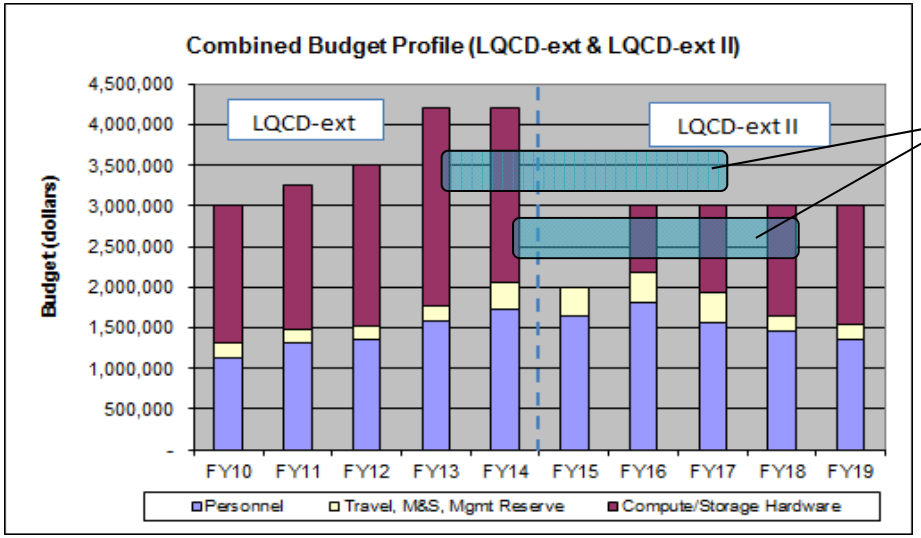

Indicates 4-yr system lifecycle.

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

Performance To Date

LQCD-ext II Project Status

- We're in the second year of the 5-year extension (funded from Oct 2014 thru Sep 2019)
- We've received $5M of our planned $14M in funding, in accordance with our baseline funding profile ($2M in FY15; $3M in FY16).
- The computing we've delivered to USQCD through May 2016 has exceeded our baseline goals!

	FY15			FY16 [1]		
	Goal	Actual	% of Goal	Goal	Actual	% of Goal
Conventional Resources[2]	95.1	105.7	111%	63.9	68.8	108%
GPU-accelerated Resources[3]	142.8	144.1	101%	53.8	56.9	106%

1) FY16 performance through May 2016.
2) Conventional resources: Ds, Bc, Pi0, 9q, 10q, 12s, BG/Q, 10% of DD2 prototype BG/Q rack
3) GPU-accelerated resources: Dsg, Pi0g, 10g, 11g, 12k (9g retired Jun 27, 2015).

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

- Goal = 95.1 TFlops-yrs; Actual = 105.7 TFlops-yrs (111% of goal)
  - The uptime goal is 8000 hours per year (91.3%), which equates to a conventional hardware goal for FY15 of 95.1 TFlops-yrs.

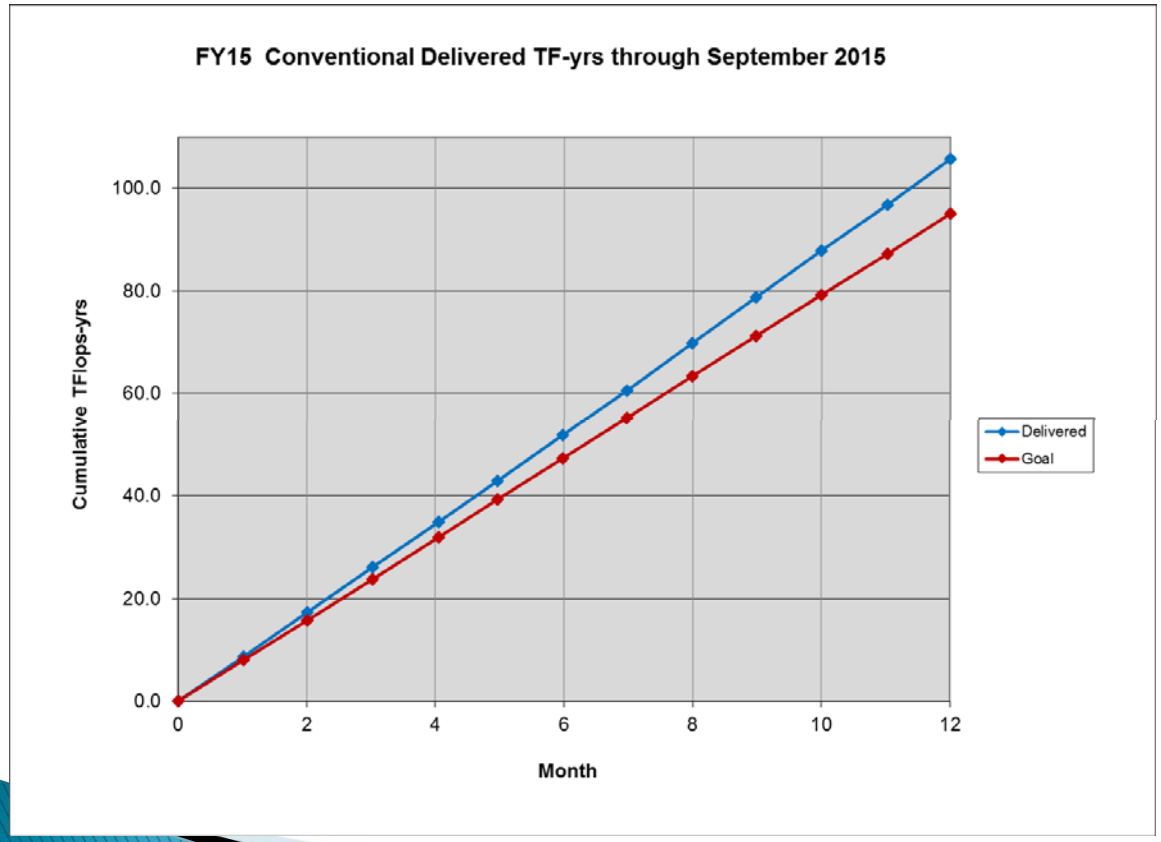

Computing resources included are the FNAL and JLab Infiniband clusters, the BNL BG/Q DD2 rack, and 10% of the BNL BG/Q DD2 prototype rack.

All sites above goal due to high uptime and smooth operations on all clusters.

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

- Goal = 142.8 effective TF-yrs; actual = 144.1 effective TF-yrs (101% of goal)
  - 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.

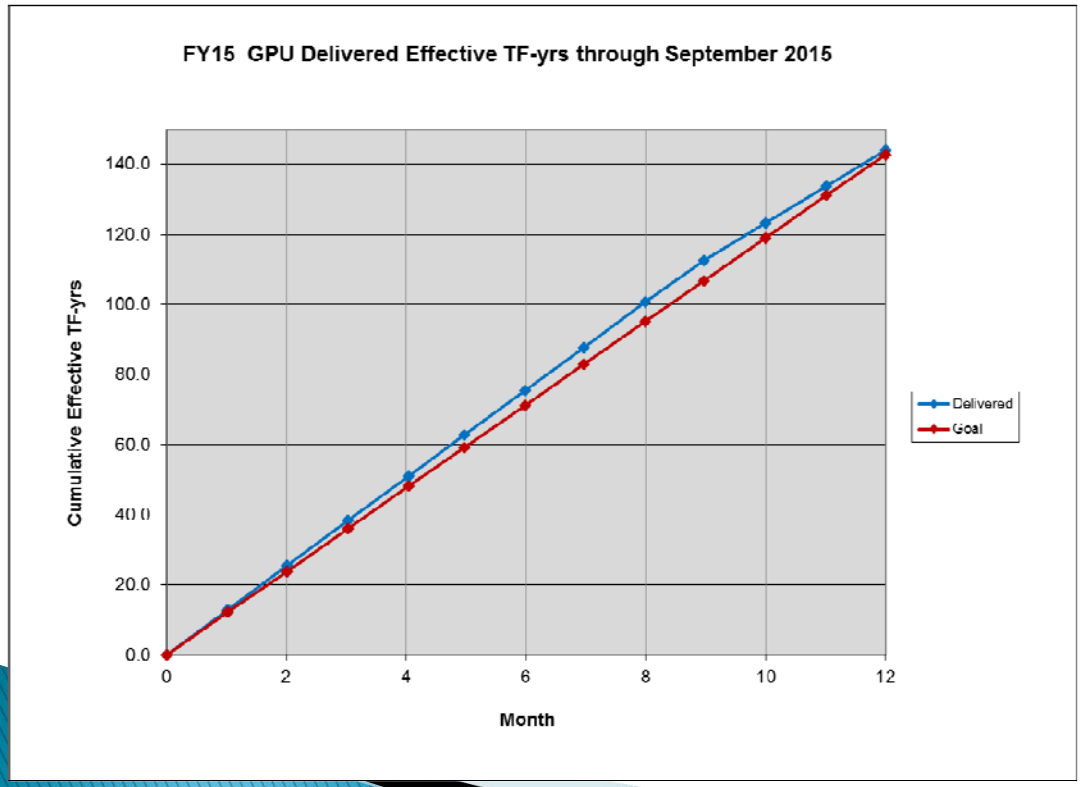

Computing resources included are the FNAL Dsg and Pi0g clusters, and the JLab 9g, 10g, 11g, and 12k clusters.

All sites exceeded goals due to smooth operations and high uptime on all clusters.

- JLab 9g retired June 27, 2015
- FNAL Dsg out of warranty, suffering GPU failures.
  - Evaluated which 1 of 2 GPUs on these nodes failed, as time permitted.

BG/Q Utilization – BNL

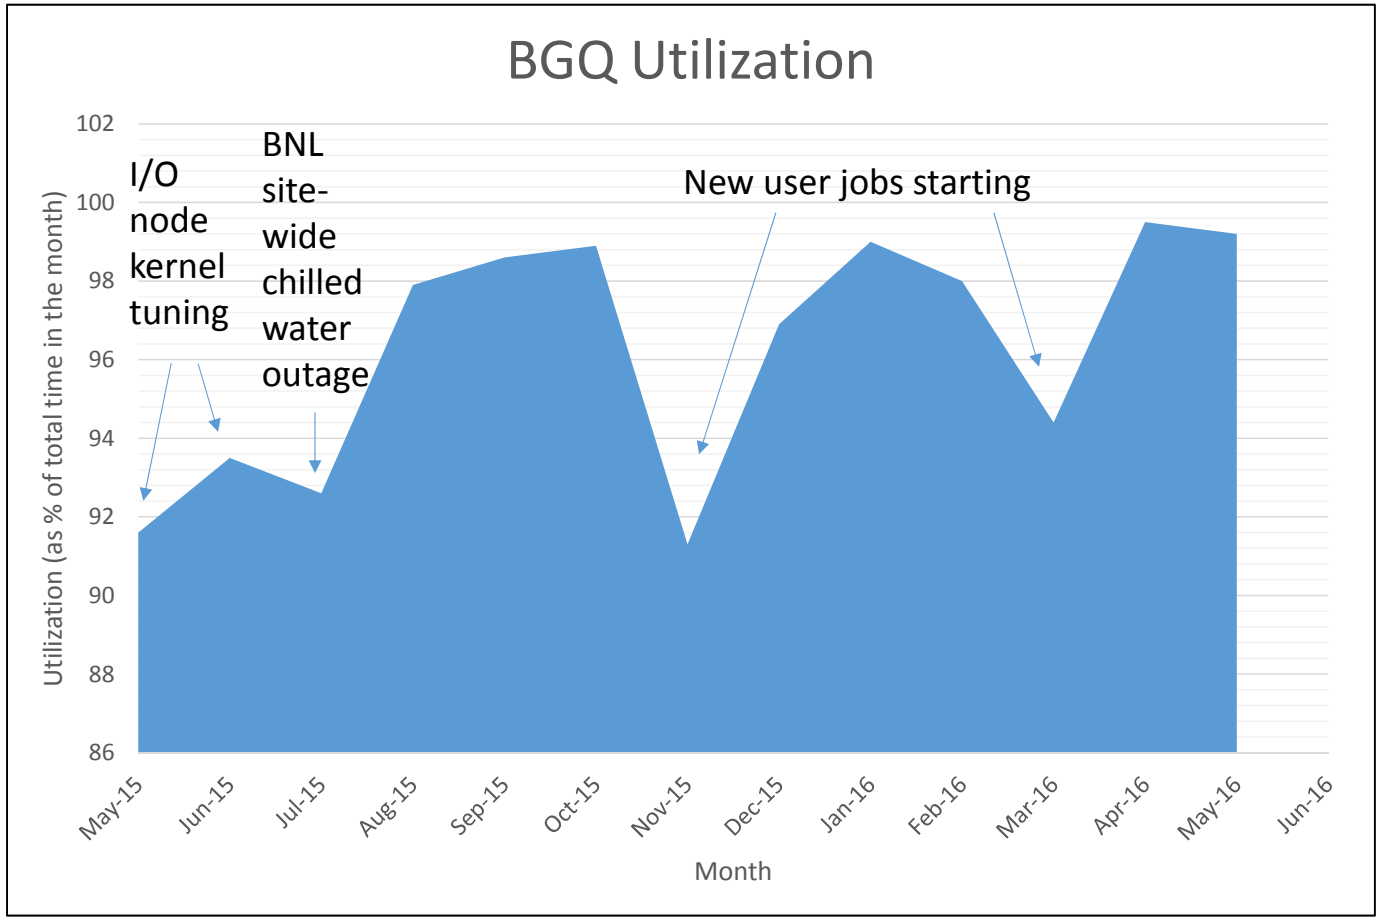

Uptime on the BG/Q is based on percentage of total possible uptime every month.

Uptime is not based on an 8000 hour year, but rather a 365*24=8760 hour year.

FY16 Cluster Utilization – FNAL

Ds

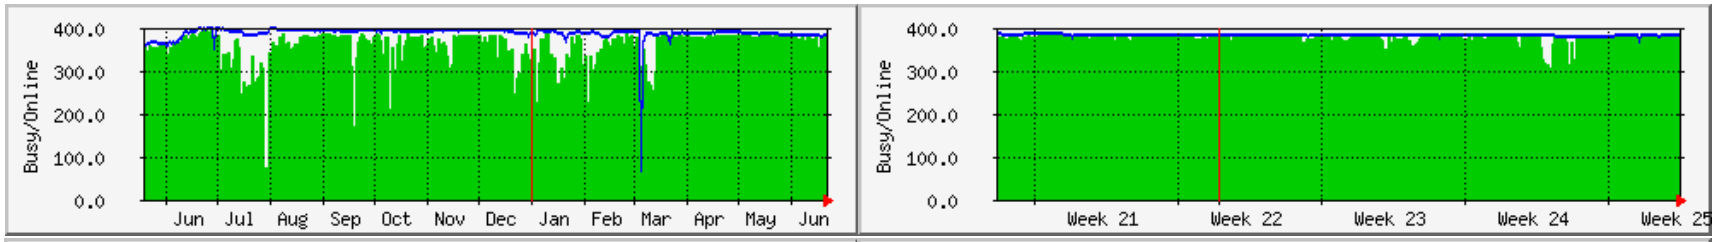

Bc

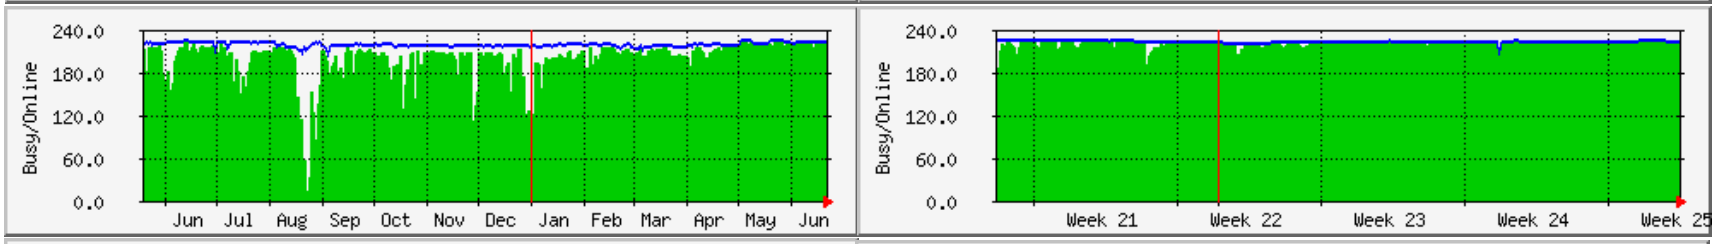

Pi0

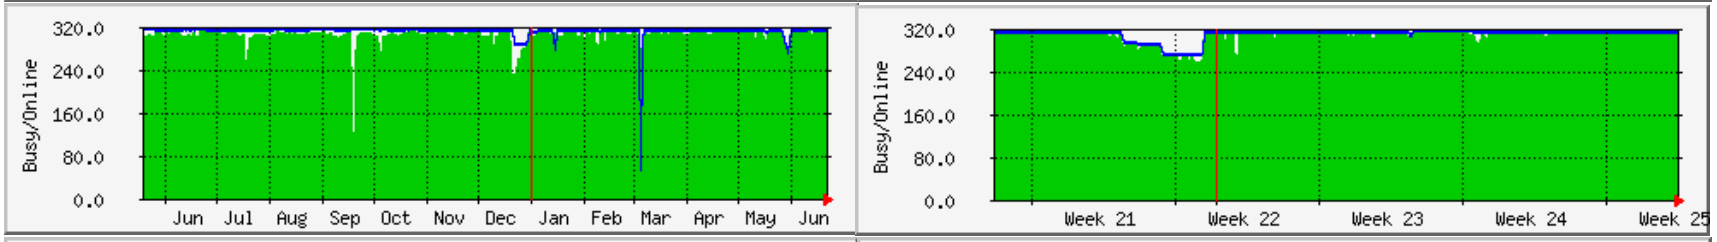

Dsg

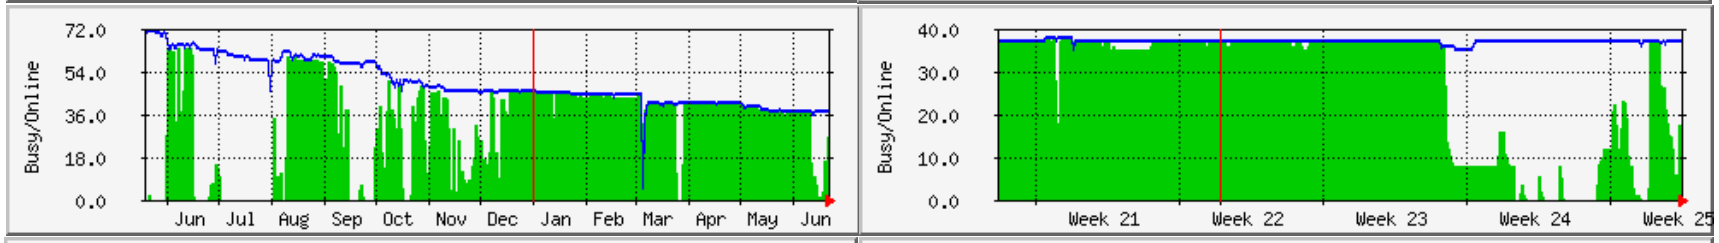

Pi0g

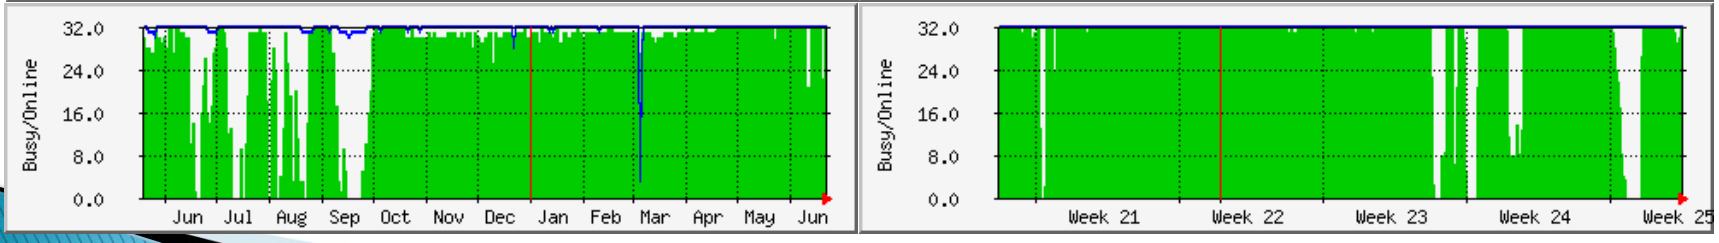

Cluster Utilization – FNAL

Fermilab Bc Conventional Cluster Utilization

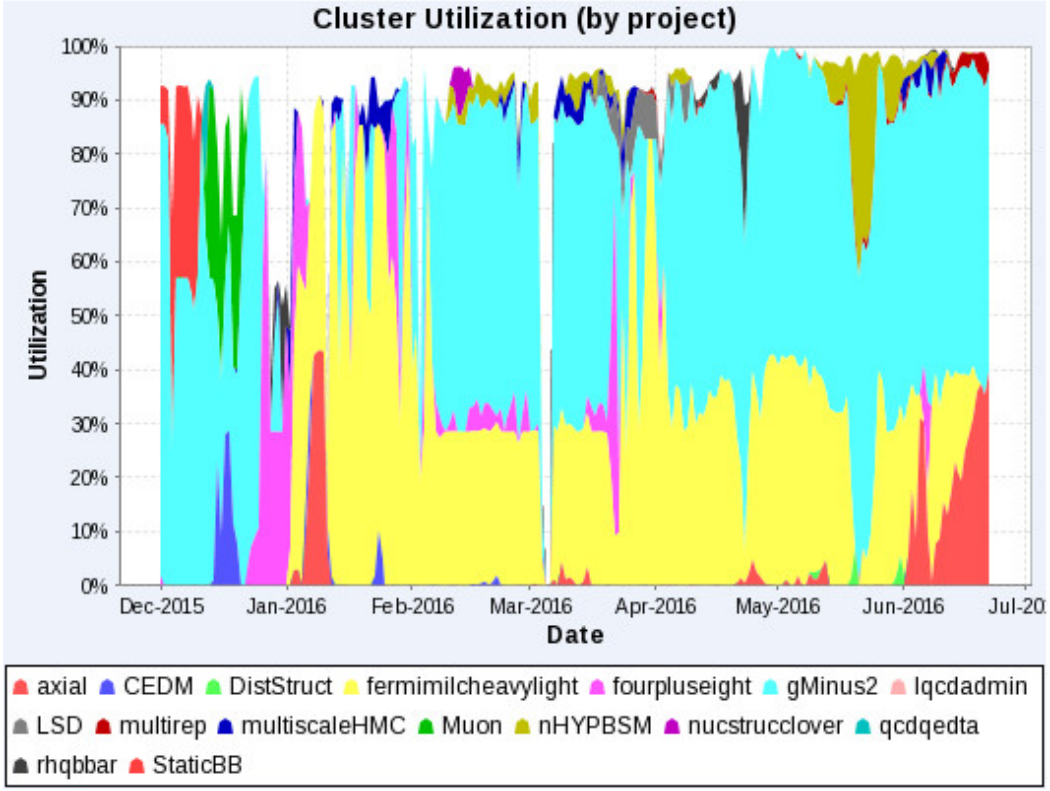

The FNAL GPU clusters had low to sporadic usage during the early months of the 2015-16 allocation period due to on going work on GPU code optimizations.

Cluster Utilization – JLab

Resource sharing between LQCD and Experimental Physics continues

- An early loan last summer from Physics, paybacks to GlueX, HPS, PRad, QWeak this spring
- Flexibility is key; debt ranged from 3.5m core hours last summer, up to 5m, to 2.1m now as projects ran
- Reduced idle time for both LQCD and Experimental Physics, as projects cycle through their periods of computing

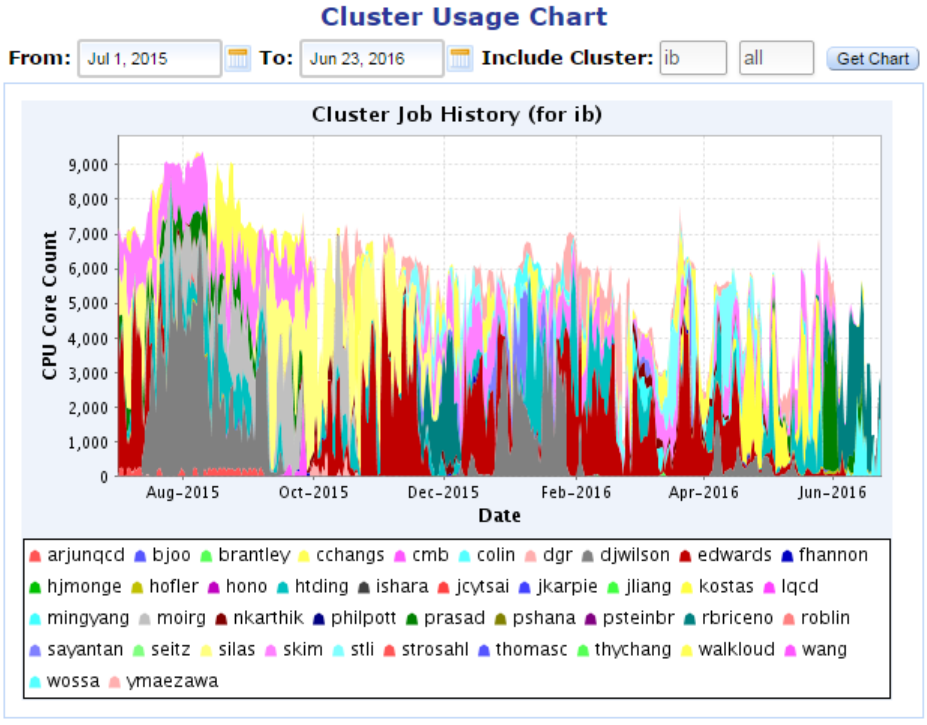

Note: In this chart, all cores are treated identically; differences in performance are ignored.

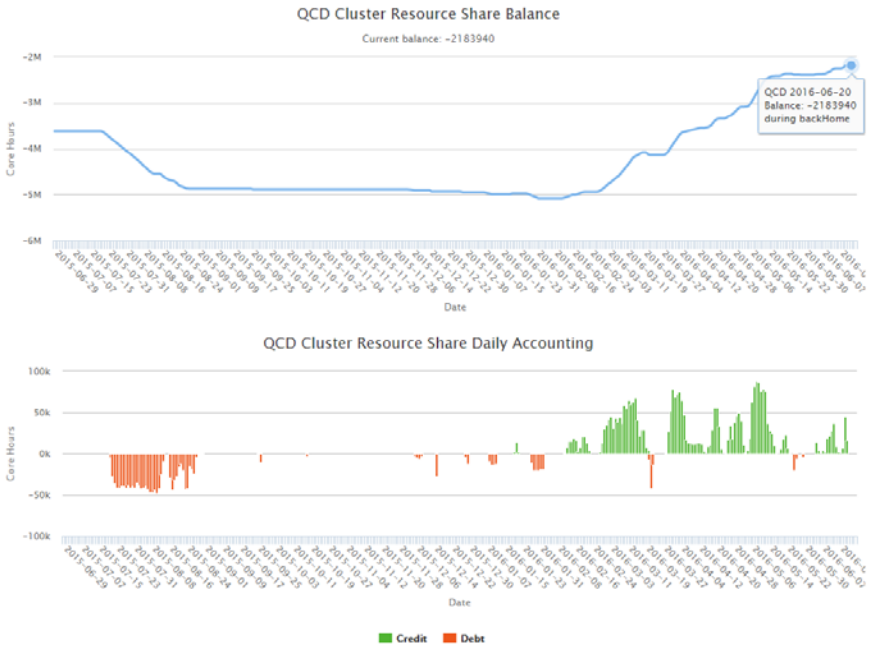

FY15 Financial Performance

FY15 Project Cost Summary

Status through September 2015; fiscal year complete: 100%

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	$ 2,442 K	$ 7 K	89%
Sub-total	$ 809 K	$ 1,954 K	$ 2,763 K	$ 2,442 K	$ 7 K	89%
Mgmt Reserve	$ 0 K	$ 46 K	$ 46 K	---	---	0%
TOTAL	$ 809 K	$ 2,000 K	$ 2,809 K	$ 2,442 K	$ 7 K	89%

Cost Performance Analysis

- FNAL Pi0 Expansion purchase cost was $451 K (from FY14).
- Spend rate across the three laboratories was on track with operations budget plan when FY15 OPS costs are compared to FY15 budget: ($2,442 K – $451 K)/$1954 K = 102% FY15 budget spent
- There was no draw on the management reserve.

Part 2: User Survey, FY16 to Date

FY15 User Survey Results

FY15 Survey Results

- The FY15 User Survey measured user satisfaction during the 12 month period from Oct 2014 through Sep 2015 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 all individuals "known" to USQCD
  - Responses were received from 66 individuals
    - By comparison, 61 individuals responded to the FY14 survey
  - 30 of 35 PI's responded: 86% response rate (74% in FY14)
  - 32 of top 64 Active Users responded: 50% response rate (50% in FY14)
- FY14 overall satisfaction rating with Compute Facilities = 97%
  - Exceeds our KPI goal of 92%
- FY14 overall satisfaction with Resource Allocation Process = 91%
  - Improvement, was 84% in FY14

Compute Facility Satisfaction Trends

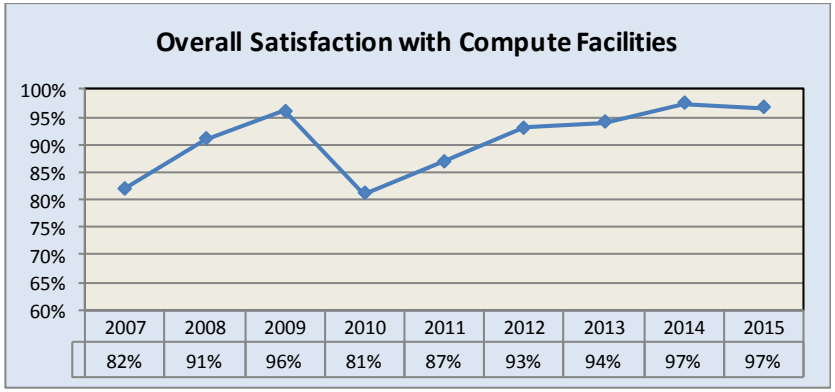

Project KPI

FY15 Computing Facilities	All Sites	BNL	FNAL	JLab
Overall Satisfaction	**97%**	89%	100%	92%
Documentation	93%	83%	96%	94%
User Support	99%	100%	99%	100%
Responsiveness	99%	100%	99%	100%
Reliability	93%	100%	94%	89%
Ease of Access	93%	100%	95%	88%
Other Tools	95%	100%	93%	97%

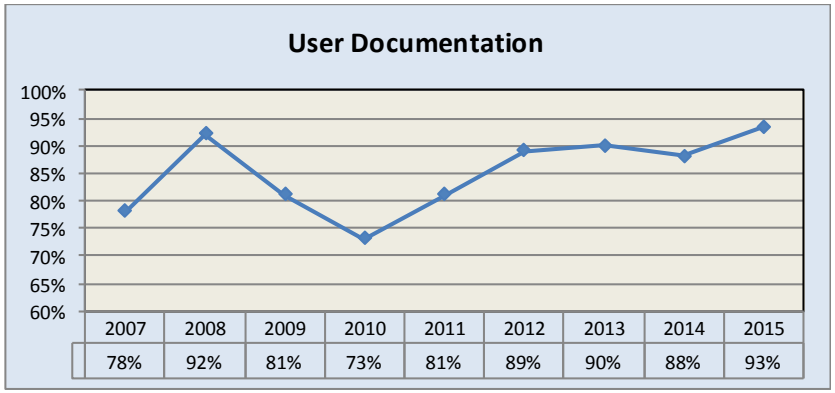

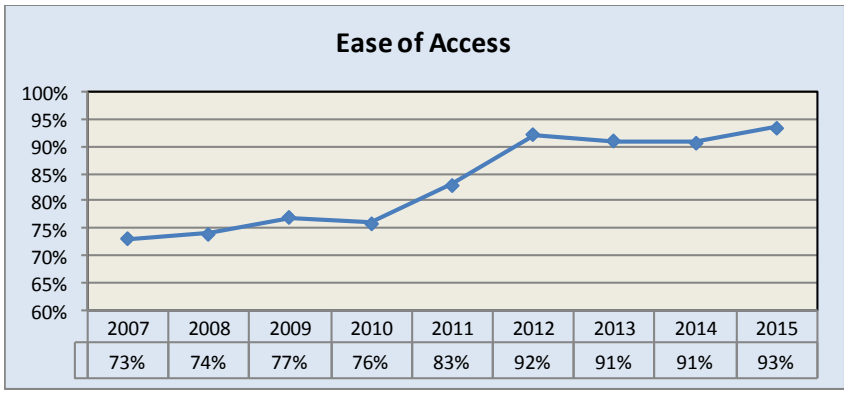

- FY15 Overall Satisfaction rating: 97% exceeds project goal of 92%, similar to recent past.
- *Yellow*: BNL's rating for User Documentation was still below par (88%), but improving.
  - Action Plan defined to improve BG/Q documentation handling and prepare for possible cluster-oriented documentation
- *Green*: Areas of improvement from mid-80% in FY14 to 100% in FY15

Allocation Process Satisfaction Trends

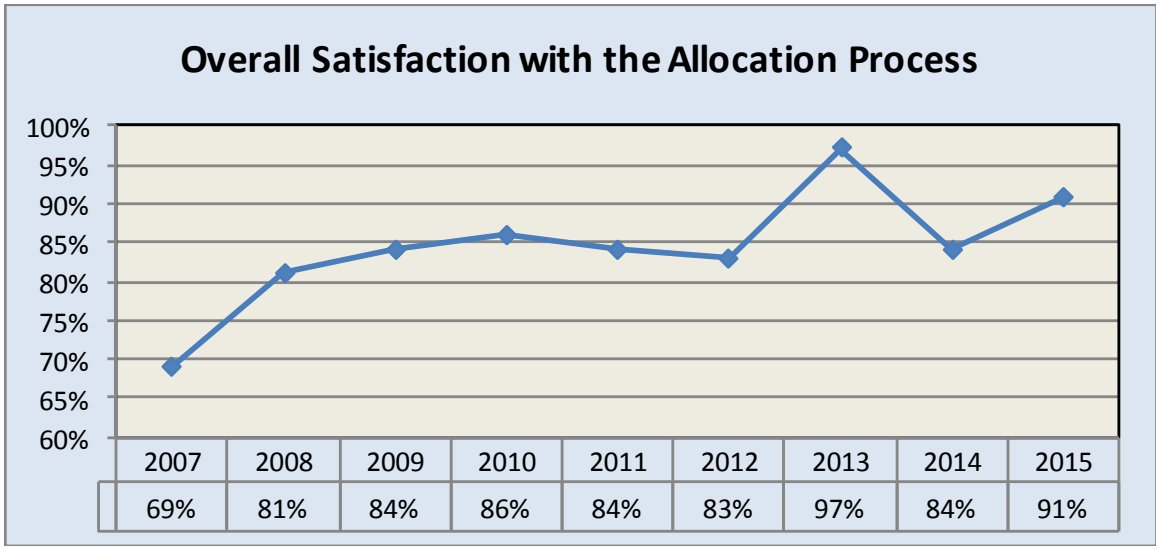

- The overall satisfaction rating for the Allocation Process was 91% in FY14.
  - This is a noticeable improvement over the 84% rating in FY14.
- Related user feedback included:
  - Acknowledgement of the challenges of allocating over-subscribed resources
  - Concern about some allocations not being used for a large part of the year while proposals that had been turned down were ready to run
    - EC instituted policy in 2016, similar to NERSC: reduces future allocations if allocations go unused in a quarter.
  - Concern about the EC and SPC having no elected members
    - USQCD elected a new member (Will Detmold) to EC in late April 2016.
  - Suggestions to streamline or improve the allocation process

FY15 User Survey Results – More Detail

Compute Facility Satisfaction Trends

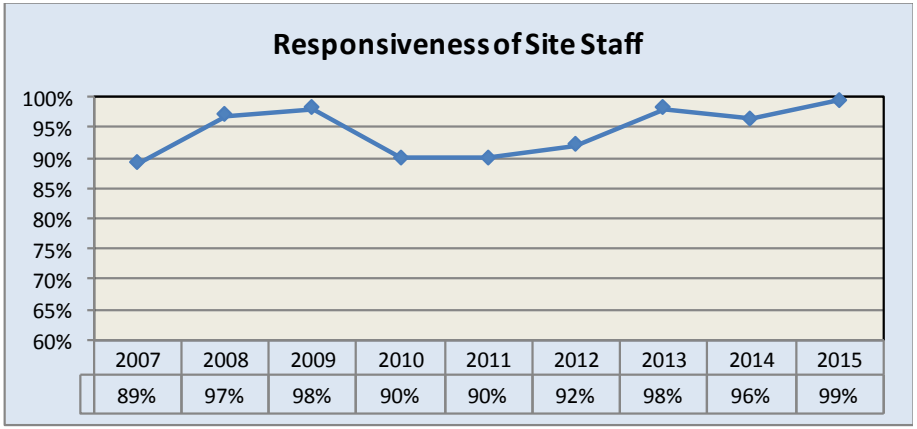

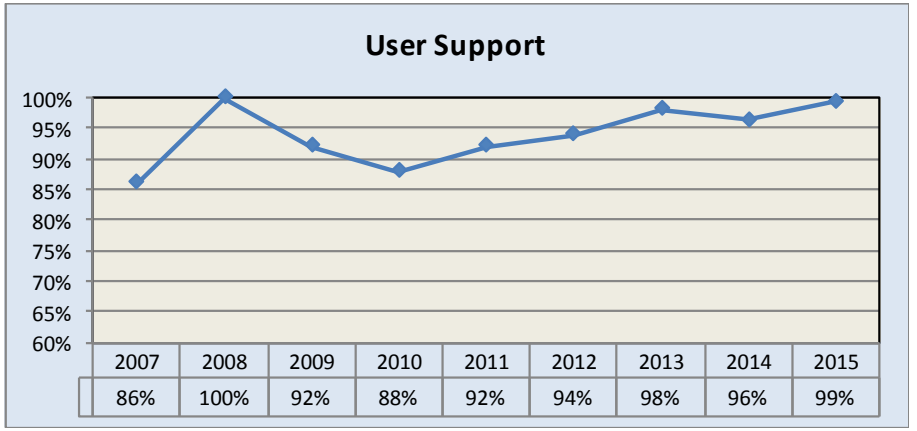

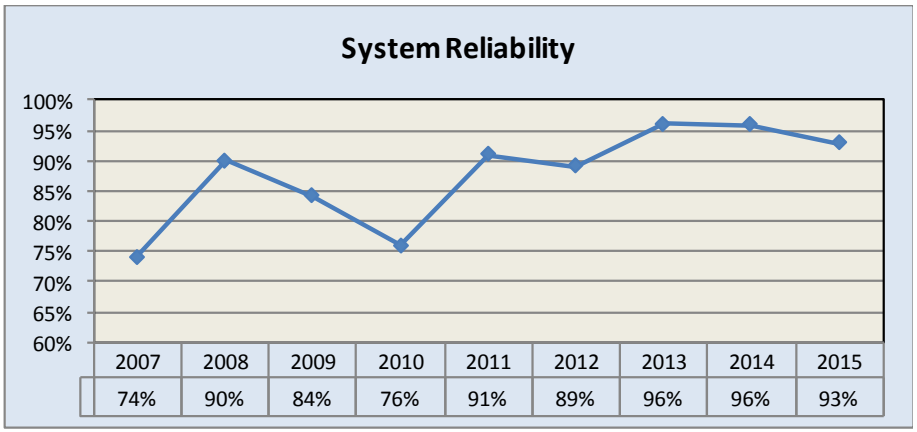

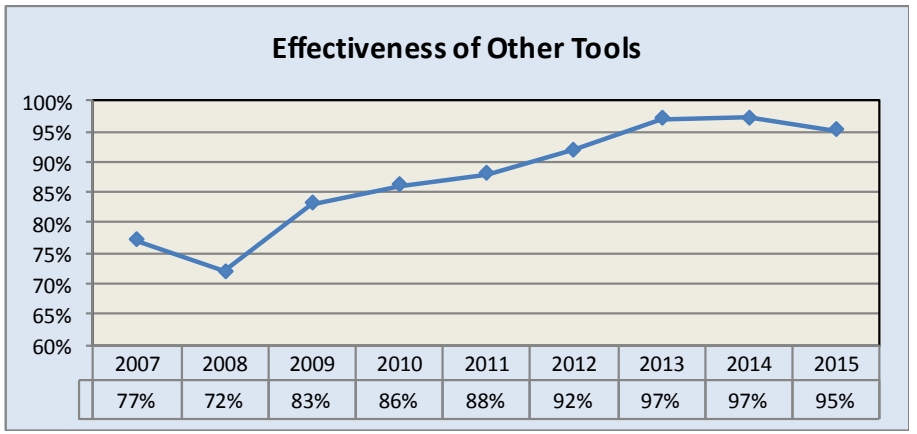

- Responsiveness of Site Staff and User Support maintain high satisfaction ratings.
- System Reliability and Online Tools also continue to maintain high satisfaction ratings.
  - Some systems are aging, past warranty, but still in use, which may explain slight downtick in System Reliability

Allocation Process Satisfaction Trends

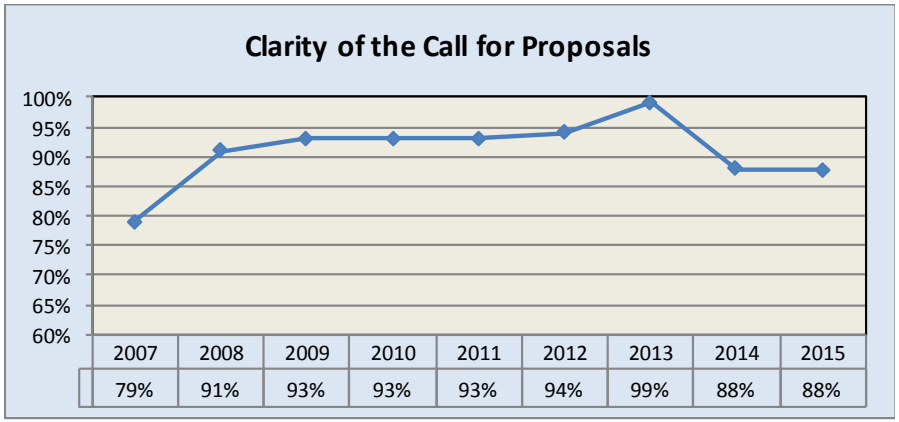

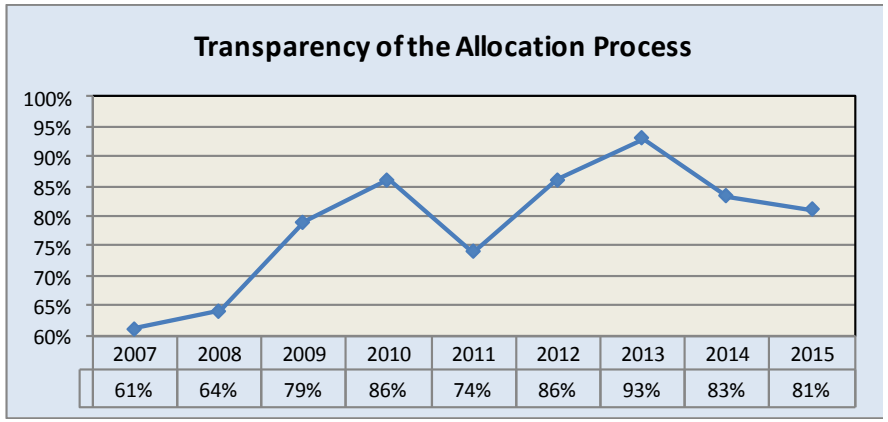

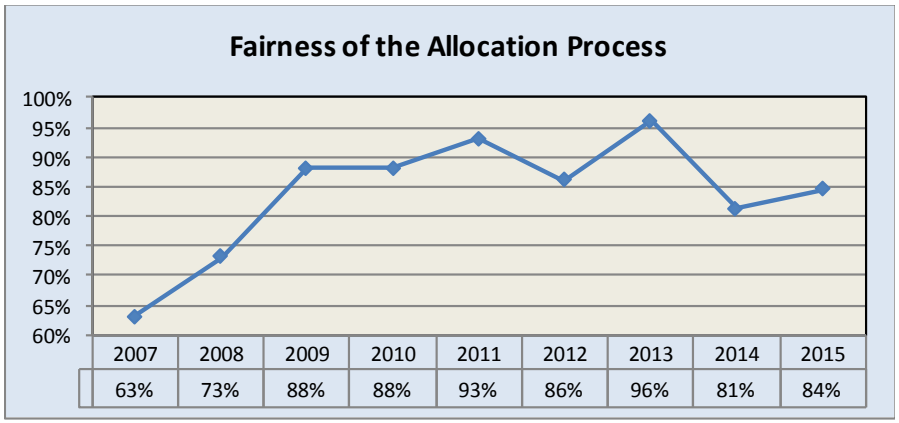

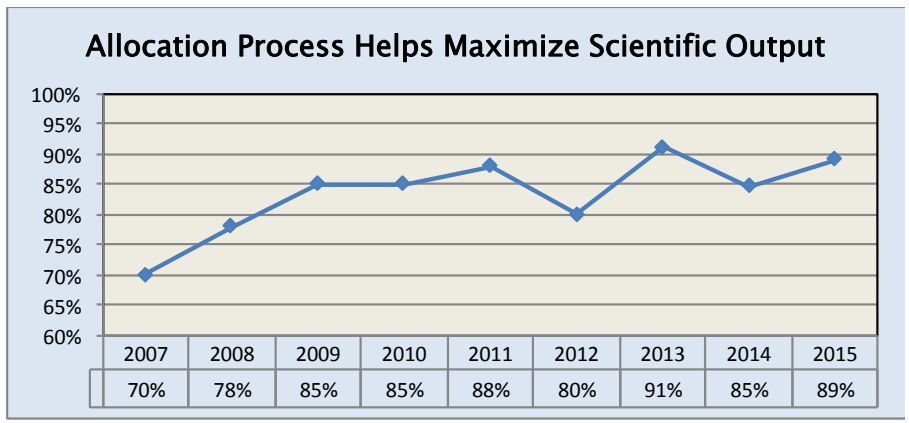

- Clarity and Transparency ratings remain near FY14 levels.
- Fairness and Maximize Scientific Output ratings rose a bit from FY14 levels.

FY16 Performance Results to Date

FY16 Conventional Resource Performance *(TFlops-yrs delivered)*

Computing resources included are FNAL and JLab Infiniband clusters, BNL BG/Q LQCD half-rack and10% of BG/Q DD2 prototype rack. The project is on-target to exceed both baseline KPI and forecast Goal.

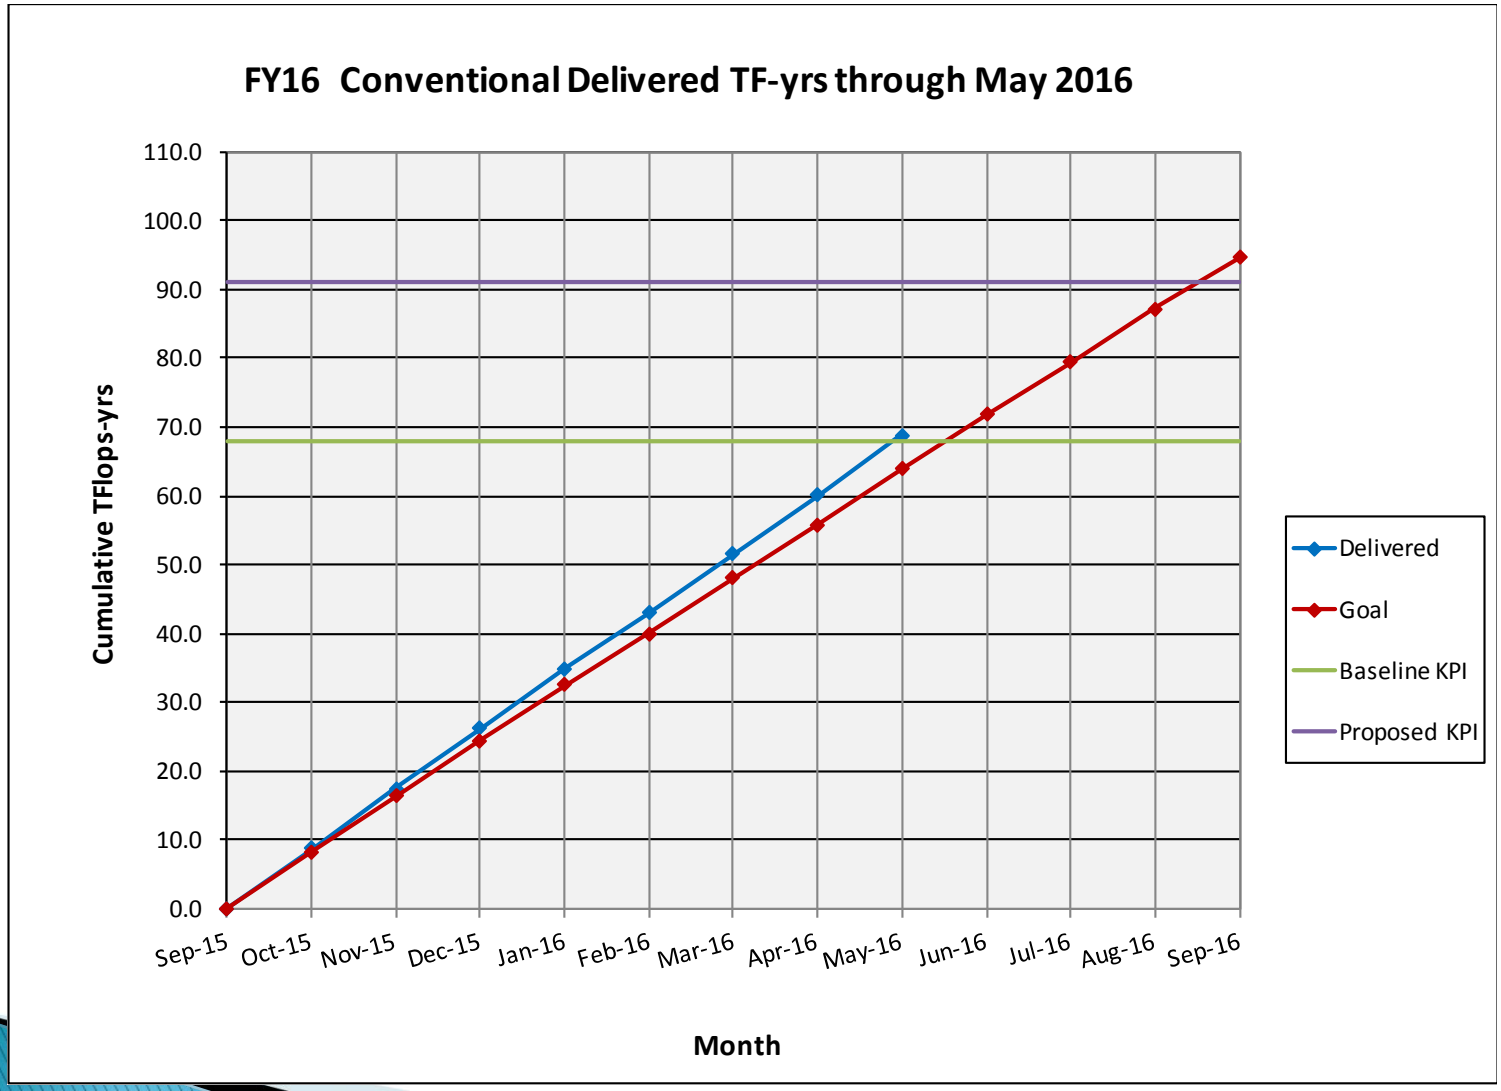

FY16 Conventional Resource Performance *(TFlops-yrs delivered)*

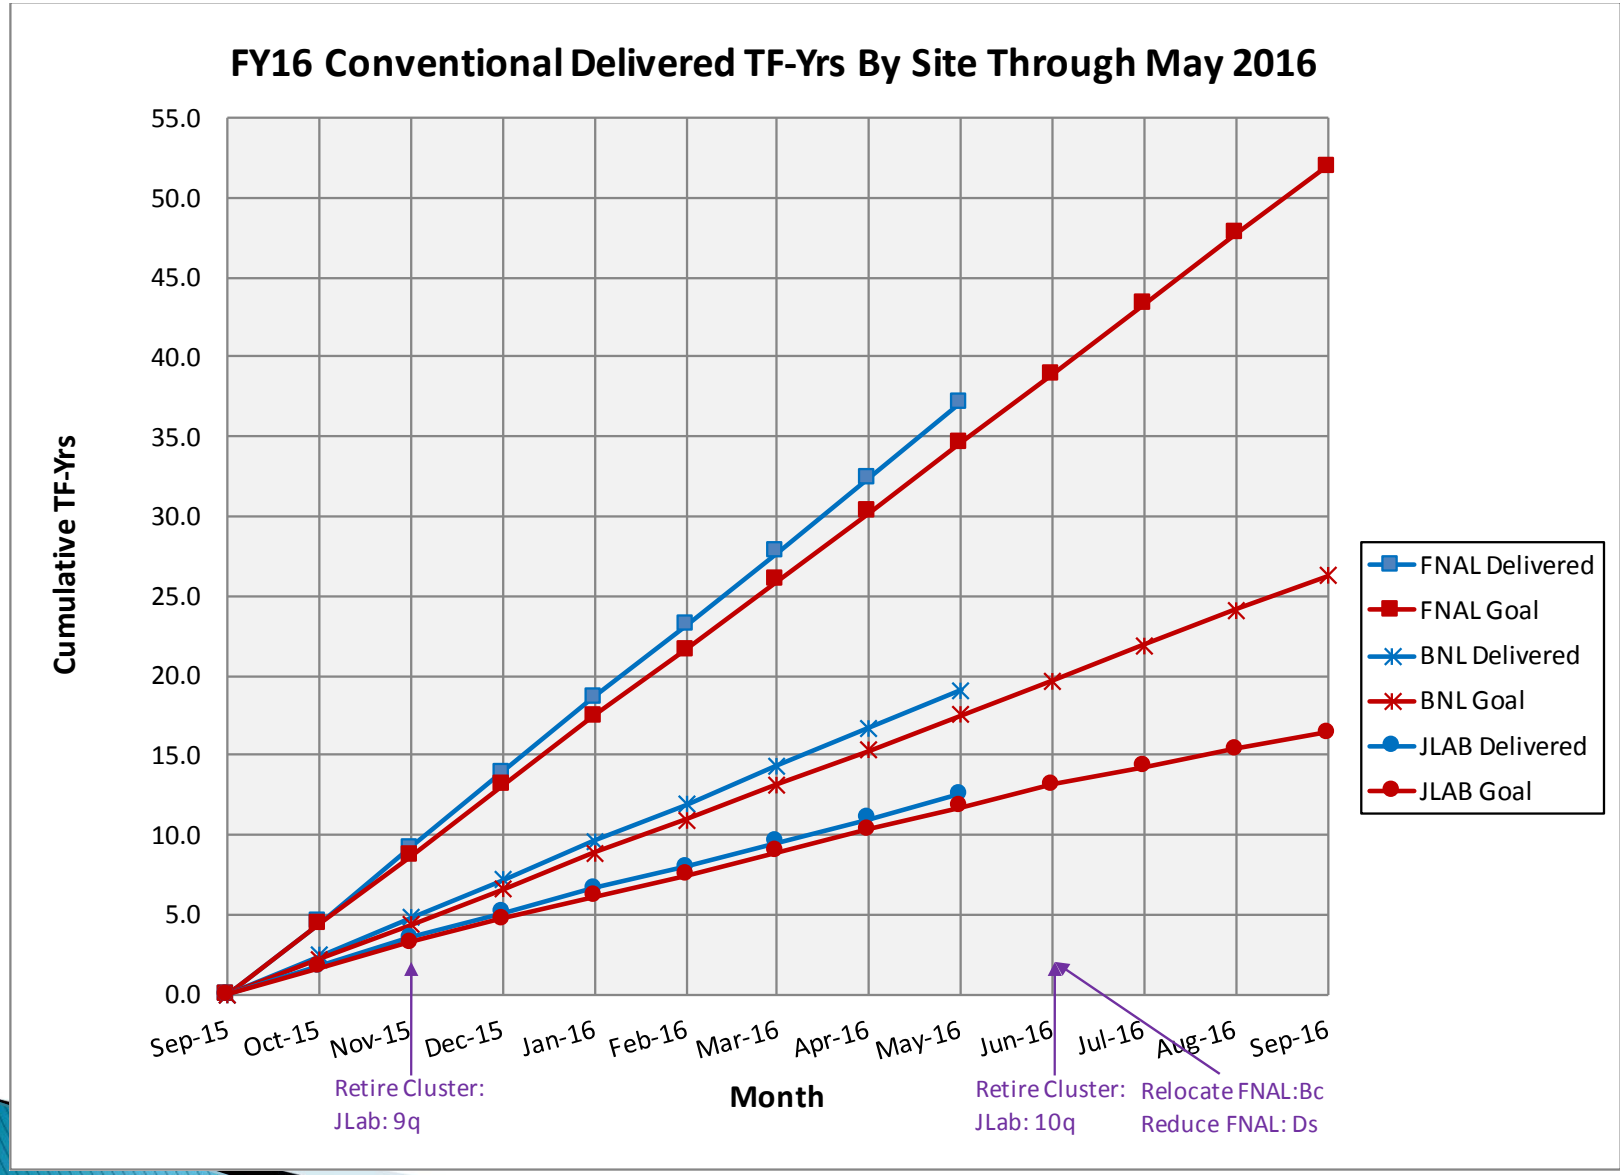

FY16 GPU Cluster Performance *(Effective TFlops-yrs delivered)*

Resources included are the FNAL Dsg and Pi0g clusters, and JLab 10g, 11g, and 12k clusters.
The project is on-target to exceed both the accelerated computing KPI as well as updated Goal.

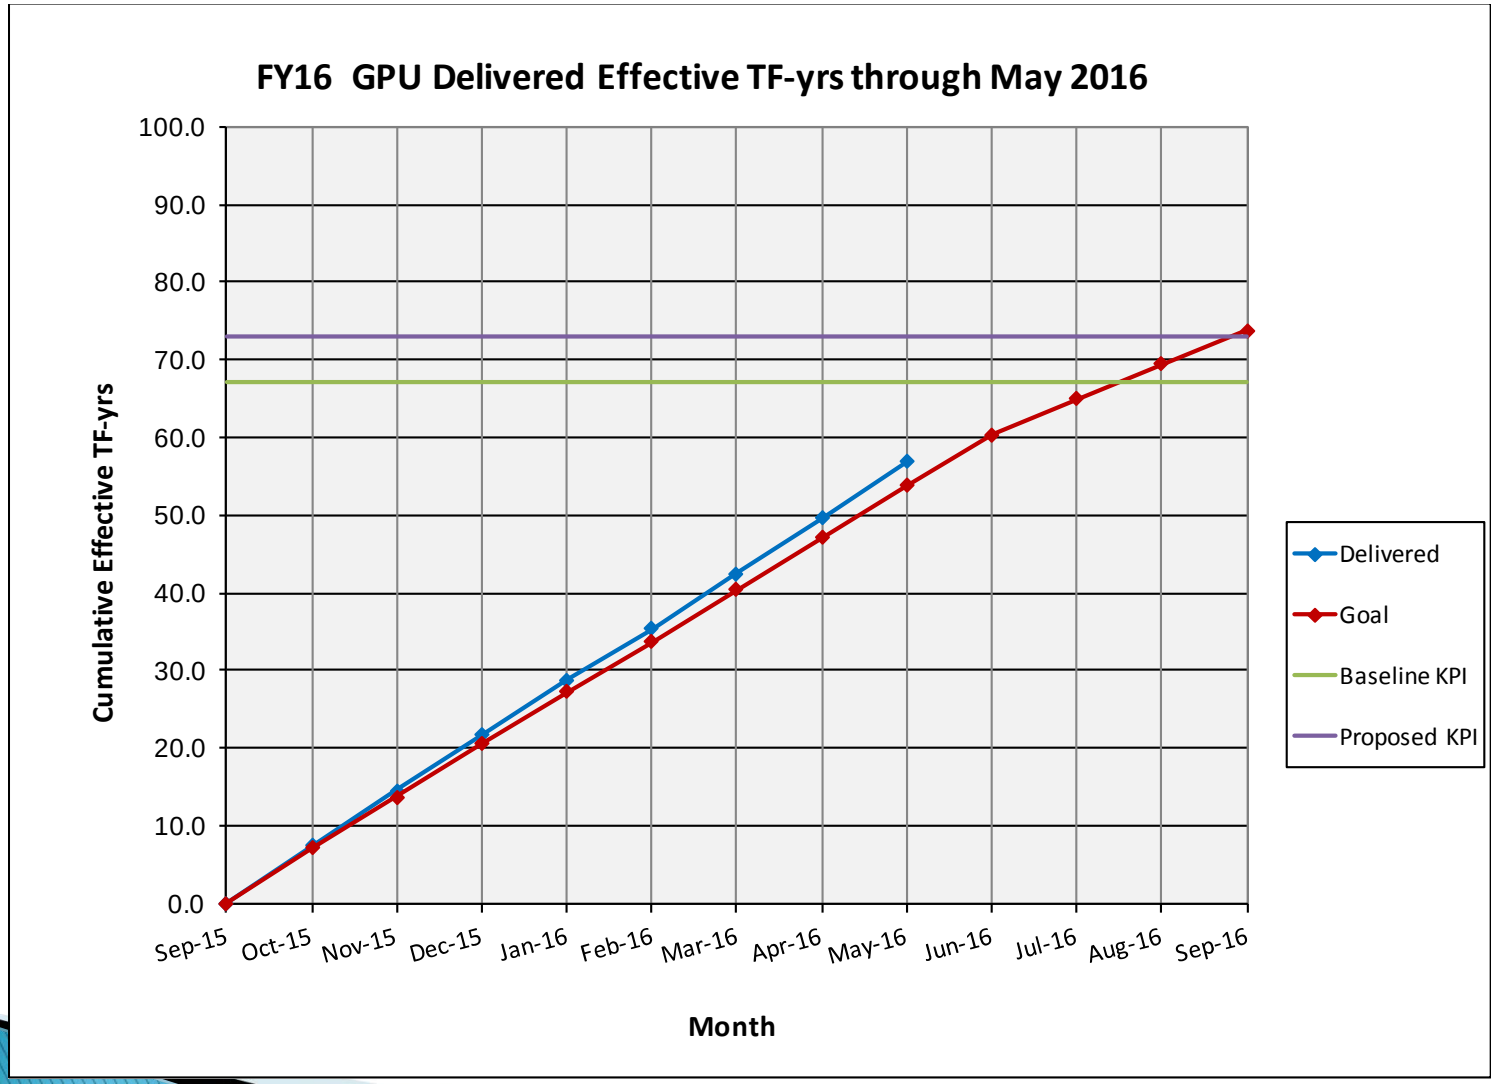

FY16 GPU Cluster Performance *(Effective TFlops-yrs delivered)*

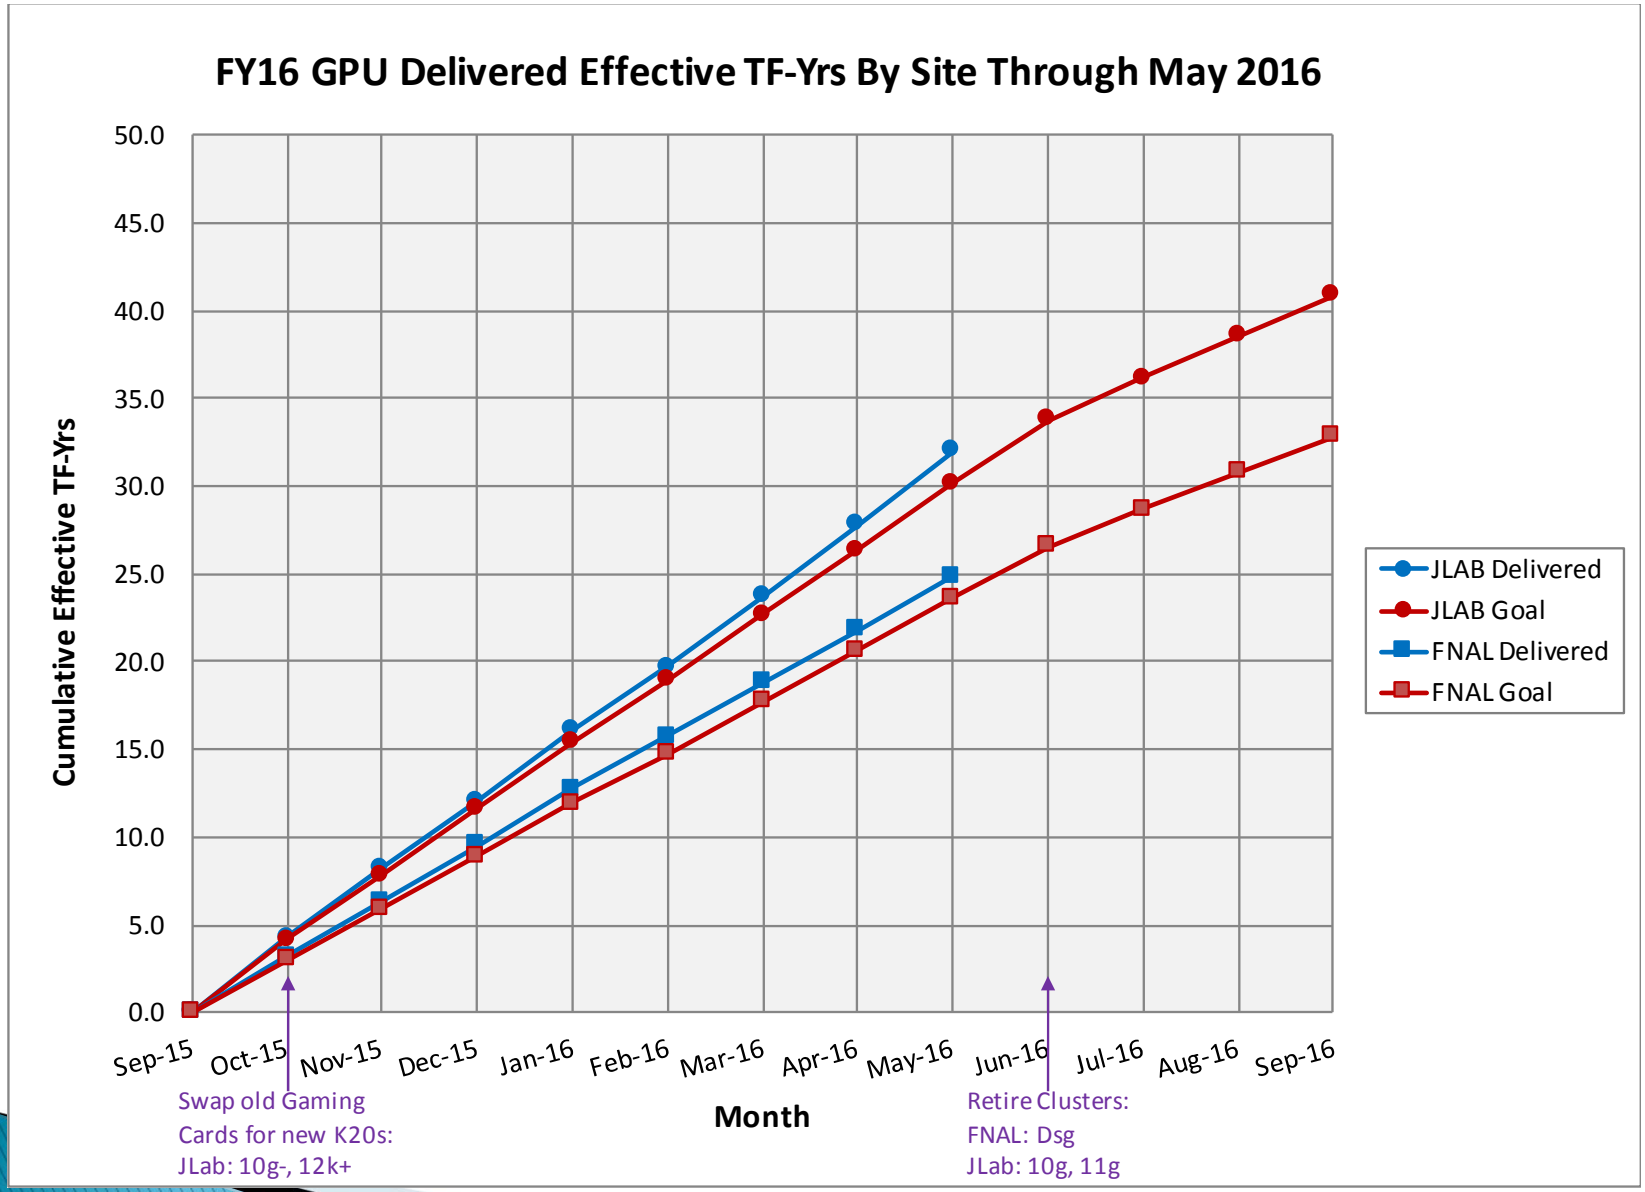

FY16 Financial Performance To Date

FY16 Project Cost Summary to Date (May 2016)

Status through May 2016; fiscal year complete: 67%

Fund Type	FY15 Carry-over	FY16 Budget	Total FY16 Funds Available	FY16 Actual Costs	FY16 Obligations	% Spent & Obligated
Equipment	---	$ 900 K	$ 900 K	---	--- (*)	0%
Operating	$ 375 K	$ 2,017 K	$ 2,393 K	$ 1,410 K	$ 2 K	59%
Sub-total	$ 375 K	$ 2,917 K	$ 3,293 K	$ 1,410 K	$ 2 K	43%
Mgmt Reserve	---	$ 83 K	$ 83 K	---	---	0%
TOTAL	$ 375 K	$ 3,000 K	$ 3,376 K	$ 1,410 K	$ 2 K	42%

Cost Performance Analysis

- FY16 Financials: *a detailed monthly time profile forecast to better track costs.*
- Our spending is on target per forecast ($1,410k spent vs $1,452k forecast)
  - IBM Maintenance (annual cost) posted in May.
  - (*) FY16 Acquisition will be posted in June, so not yet in this table, will consume project EQ funds.
  - Complex FY16 Acquisition is causing higher acquisition staff costs than forecast, offset by lower operations costs.
- There has been no draw on the Management Reserve.

FY16 Project Cost Performance to Date (May 2016)

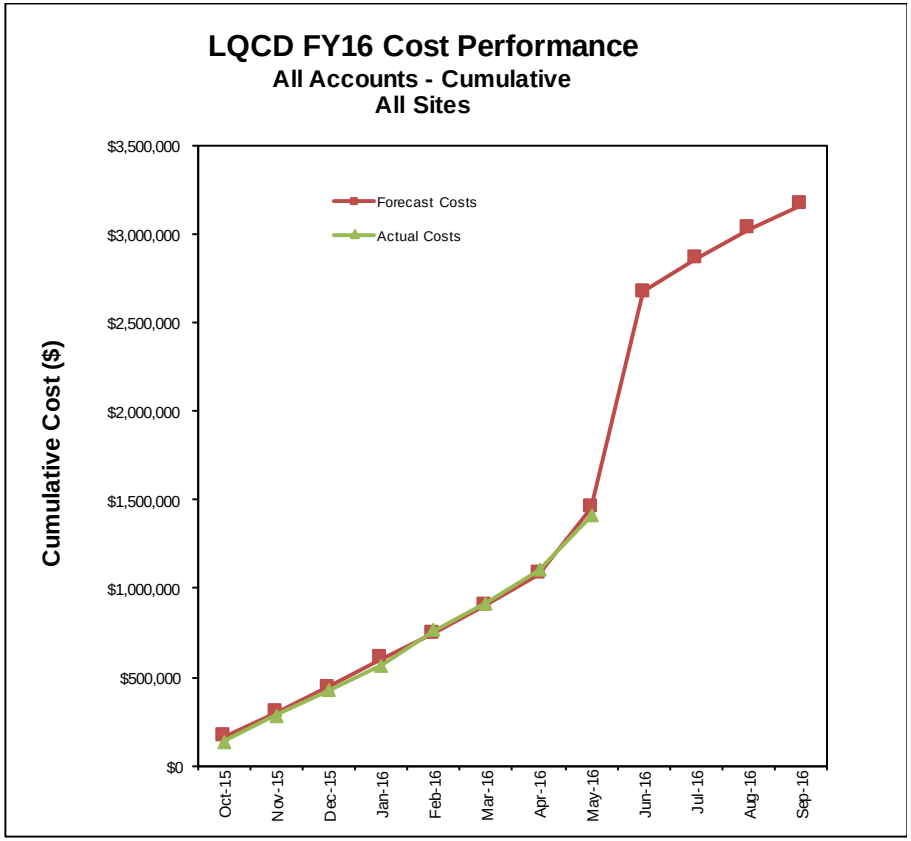

Cumulative Project Costs

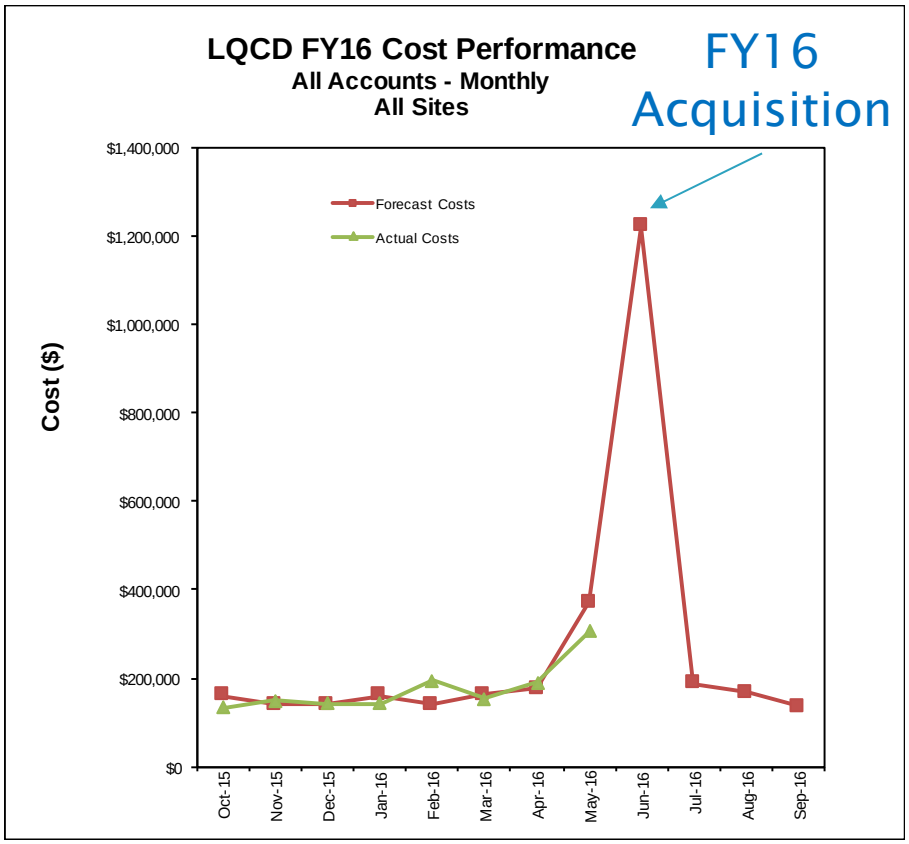

Monthly Project Costs

Summary: to date…

* Project is mildly under-spent so far in FY16
* Forecast: modest carry-over at end FY16 *possible*

Hardware Acquisition Strategy

Acquisition Strategy and Plans

- Baseline Acquisition Strategy
  - Our current acquisition strategy calls for a total of four procurements across fiscal year boundaries at JLab in FY16-17 and FNAL at FY18-19 (Baseline Plan).
  - On an annual basis, we develop and execute an acquisition plan that considers viable options and informs our purchasing decisions.
    - FY16 hardware acquisition activities are well underway.
    - Effort is being led by Chip Watson – new system will be deployed at Jefferson Lab.
    - Details will be discussed in a later session.
- Revised Acquisition Strategy
  - BNL has expressed an interest in cluster hosting, which would broaden and strengthen our base of technical expertise and compute cluster offerings.
  - Result is a proposed change in our acquisition strategy and operating model. Details of this will be discussed in a later session.
  - We appreciate your consideration of this proposed change and look forward to your comments and recommendations.

Summary

- Our hardware facilities are running well and we are successfully executing against our plans.
- We exceeded our FY15 performance goals.
- We are on target to meet or exceed our FY16 goals.
- We have worked hard to maximize our hardware portfolio and we have developed and executed strategies that optimize our procurements.
  - We are extending lifecycles to 5 years of service, where appropriate.
  - We have followed a proven acquisition process in past years with successful results; we will continue this approach going forward, with an eye towards identifying opportunities for improvement.
  - Our FY16 acquisition is well underway, with a variety of architectures considered.
- We remain focused on effectively optimizing the use of our resources to meet USQCD needs and maximize scientific output.