

Report from the Project Manager

Don Holmgren for Bill Boroski
USQCD All-Hands Meeting
BNL
April 19-20, 2013

Because of inclement weather in Chicago (and a flooded basement!) Bill Boroski was not able to travel to BNL and sends his regrets.

LQCD-ext Project Scope and Budget

- Acquire and operate dedicated hardware at BNL, JLab, and FNAL for the study of QCD during the period FY10-FY14.
- Currently running against the FY10 baseline, with these exceptions:
 - FY11 and FY12 procurement were mixes of conventional and GPU-accelerated clusters, rather than conventional clusters
 - LQCD-ext in FY13 assumed the responsibility of operating the machines purchased at JLab during the LQCD-ARRA project (9q, 10q, 9g, 10g, 11g)
 - Hardware budgets in FY13 and FY14 were adjusted to accommodate ARRA operations, provide salary and M&S support at BNL in return for 10% usage of the BG/Q DD2 prototype rack, and provide more disk and tape storage project-wide.

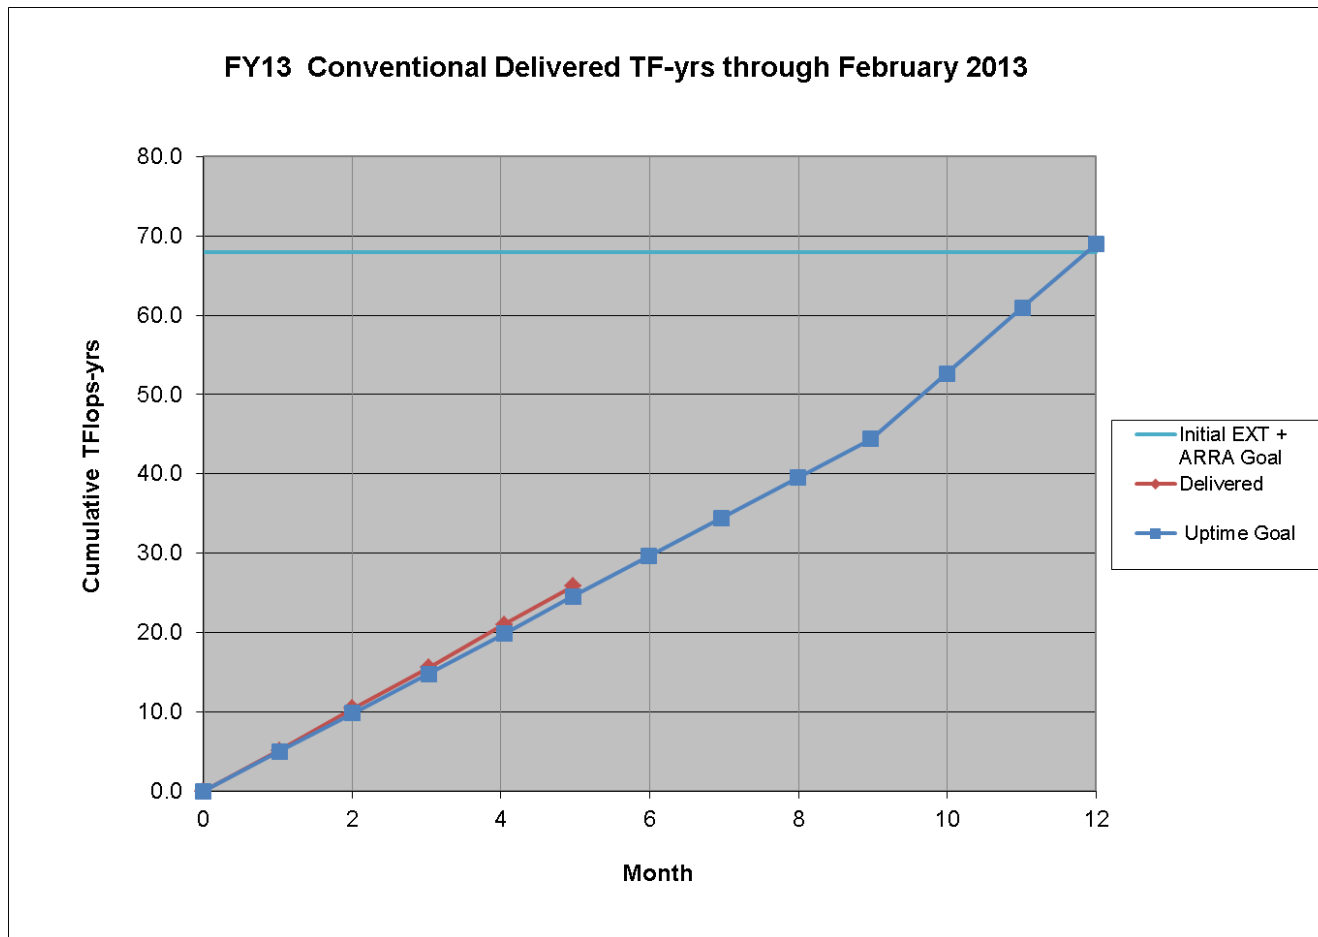
FY12 and FY13 Performance

- LQCD-ext stayed within budget in FY12
 - Delivered 37.6 TF-yrs (baseline goal: 34.0 TF-yr)
 - Delivered 667 GPU-KHrs on the Dsg at FNAL
 - Deployed 12s and 12k clusters at JLab
 - Met uptime, customer satisfaction rating, and closing of helpdesk ticket goals
- Also project is on track in FY13
 - Planned deployments (BG/Q, FNAL cluster) are in progress and will easily meet schedule goals
 - Meeting delivered TF-yrs (uptime) goals on both conventional and GPU hardware

Conventional Resource Performance

(TFlops-yrs delivered)

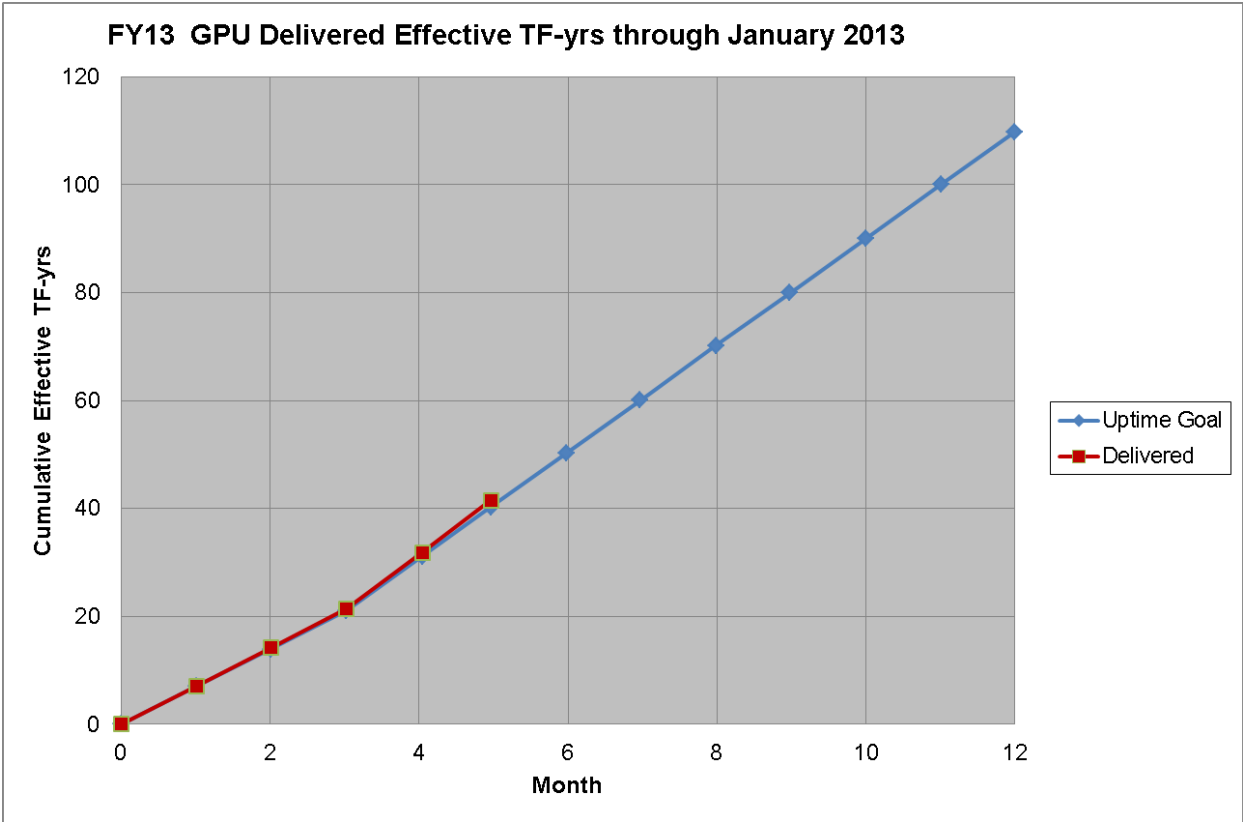
Data for FY13 through February 2013 are shown. The uptime goal is 8000 hours per year (91.3%). Computing resources included are the FNAL and JLab Infiniband clusters and the BNL BG/Q DD2 rack. New FY13 resources (BG/Q at BNL and cluster at FNAL) are assumed to come online June 30.



GPU Cluster Performance

(Effective TFlops-yrs delivered)

Data for FY13 through February 2012 for GPU-accelerated clusters is shown. 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 since July 1, 2012. Resources included are the FNAL Dsg and JLab 9g, 10g, 11g, and 12k clusters. The JLab 12k cluster came online Jan 3, 2013.



User Survey

- Received input from 76 individuals. *Thank you very much to all participants! The results are carefully reviewed by not only project personnel, but also our DOE program managers and reviewers.*
 - 61 employed or studying at a university or college, the rest mostly at one of the three host laboratories
 - 28 submit jobs daily, 22 submit occasionally or never
 - 43 had submitted jobs at FNAL, 23 at JLab, 7 at BNL
 - *Your free-form comments are particularly useful to the project*
- We used *Survey Monkey* software this year
 - Allowed us to track by e-mail address which users had not responded (but still kept responses anonymous)
 - Assuming there are no objections, we will use this tool again for the FY13 survey in the early Fall

User Survey – Facility Ratings

Category	FY07	FY08	FY09	FY10	FY11	FY12
Overall Satisfaction	82%	91%	96%	81%	87%	93%
System Reliability	74%	90%	84%	76%	91%	89%
Ease of Access	73%	74%	77%	76%	83%	92%
User Support	86%	100%	92%	88%	92%	94%
User Documentation	78%	92%	81%	73%	81%	89%
Responsiveness of Site Staff	89%	97%	98%	90%	90%	92%
Effectiveness of Other Tools	77%	72%	83%	86%	88%	92%

- All ratings except System Reliability improved over FY11
- Final report will show breakdowns by laboratory
 - In some categories there are considerable differences in ratings between the labs
 - The project is working to understand ratings of concern and will attempt to address them
- We will work to have higher participation in FY13 (in FY12, only 52% of active users responded to the survey)

User Survey – Allocations Process

Category	FY07	FY08	FY09	FY10	FY11	FY12
Overall Satisfaction with Proposal Process	69%	81%	84%	86%	84%	83%
Clarity of the Call for Proposals	79%	91%	93%	93%	93%	94%
Transparency of Allocation Process	61%	64%	79%	86%	74%	86%
Apparent Fairness of Allocation Process	63%	73%	88%	86%	93%	86%
Belief that Allocation Process helps maximize scientific output	70%	78%	85%	79%	88%	80%

- Slight drop in overall satisfaction
- Of concern are the notable drops in “apparent fairness” and “maximizes scientific output” ratings

Utilization of LQCD-ext Hardware

- In the first months of the current allocations year, utilization of the FNAL and JLab facilities was lower than normal (see FNAL and JLab site talks)
- Because we are in the process of applying for funds for a possible FY15-FY19 extension of LQCD-ext, it is very important that the community utilize these resources fully (assuming, of course, that they are still needed)
- We very much appreciate the community's quick increase in usage last Fall when the project raised concerns about low utilization
- Utilization (rather, low utilization) is a frequent topic of discussion on monthly calls with the DOE NP and HEP program managers

Upcoming Activities

- Planning for FY14 hardware purchases has begun
- As in prior years, we need to divide funds among the various hardware architectures in the way that best optimizes our portfolio of hardware for the scientific requirements of USQCD (BG/Q, conventional, GPU-accelerated, Phi-accelerated)
- We will follow a similar schedule to last year:
 - Strategy for making a decision will be presented to the LQCD-ext annual review panel May 9,10
 - Executive Committee and SPC are asked to provide the projected scientific program requirements for the various architectures by June 15
 - LQCD-ext will prepare an Alternatives Analysis with recommendation for hardware purchases by the end of July. EC provides advice for how to proceed by Aug 10.
 - LQCD-ext provides FY14 financial plan to DOE by Aug 20 (or earlier if requested)

Questions?