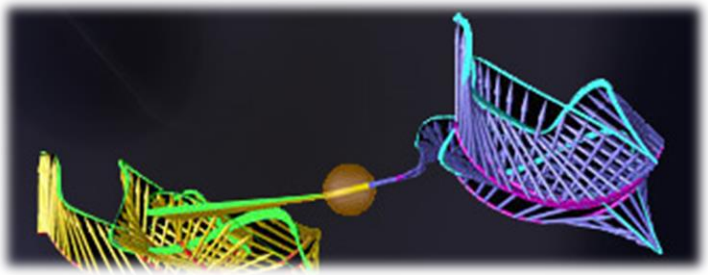

Report from the Project Office

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

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

USQCD All-Hands Meeting
Jefferson Lab
April 28-29, 2017

Outline

- LQCD-ext II progress to date
- Updates to our baseline operations plan
- Organizational changes
- Planning for the annual DOE review
- FY17 hardware acquisition activities
- FY18 acquisition plans
- User survey results

LQCD-ext II Project Status

- We're in the third year of the 5-year extension (Oct 2014-Sep 2019)
- We've received $8M of our planned $14M in funding (57%), in accordance with our baseline funding profile
  - ($2M in FY15; $3M in FY16, $3M in FY17).
- The computing we've delivered to the collaboration through March 2017 continues to exceed our baseline goals (TF-yrs delivered).

	FY17 [1]			Cumulative (Oct '15 Thru Mar '17		
	Goal	Actual	% of Goal	Goal	Actual	% of Goal
Conventional Resources[2]	68.2	73.4	108%	257.9	279.2	108%
Accelerated Resources[3]	40.9	43.5	106%	257.5	269.5	105%

1) FY17 performance through March 2017.
2) Conventional resources operational in FY17: Bc, Pi0,12s, 16p, BG/Q, 10% of DD2 prototype BG/Q rack (Bs retired Dec 2016)
3) Accelerated resources operational in FY17: Pi0g, 12k, (10g and 11g retired Dec 2016, BNL-IC brought online Jan 2017).

FY17 LQCD-ext II Project Performance

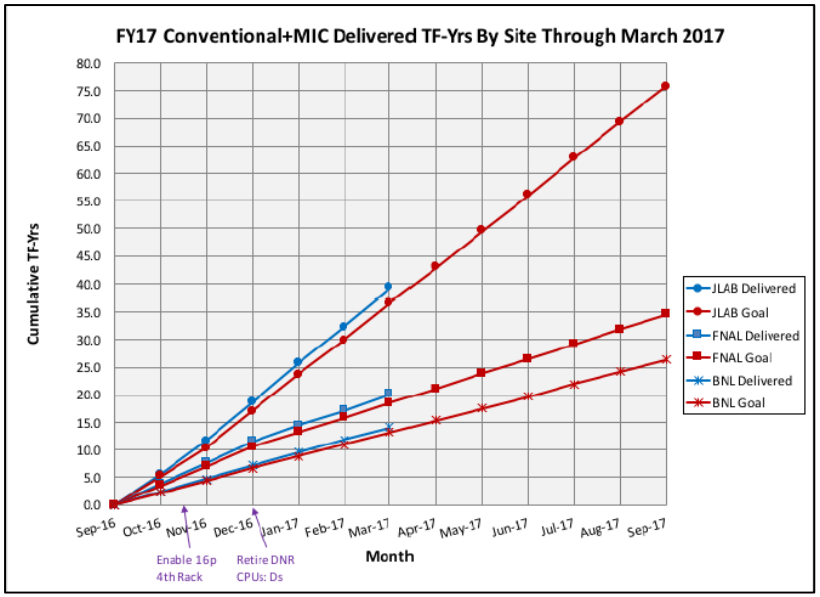

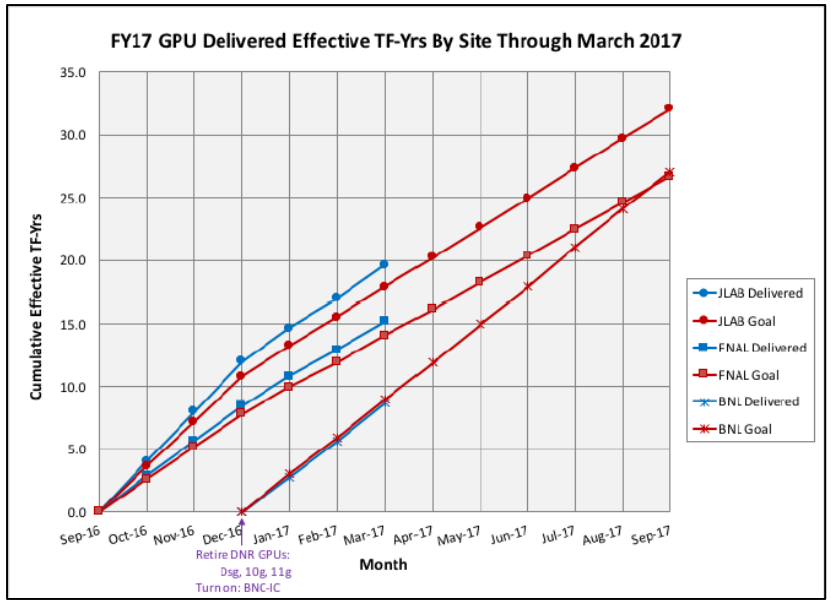

FY17 data for conventional resources are shown.

Goals are being exceeded because of excellent uptime at all three sites and running Ds beyond planned retirement date.

- The uptime goal is 8000 hours per year (91.3%).
- Performance goal is based on an average of the sustained performance of domain wall fermion (DWF) and highly improved staggered quark (HISQ) algorithms

FY17 data for accelerated clusters is shown.

Goals are being exceeded due to excellent uptime at all three sites and running Dsg, 10g and 11g beyond planned retirement dates.

- The uptime goal is 8000 hours per year (91.3%).
- Conversion from GPU-hrs. to effective TF-yrs is 140 GF/GPU, based on allocation-weighted performance of GPU projects running from July 1, 2012 through Dec 2012.

Approved Changes to our Baseline Plan

Site Operations (CR16-01)

- Baseline operations plan called for cluster hosting at FNAL and JLab through Sep 2019, and operation of the BG/Q half-rack at BNL through Sep 2017.
- Change Request 16-01 was approved by Change Control Board (CCB) and Federal Project Director as required.
  - BNL began delivering cluster computing resources in Jan 2017.
  - BNL will purchase, deploy and operate new LQCD clusters in future years (planning for the FY17 acquisition is in process).

Performance Goals (CR16-02)

- The approved baseline defined performance goals separately for conventional and GPU-accelerated machines.
- New computing architectures required us to redefine and combine these performance goals.
  - New MIC technologies do not neatly fit into either category, constraining the computing project to only invest in Conventional and Accelerated Computing at a certain level each year in order to be judged successful.
- Change Request 16-02 was approved by the CCB and Federal Project Director as required.

Changes in our LQCD-ext II Project Team

Organizational changes:

- We have included Ted Barnes (DOE-ONP) to acknowledge the very active role he continues to play on the project.
- Alexandr Zaytsev has replaced Shigeki Misawa as co-Site Architect at BNL.

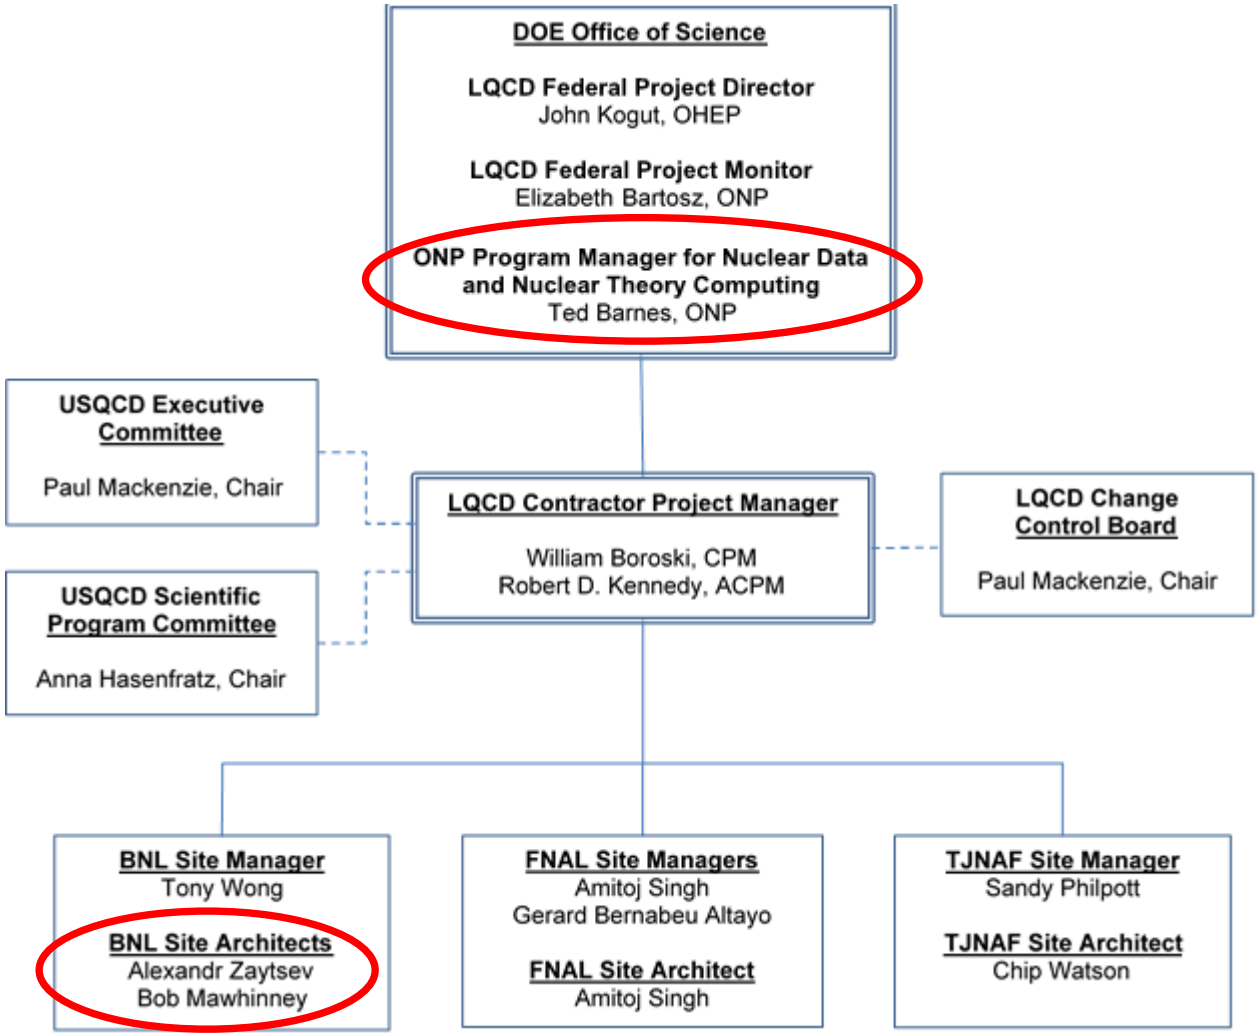

Planning for our Upcoming DOE Annual Review

- 2017 Annual Review scheduled for May 16-17 at Fermilab
- Review charge very similar to previous years….
  - Continued significance and relevance of the LQCD-ext II project, with an emphasis on its impact on the experimental programs' support by OHEP and ONP
  - Progress towards scientific and technical milestones
  - Status of technical design and proposed technical scope for FY17
  - Feasibility and completeness of proposed budget & schedule
  - Responsiveness to recommendations from last year's review
  - Effectiveness of USQCD in allocating LQCD-ext II resources to its community of lattice theorists
- …but with a formal request for USQCD to present its plans for further capacity computing
  - Will USQCD be requesting a further extension of the IT hardware project beyond FY19?
  - If so, what is the status of a whitepaper presenting the research plan?
  - If not, what are the plans for ramping down the current project?

FY17-18 Acquisition Activities

Rob Kennedy
LQCD-ext II Associate Contractor Project Manager

Context: Acquisition Schedule Post-CR

Plan Name	FY16	FY17 Deployments	FY18 Deployments	FY19 Deployment
Former Baseline	JLab	JLab (FY16 options)	FNAL	FNAL (FY18 options)
3-Site Cluster Hosting Baseline	JLab	1/3 JLab (FY16 options) 2/3 BNL	2/3 BNL (FY17 options) 1/3 FNAL (initiate procurement)	FNAL (execute procurement)

- 3-Site Cluster Hosting revised Acquisition Schedule
  - Split 4 acquisition budget years across 3 sites
  - Constraint: Maintain same level of delivered computing
- 40-node allocation on BNL-IC (K80 GPUs)
  - Production 1/4/2017. Allocation through end FY19
  - "40 nodes" is time-averaged. Can be more or less anytime.
  - Not traditional acquisition, but adds computing to portfolio
    - Also, implementing access to storage, tape archive there

FY17 Acquisition Status

- JLab: ~1/3 of Computing Acquisition Funds
  - Options purchase based on FY16 acquisition contract.
  - Expanded 16p to 256 KNL nodes (plus spares) very early in FY17.
- BNL: ~2/3 of Computing Acquisition Funds
  - Led by Bob Mawhinney, Alex Zaytsev. *Details: Bob M's Saturday talk*
  - Acquisition team working with Acquisition Review Committee
- FY17 Acquisition Review Committee – formed earlier this year
  - Review proposed FY17 (BNL) computing hardware acquisition plan
    - Chair: Rob Kennedy
    - Focus: develop more USQCD-specific software benchmarks for RFP process
  - Members include Site Architects, Site Managers, *Collaboration Reps*:
    - Carleton Detar, Steve Gottlieb, Chulwoo Jung, James Osborn, Frank Winter
  - Draft report available May '17. Early Notables from Acquisition team:
    - Target job size range: jobs using up to ~16 nodes
    - Dual-rail with KNL is not cost-effective vs Single-rail KNL for target job sizes
    - SPC: much higher "over-request" % for CPU and KNL than for GPUs

FY18 Acquisition Plans

- BNL: ~2/3 of Computing Acquisition Funds
  - Options purchase based on FY17 acquisition contract.
    - Most likely, this will lead to more of the FY17 choice.

- FNAL ~1/3 of Computing Acquisition Funds
  - Hold this portion of FY18 funds for a purchase in FY19.
  - Initiate the FY18-FY19 acquisition process in FY18.
    - Take as far as possible without FY19 funds on hand.
  - FY19 Funds arrive: FNAL executes FY18/19 RFP ASAP for "early" deployment of FY19 computing.
    - Plans for FY20 and later operations may impact this.

User Survey Results

Rob Kennedy
LQCD-ext II Associate Contractor Project Manager

FY16 Survey Results

- The FY16 User Survey:
  - Measured user satisfaction from October 2015 through September 2016
  - Survey open from through December 16, 2016 to March 10, 2017
  - Same format as in recent years, 29 questions designed to measure satisfaction with:
    - LQCD Compute Facilities
    - USQCD Resource Allocation Process
- The User Survey was distributed to all scientific members of USQCD
  - Responses were received from 73 individuals vs. 66 in FY15
  - 26 of 27 PI's responded: 96% response rate vs. 86% in FY15
  - 33 of 50 most Active Users responded: 66% response rate vs. 50% in FY15
- FY16 overall satisfaction rating with Compute Facilities = 93%
  - Exceeds LQCD Computing Project KPI goal of 92%. Was 97% in FY15.
- FY16 overall satisfaction rating with Resource Allocation Process = 85%
  - Down from FY15's rating (91%). At the level in FY11,12,14 (ratings in mid-80's).

FY16 Survey Results

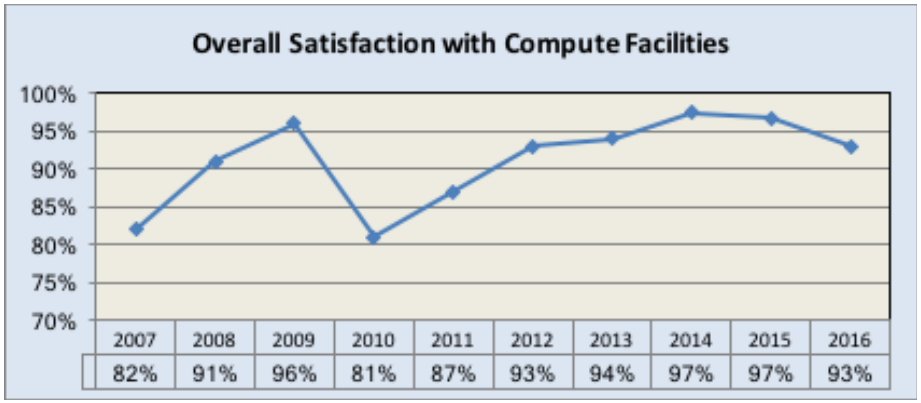

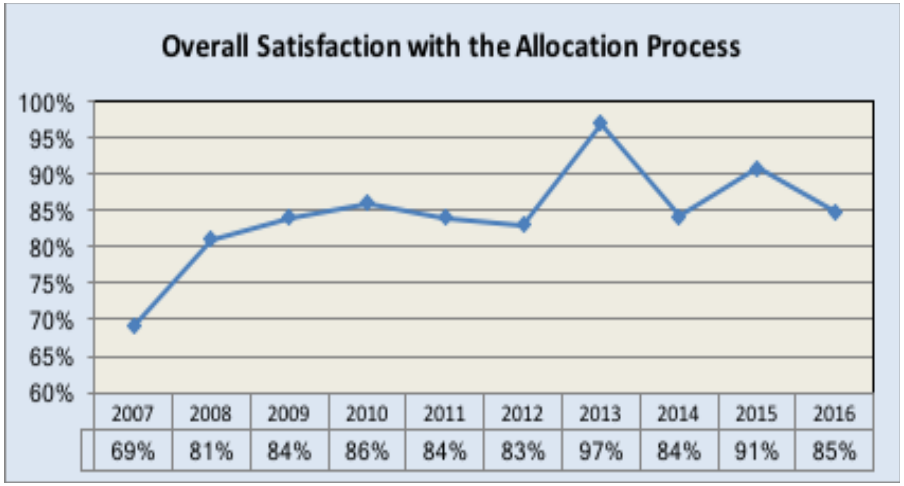

- User Comment Topics: suggested by >= 2 user comments
  - LQCD: User Documentation at BNL, JLab – *action plan documented*
  - LQCD: Simplify Moving Projects from Site to Site – *discussing*
  - USQCD: Concern about turn-around time for Class B, C proposals – *discussing*
  - USQCD: Link between science priorities, top allocations, outcomes – *discussing*
- User Survey Report: near-final draft… but not final yet.
  - Please, talk to Bill or Rob at break if you have comments. Still time to provide input to report.
  - And you can always send email to Bill or Rob… do not have to wait for annual survey.

Other LQCD Computing Topics

Rob Kennedy
LQCD-ext II Associate Contractor Project Manager

Data Preservation Policy

- Please plan how to preserve your data after your allocation ends.
- Project developed a Policy on Data Preservation
  - Sites may implement policy a little differently to adapt to local environment
- Data Preservation Policy for Disk Storage
  - Disk data that is not covered by a storage allocation and not community-owned can be moved to tape after 1 month from the end of your allocation *at a site's discretion* unless prior arrangements have been made.
- Data Preservation Policy for Tape Storage
  - Tape data that is not community-owned and not used for 3 years after your allocation ends can be shelved *at a site's discretion* unless prior arrangements have made.
- Related: Managing Data Storage
  - Sites fit current project allocations AND community data in available resources.
    - Community Data status as defined by USQCD-EC.
    - Policy empowers sites to "clean up" data from past allocations that consumes resources
  - Sites have had to scale down past disk, tape allocations to fit available resources
  - More tape and/or disk storage resources = less CPU and/or memory resources.

Questions?