

# Report from the Project Manager

Bill Boroski  
LQCD-ext Contractor Project Manager

USQCD All-Hand's Meeting  
Fermi National Accelerator Laboratory  
May 4-5, 2012

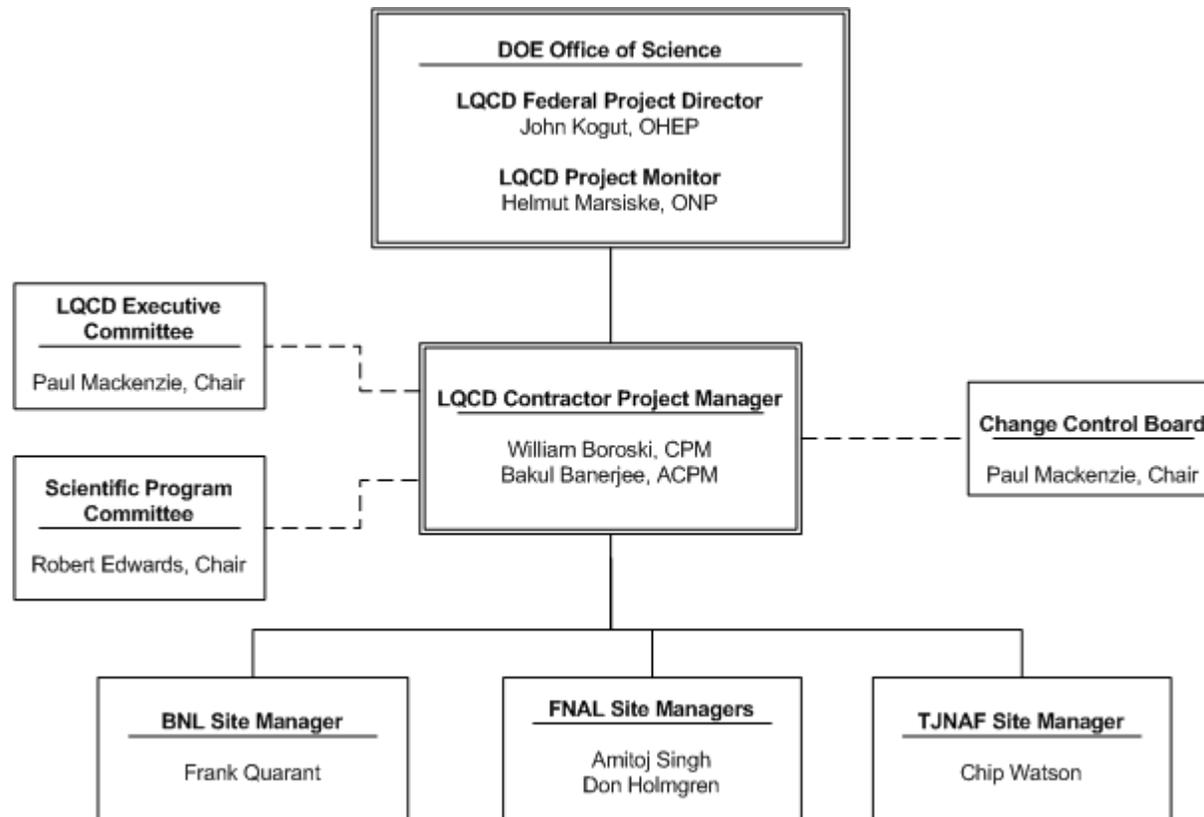
# Outline

- ▶ Updates to project scope, organization, and budget
- ▶ FY11 / FY12 performance results
- ▶ User survey results
- ▶ Facility utilization
- ▶ Hardware acquisitions
- ▶ 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.
- ▶ Currently executing against baseline plan, with a few exceptions
  - QCDOC at BNL was operated through August 2011
  - Kaon (FNAL) and 7n (JLab) are being operated beyond planned lifetimes
  - FY11 procurement included a mix of conventional Infiniband cluster nodes and GPU-accelerated nodes. FY12 procurement will also contain a mix.
  - Planning to provide 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).
  - Will assume responsibility for operating and supporting the compute hardware at JLab acquired under the LQCD-ARRA project (FY13–14).

# Management & Oversight



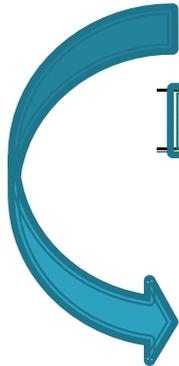
- ▶ Changes since last year
  - Robert Edwards replaced Frithjof Karsch as SPC Chair
  - Frank Quarant replaced Eric Blum as BNL Site Manager

# 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
<b>Total</b>	<b>3,000</b>	<b>3,250</b>	<b>3,600</b>	<b>4,100</b>	<b>4,200</b>	<b>18,150</b>



## 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
<b>Total</b>	<b>9,885</b>	<b>520</b>	<b>10,405</b>

Baseline storage budget was set at ~5% of total hardware budget

# 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
<b>Total</b>	<b>3,000</b>	<b>3,250</b>	<b>3,600</b>	<b>4,100</b>	<b>4,200</b>	<b>18,150</b>



## 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
<b>Total</b>	<b>9,885</b>	<b>520</b>	<b>10,405</b>

Baseline storage budget was set at ~5% of total hardware budget

# LQCD-Ext Project Budget

- ▶ We are currently half-way through the LQCD-Ext project.
  
- ▶ Changes in the budget forecast, relative to the baseline.
  - TPC reduced by \$100K due to tight budget constraints in FY12.
    - Was \$18.15 million; Now \$18.05 million.
  - Personnel Budget Changes
    - Updated salary cost basis for FY13-14
    - Modified staffing model based on operating experience
    - Increased staffing support to operate BG/Q and ARRA facilities in FY13-14
  - Storage Hardware Budget Changes
    - Increased to accommodate growing storage needs
  - Compute Hardware Budget Changes
    - Reduced to accommodate staffing support for BG/Q and ARRA in FY13-14
    - Reduced to accommodate increased storage needs
  - \$94K of unspent management reserve from FY10-11 has been applied to FY12 hardware procurement and deployment budget

# LQCD-Ext Project Budget

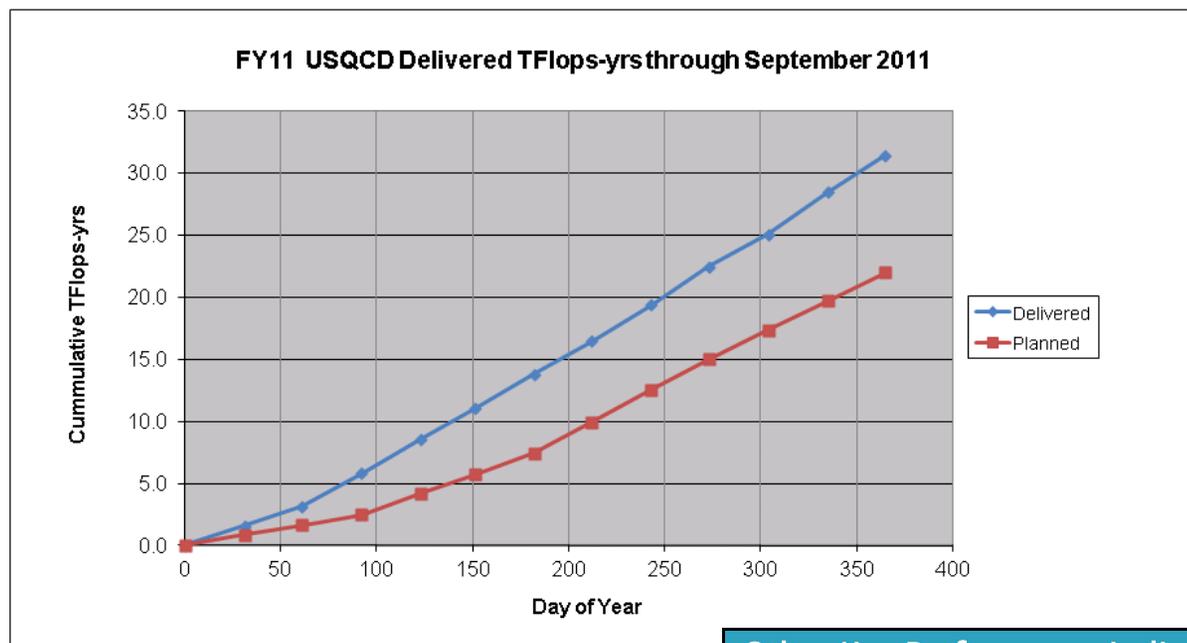
- ▶ Comparison of current forecast to baseline budget (\$K)

Expenditure Type	Baseline Budget	Current Forecast	Change Relative to Baseline	% Change
Personnel	6,885	7,038	153	2%
Travel	60	60	0	1%
M&S (spares, tape, etc.)	440	465	25	5%
Compute Hardware	9,885	9,526	(359)	(4%)
Storage Hardware	520	691	171	25%
Management Reserve	360	269	(91)	(25%)
<b>Total</b>	<b>18,150</b>	<b>18,050</b>	<b>(100)</b>	<b>(0.6 %)</b>

# Performance Measures & Metrics

# FY11 Performance Summary

- ▶ FY11 Goal = 22.0 TFlops-yrs
- ▶ Actual = 31.48 TFlops-yrs (143% of goal)



FY11 Acquisition Plan called for both Infiniband and GPU cluster deployments.

Milestone target dates for both IB and GPU cluster deployments were missed due to impact of Continuing Resolution and Thailand flooding.

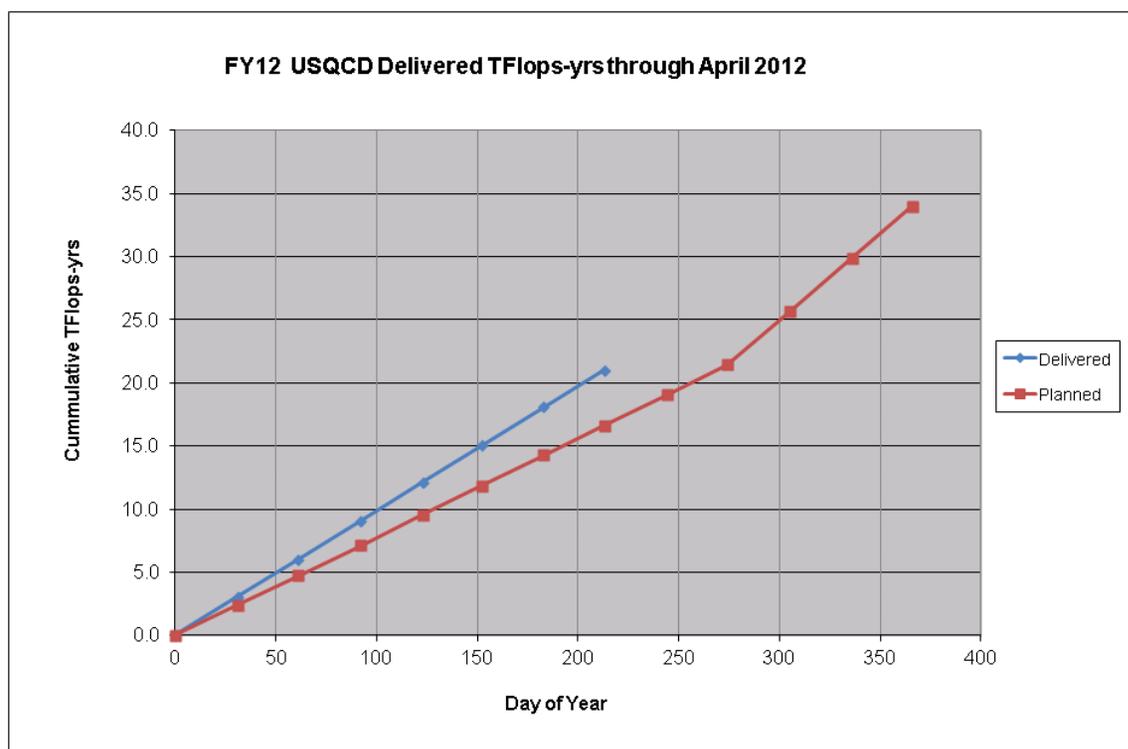
Project operated QCDOC (BNL), Kaon (FNAL), and 7n (JLab) beyond planned lifetimes.

Other Key Performance Indicators (KPIs)	Target	Actual
TFlops deployed	12 TF	17.5 TF*
Customer satisfaction rating	≥92%	87%
% tickets closed within 2 business days	≥95%	95%
% average machine uptime	≥95%	97%

\*Infiniband cluster = 9 TF; GPU cluster = 8.5 TF (effective)

# FY12 Milestone Performance *(TFlops-yrs delivered)*

- ▶ Data for FY12 conventional Infiniband clusters thru April 2012 are shown.
- ▶ The unmodified goal for FY12 is 34.0 TFlops-yrs.
- ▶ Goal through April = 16.6 TFlops-yrs
- ▶ Actual = 21.0 TFlops-yrs (126% of goal)



“Unmodified” project goal assumes only conventional Infiniband clusters

- Project is operating both Kaon (FNAL) and 7n (JLab) clusters beyond planned lifetimes
- At the current pace, even without contributions from the planned JLab IB cluster starting in FY12Q4, we will still meet the unmodified goal, because of strong uptimes and contributions from Kaon and 7n

We are beginning to formulate new project goals that take into account both conventional and GPU-accelerated clusters.

# User Survey Results

# FY11 Survey Results

- ▶ Following the suggestions made by the 2011 DOE Progress Review Committee, we modified the user survey in an attempt to encourage a higher response rate.
  - Reduced the total number of questions from 44 to 22.
  - Revised the wording of some questions.
  - Retained the ability for users to provide free-form comments.
- ▶ Received input from 61 users (small statistical sample).
  - Approximately 102 users submitted jobs to one of the three facilities during the past year
  - FY11 response rate = ~60% (61 individuals)
  - Improvement from FY10, when only 39 users responded to the survey call.
- ▶ Thank you very much to everyone who participated in the survey.
  - In addition to the feedback and insight it provides to the project team, the results are also carefully reviewed by our stakeholders.

# Satisfaction with Facility Operations

## Satisfaction Ratings Over Time:

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

- ▶ Although significantly improved over FY10, the overall satisfaction rating of 87% is below our target goal of 92%. We believe that the timing of several external factors may have contributed to this rating.
- ▶ Ease of access rating continues to suffer due to access issues associated with the use of Kerberos authentication.
- ▶ User documentation remains an area for improvement.
- ▶ User support and responsiveness ratings appear to have suffered due to loss of key knowledgeable individuals at one of our sites, and to understaffing at another.

# Satisfaction with Allocation Process

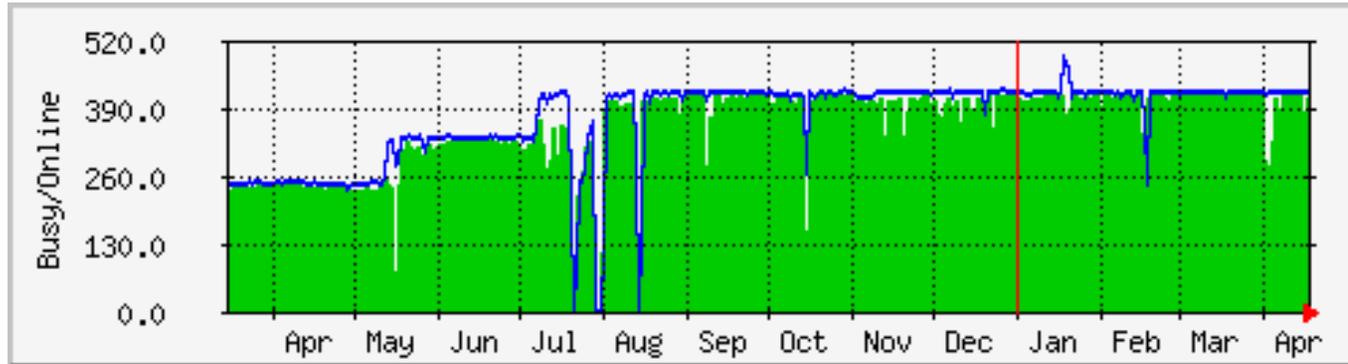
## Satisfaction Ratings Over Time:

	FY07	FY08	FY09	FY10	FY11
Overall satisfaction with the proposal process	69%	81%	84%	86%	84%
Clarity of the Call for Proposals	79%	91%	93%	93%	93%
Transparency of the allocation process	61%	64%	79%	86%	74%
Apparent fairness of the allocation process	63%	73%	88%	86%	93%
Belief that the allocation process helps maximize scientific output	70%	78%	85%	79%	88%

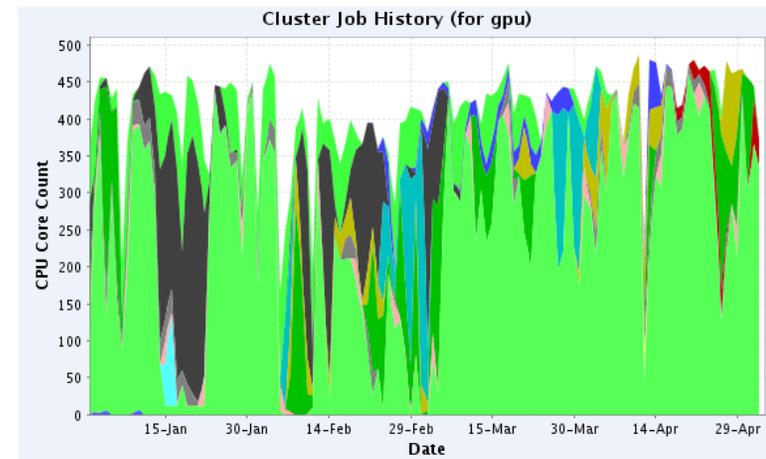
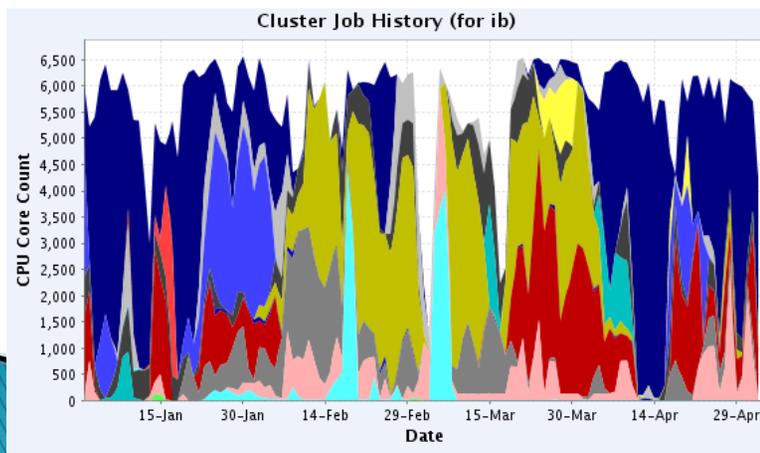
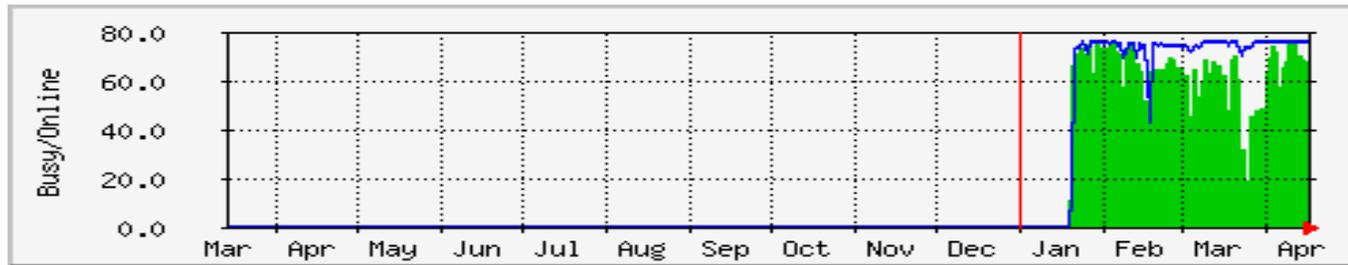
- ▶ User satisfaction ratings nearly met or exceeded prior year ratings in all categories except one: transparency of the allocation process.
- ▶ Several concerns were voiced by survey respondents regarding the allocation process.
  - Not clear why certain proposals appear to be preferred over others
  - Would be useful to have a clear statement of the scientific criteria under which proposals are to be evaluated, and of the scientific goals of USQCD
  - The CFP is getting too long, so subtle changes in a given year may go unnoticed. Perhaps changes should be noted early in the CFP message.

# Facility Utilization

Ds



Dsg



# Hardware Acquisition Planning

# FY13 Acquisition Planning

- ▶ With the emergence of new platforms such as GPU-accelerated clusters, we outlined a new strategy at the FY11 review that we are continuing to follow:  
*Procure systems that will best optimize our portfolio of hardware (including anticipated supercomputer time) against our portfolio of applications (including configuration generation).*
- ▶ In FY13, we once again have several hardware options to consider:
  - Infiniband clusters, GPU-accelerated clusters, BG/Q
- ▶ In order to maximize the use of hardware funds, we are in the process of gathering critical information
  - We will be gathering information on various hardware options, including the IBM BG/Q
    - Pricing and availability of production BG/Q hardware
    - Cost model for operating a BG/Q at BNL
  - We need your input to help us optimize the use of hardware funds and best meet scientific computing needs.
    - What applications will be able to be run on GPUs at that time?
    - What portion of the analysis computing can be done more cost effectively on GPUs vs. IB clusters?
- ▶ We have established a process for finalizing the FY13 acquisition plan that closely follows the FY12 planning process. We propose to use this process to gather information and make an informed decision regarding the planned hardware choice for FY13. Target decision date is mid-August.

# FY13 Acquisition Planning: Key Dates

Activity	Target Due Date
Project provides Executive Committee (EC) with data summarizing distribution of job types and sizes over the past year	Apr 15
Project presents acquisition strategy to external committee at DOE annual review	May 16
EC & Scientific Program Committee provides project with anticipated scientific program requirements for various architectures	Jun 15
Project prepares Alternatives Analysis document, which summarizes consideration of various options and proposes cost-effective solution for FY13 hardware deployment.	Jul 29
EC reviews Alternatives Analysis document and proposed solution, and provides advice to the Project on how to proceed.	Aug 10
Project prepares FY13 hardware acquisition plan and informs stakeholders	Aug 15
Project Manager provides Federal Project Director (OHEP) and Federal Project Monitor (ONP) with the FY13 Financial Plan, which contains information on the allocation of hardware funds to the host laboratories.	Aug 20

# Summary

- ▶ We are now half-way through the LQCD-ext project. Facilities are running well, we're executing well against our plans, and we're expanding the scope of the LQCD-ext project to include the BG/Q and ARRA machines.
- ▶ We successfully met or exceeded all but one of key performance goals in FY11. We did not meet our target deployment dates.
  - User survey results indicate areas for potential improvement.
  - We missed deployment milestones due to Continuing Resolution and other factors.
- ▶ We are on target to meet nearly all of our FY12 performance goals.
  - Our site managers continue to do a very good job of operating their respective systems for minimize downtime and maximize output.
  - We've been affected by the budget situation in Washington; Continuing Resolutions impact the timing of our procurement and deployment activities.
- ▶ We have significant opportunities to maximize our hardware portfolio going forward and are working to optimize our procurement strategies in order to make the most effective use of project resources.