

## Lattice QCD Computing Project

# Project Management and Performance

Bill Boroski  
LQCD Contractor Project Manager

2009 Annual Review  
Fermi National Accelerator Laboratory  
June 4-5, 2009





# Outline

- Project Scope
- Management & Oversight
- Work Organization and Planning
- FY08/09 System Deployment Summary
- Performance Measures and Metrics
  - Technical, Scientific, Cost and Schedule
- FY08 User Survey Results

# Project Scope and Budget

- Acquire and operate dedicated hardware at BNL, TJNAF, and FNAL for the study of quantum chromodynamics

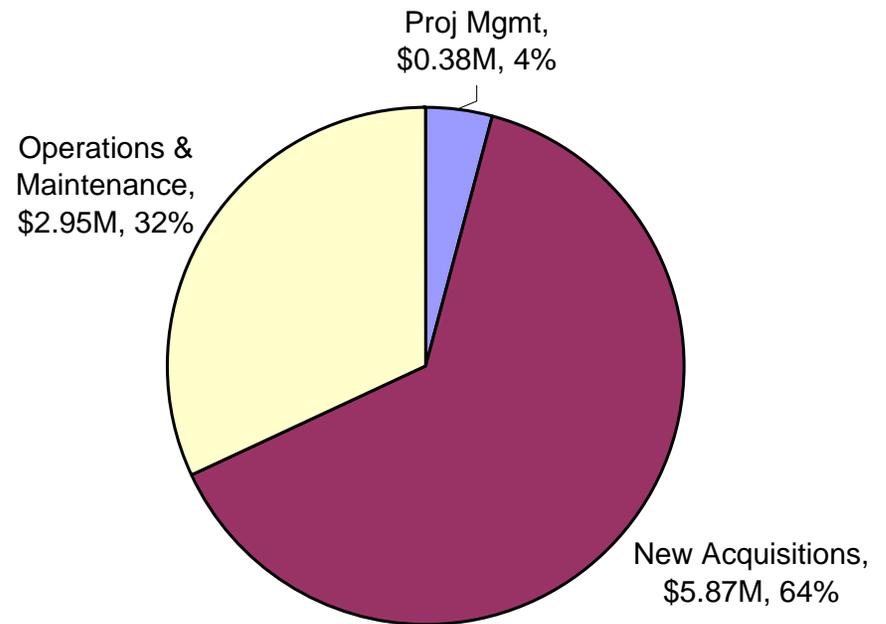
- Budget: \$9.2 million (*provided jointly by OHEP and ONP*)
- Period of performance: FY06 through FY09

- Project funding covers:

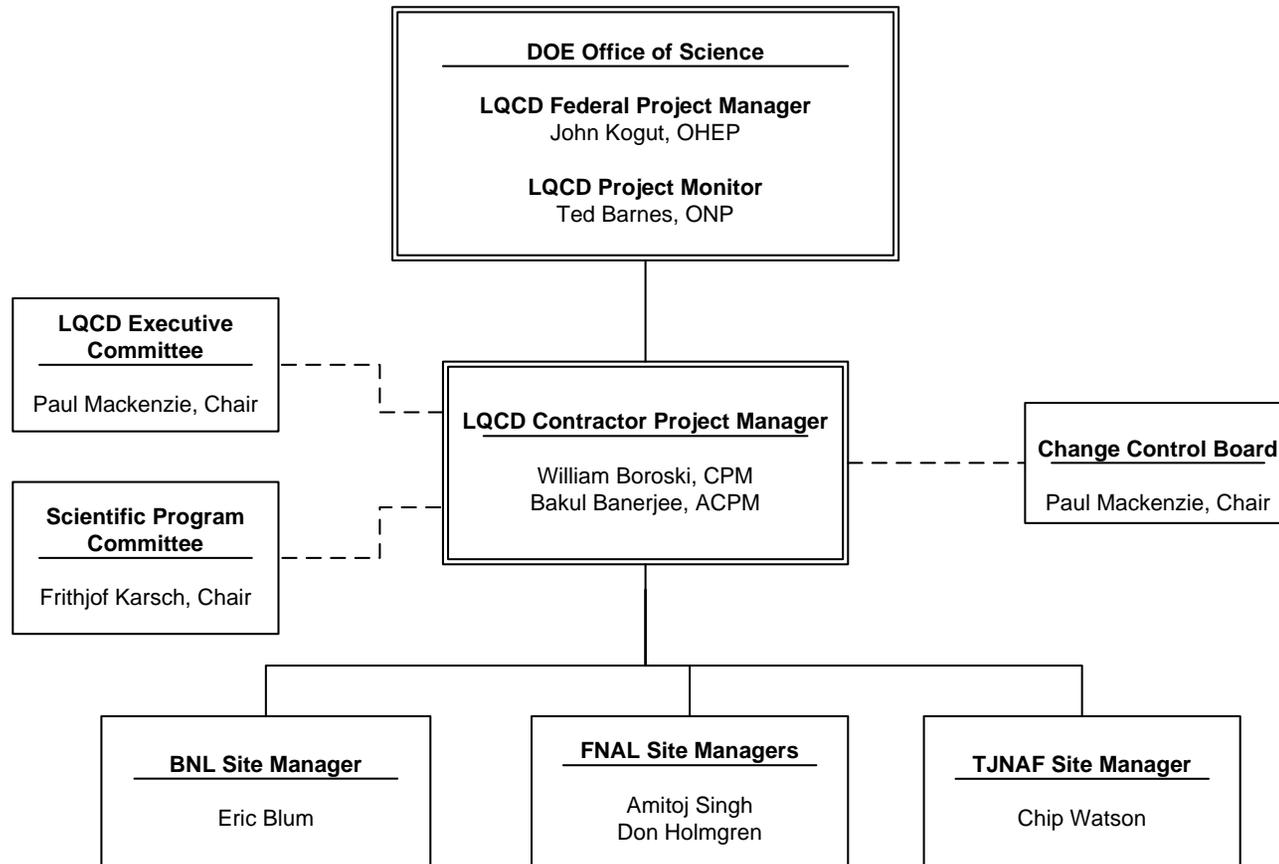
- Project management
- Operations and maintenance of existing systems
- Acquisition and deployment of new hardware
  - FY06: Kaon cluster at FNAL; 6n cluster at JLab
  - FY07: 7n cluster at JLab
  - FY08/09: JPsi cluster at FNAL

- Not in scope

- Software development
- Scientific software support



# Management Organization

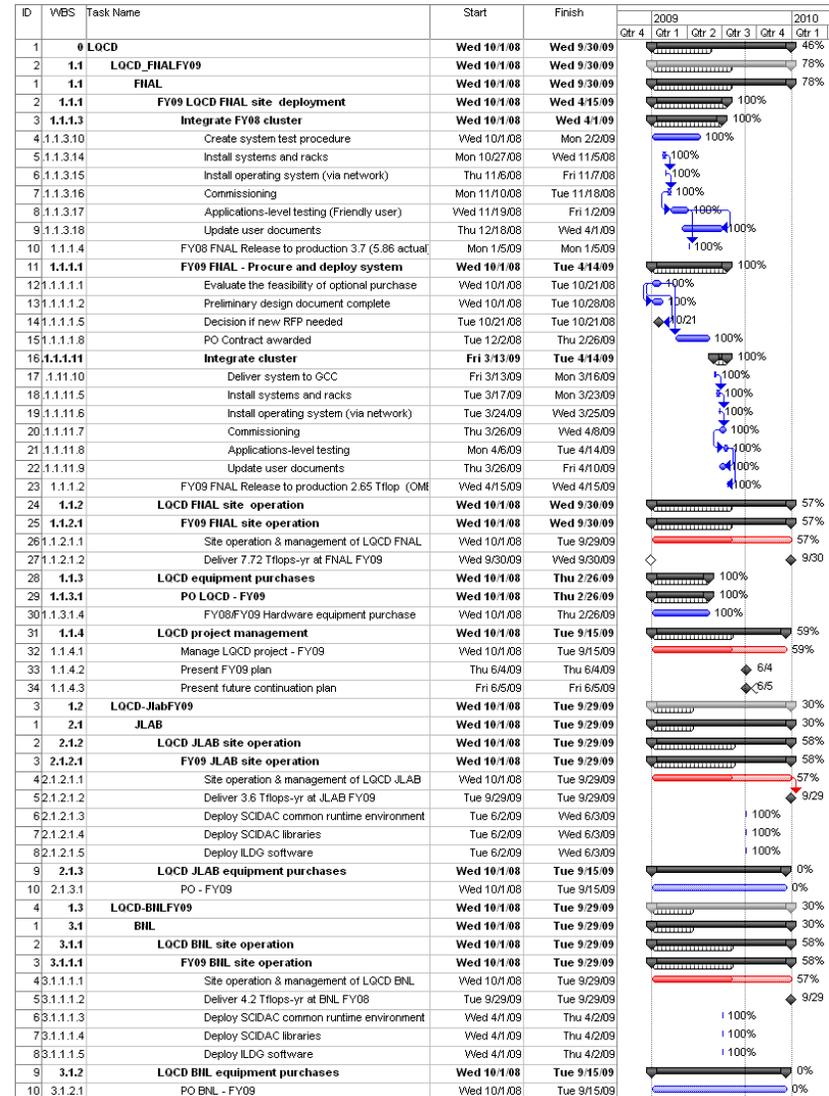


All federal and contractor project managers are certified “Level 1 Qualified IT Project Managers.”

Org changes since FY08 review include the Chairs of the Executive Committee, Scientific Program Committee, and Change Control Board

# Work Planning and Organization

- Project Execution Plan (PEP)
  - Controlled document defining project need, requirements, scope, management, cost and schedule, change control, etc.
- MOUs with host institutions
- 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, Alternatives Analysis, Annual Acquisition Plans, C&A Documentation





# Steady-state Operations & Maintenance

- Site Managers are responsible for day-to-day operations of their respective sites
- User allocations are determined annually by the Scientific Program Committee and provided to each site manager for implementation
- Site manager responsibilities include:
  - Establishing systems to track system performance and usage;
  - Reporting progress against goals;
  - Providing monthly cost and effort reports to the project office;
  - Ensuring that host laboratory commitments are met;
  - Identifying issues and concerns to the CPM.



# Procurement and Deployment of New Systems

- Project plan and performance goals called for a major new acquisition in each year of the project.
- Throughout the project, each procurement was treated as a sub-project
- Procurement and deployment plans, with timeline and milestones, were developed as part of the annual planning and budgeting process.
  - Each planning exercise took into account performance requirements and goals; existing facility capabilities and required facility upgrades; technical advances; etc.
  - Most recent activities were focused on the selection, procurement and deployment of the FY08/09 cluster at FNAL
- Five new systems have been deployed over the course of the 4-year project:
  - FY06: Kaon at FNAL (600 nodes) and 6n at JLab (260 nodes)
  - FY07: 7n at JLab (396 nodes)
  - FY08/09: JPsi at FNAL (864 nodes)



## Project-funded Workforce Staffing Model

- Several adjustments were made in the FY09 staffing model, in order to increase the level of systems admin support for SS operations.
  - Reduced site management from 0.25 to 0.15 FTE per site ( $\Delta = -0.30$  FTE)
  - Reduced sys admin support at BNL from 0.75 to 0.25 FTE ( $\Delta = -0.5$  FTE)
  - Reduced deployment support at FNAL from 0.75 to 0.5 FTE ( $\Delta = -0.25$  FTE)
  - Increased level of operations sys admin support at FNAL and JLab from 1.1 to 1.9 FTE per site ( $\Delta = 0.8$  FTE/site or 1.6 FTE/total)
  - Held project management support at FNAL at 0.5 FTE
- Net result was an increase in supported effort of 0.55 FTEs
- Based on differences in fully-loaded salary costs, corresponding salary cost increase was only \$8K.



# Communications and Reporting

- **Bi-weekly or Monthly Site Managers Meeting**
  - Address site-specific issues or concerns
  - Discuss procurement plans/activities
  - Exchange of other relevant information
  
- **Monthly DOE Program Office Meeting**
  - Report on progress against performance goals (TFlops-yrs delivered, cost, procurement activities, etc.)
  - General exchange of information
  
- **Quarterly Progress Reports**
  - Following OMB reporting guidelines and templates
  - Performance graded using “stoplight” system
  
- **Informal communications between federal and contractor project managers, as necessary**



# System Deployments

## FY08/09 Hardware Procurement

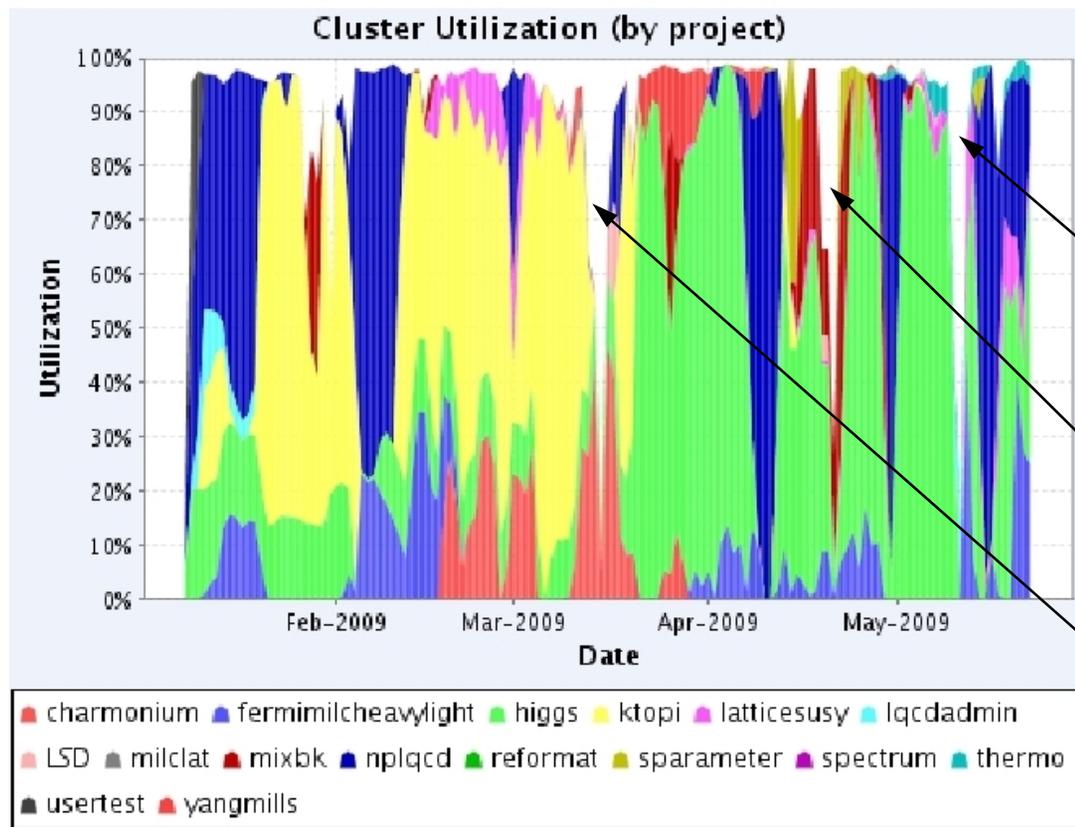
- As presented during the 2008 review, we combined the FY08/09 procurements with an option clause
  - Combining procurements resulted in cost savings associated with reduced labor costs
  - Approach endorsed by the 2006, 2007, and 2008 review committees
- The FY08 portion of the J/Psi cluster (5.75 TF, 586 nodes) was released to production on January 5, 2009
- The FY09 portion of the J/Psi cluster (2.65 TF, 270 nodes) was released to production on April 15, 2009
- Performance goals from the OMB Exhibit 300:

<b>FY08 Goals:</b>	<b>Planned Completion Date</b>	<b>Actual Completion Date</b>	<b>Actual Deployment Capacity</b>
12 Tflops-yrs aggregate computing delivered	09/30/08	09/30/08	12.1 TF-yrs
Procure and deploy 4.2 Tflops at FNAL	12/30/08	01/05/09	5.75 TF
<b>FY09 Goals:</b>			
Procure and deploy 2.0 Tflops at FNAL	06/30/09	04/15/09	2.65 TF
15 Tflops-yrs aggregate computing delivered	09/30/09	In progress	In progress

# JPsi Cluster Usage (Dec 08 -> May 09)

## Fermilab JPSI Cluster Utilization Chart

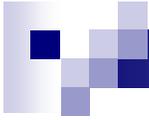
From:  (yyyy-mm-dd) To:  (yyyy-mm-dd) By:



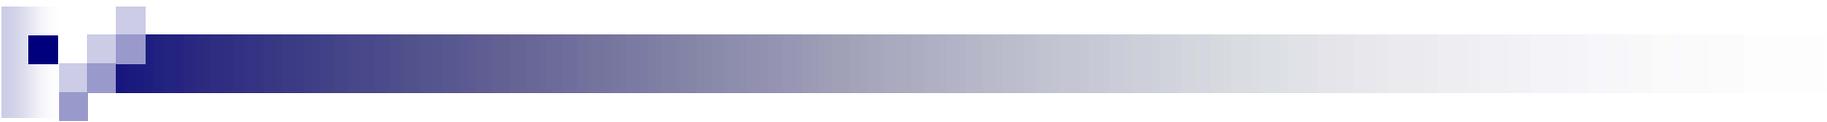
Explanation of down-time here...

Explanation of down-time here...

One heavy user completed a large job campaign, causing cluster to become temporarily vacant.



# Performance Measures and Metrics

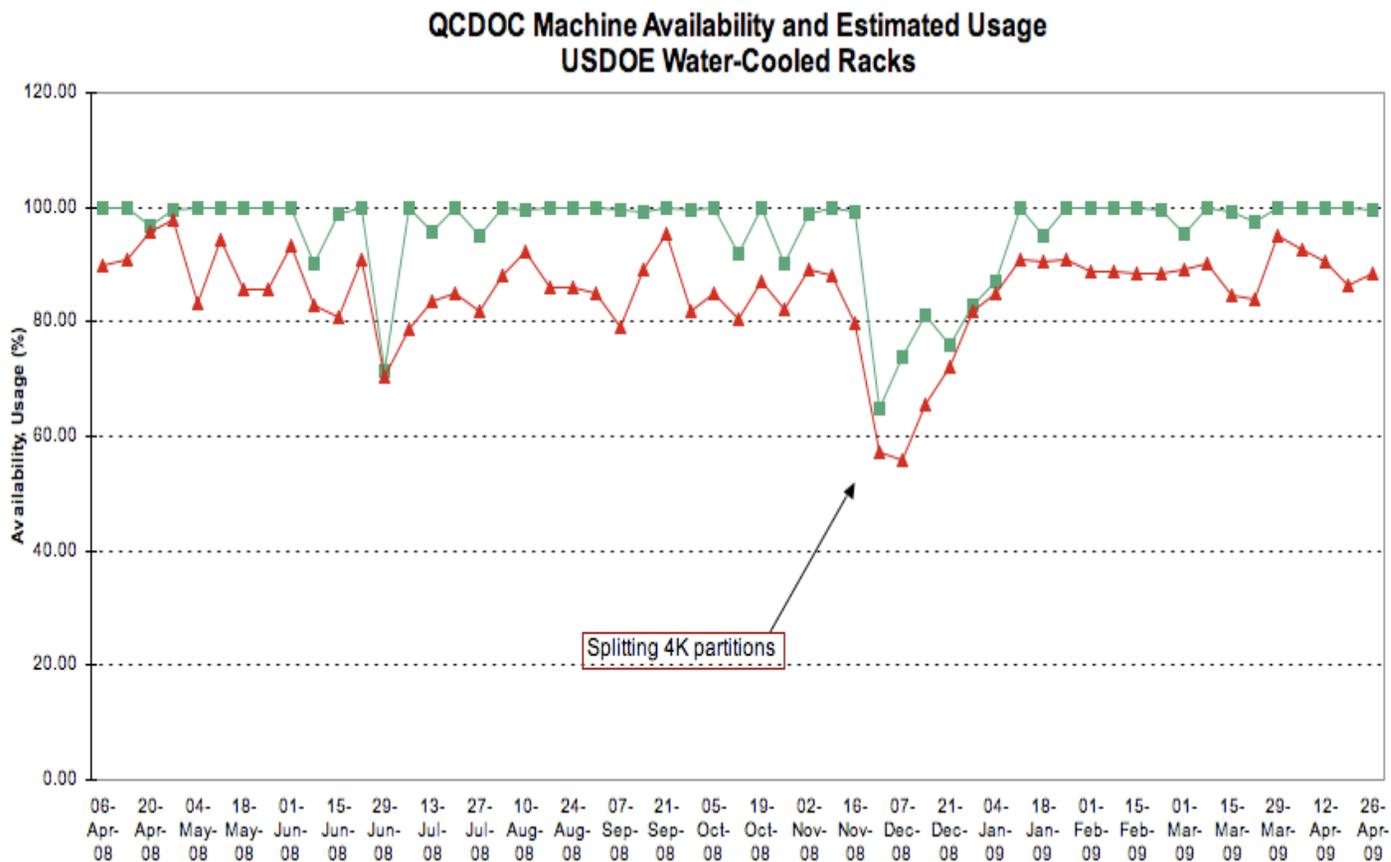


## LQCD Hardware Performance Data

- Performance and utilization data are available online for LQCD resources at all three sites (BNL, JLab, and FNAL)
  - QCDOC at BNL: <http://lqcd.bnl.gov/comp/usage/>
  - 6n and 7n at JLab: <http://lqcd.jlab.org/>
  - QCD, Pion, Kaon, and JPsi at 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.
  - Job data
    - Project allocation usage, jobs running and in queue, nodes allocated to projects.

# BNL QCDOC Utilization

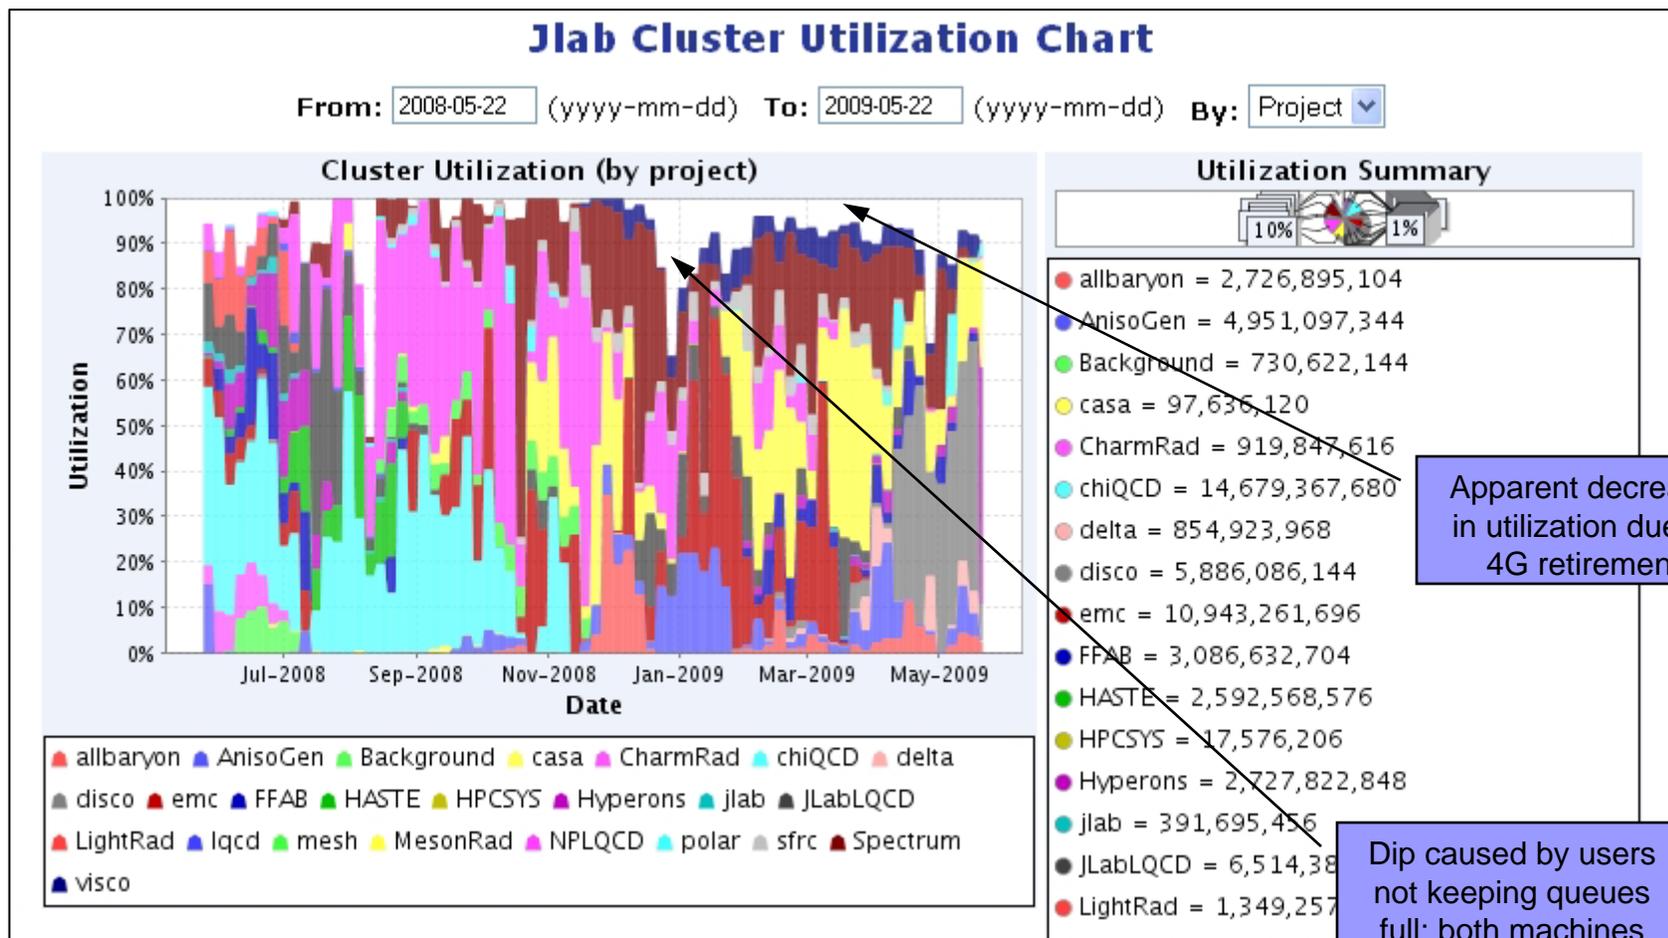
- Period of performance: April 2008 through April 2009



Green = availability; Red = estimated usage.

# JLab Cluster Utilization (6n, 7n, and 4G through 2008)

- Period of performance: May 22, 2008 through May 22, 2009

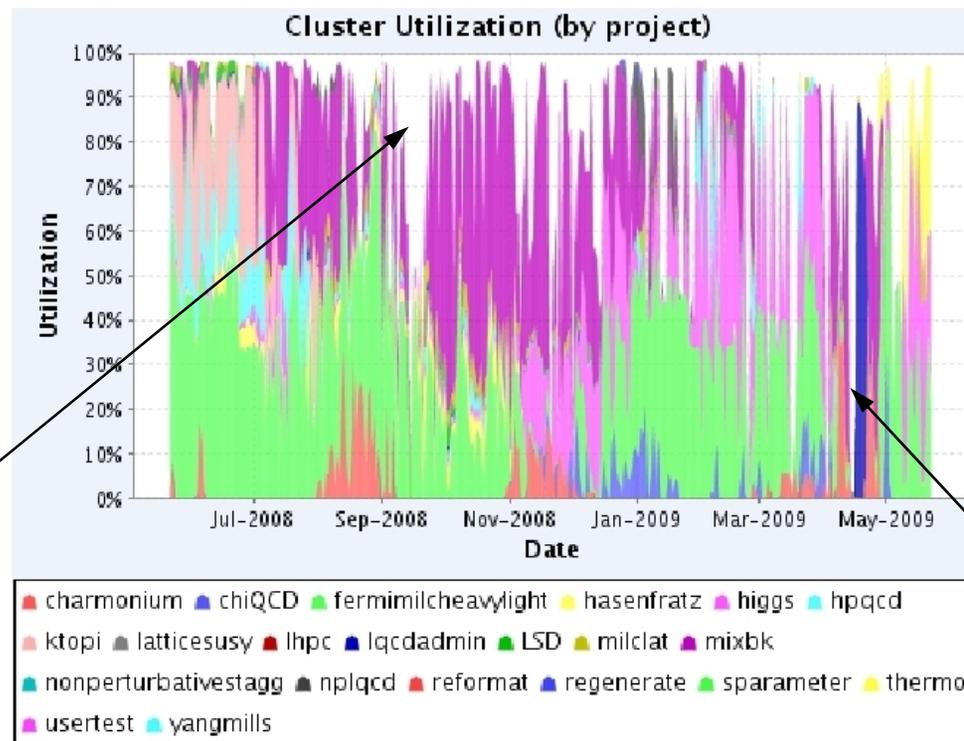


# FNAL Cluster Utilization (Kaon & Pion)

- Period of performance: 5/22/2008 through 5/22/2009

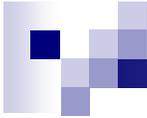
## Fermilab KAON & PION Cluster Utilization Chart

From:  (yyyy-mm-dd) To:  (yyyy-mm-dd) By:



Emergency power-down of all equipment at LCC as a result of heavy rains

Day-long power outage at LCC computing facility

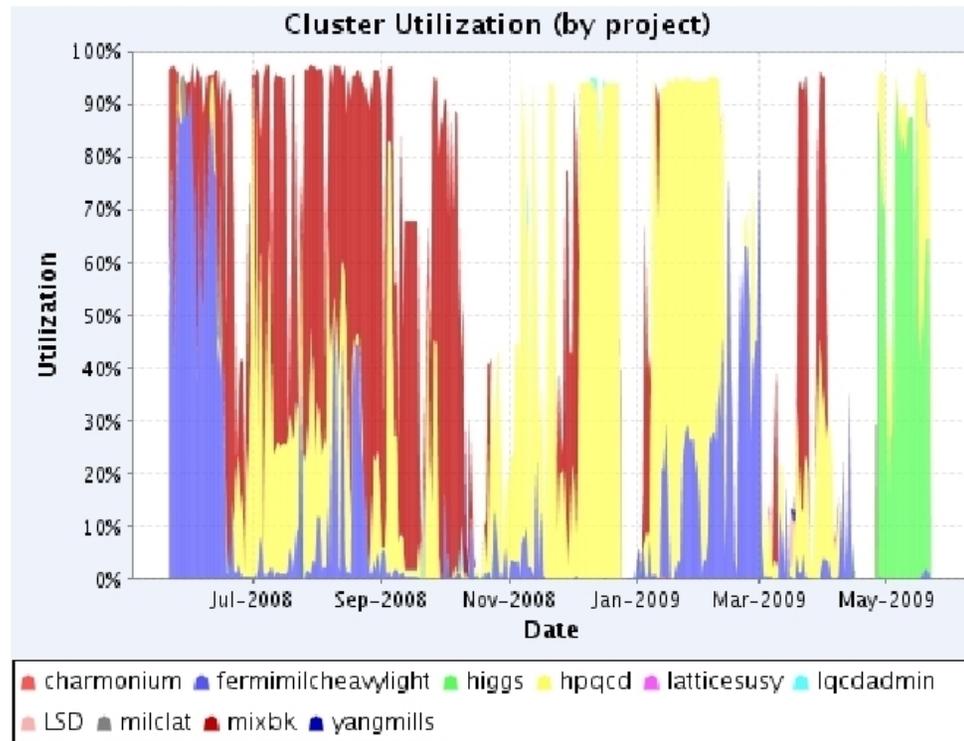


# FNAL Cluster Utilization (QCD)

- Period of performance: 5/22/2008 through 5/22/2009

## Fermilab QCD Cluster Utilization Chart

From:  (yyyy-mm-dd) To:  (yyyy-mm-dd) By:  ▾



## Utilization of Available Capacity

- Average capacity utilization across the metafacility over the period May 1, 2008 through April 30, 2009.

Machine	Capacity (TF/s)	Average Utilization (%)
QCDOC	4.20	85%
6n	0.30	92%
7n	3.00	95%
QCD	0.15	68%
Pion	0.86	86%
Kaon	2.56	96%
JPsi	8.40	92%

Average utilization calculated by dividing the total hours delivered to jobs as obtained from PBS accounting files by the measured number of node hours available on each cluster for the period May 1 2008 - April 30 2009. Measured available node hours is obtained by recording on 10-minute intervals the # of nodes in each cluster available to user jobs.



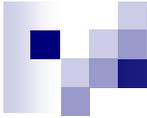
## e300 Performance Measures and Metrics

- Performance goals and milestones are explicitly defined in the OMB Exhibit 300 document.
  - 17 project milestones
    - External reviews of future procurement plans
    - Incremental procurements/Tflops-deployed
    - Aggregate Tflops-yrs delivered
  
  - 39 performance indicators
    - Science goals
    - Additional computing resource brought on-line
    - System performance (i.e., % of time system available for work)
    - Process improvements (i.e., % of tickets closed within 2 business days)
  
- Progress against these goals is tracked and reported periodically to the Federal Project Manager and through the OMB reporting process.



# Computing Performance Measures and Metrics

- Deployment and cumulative performance milestones defined for each year:
  - “Deployed Tflops”
    - Defined as incremental capacity brought on-line, expressed as average of DWF and asqtad inverter performance
  
  - “Delivered Tflops–yrs”
    - Defined as available capacity expressed as average of DWF and asqtad inverter performance
    - “1 year” = 8000 hours



## Milestone Performance *(Tflops deployed to date)*

<u>Year</u>	<u>Tflops Deployed</u>	
	<u>Baseline</u>	<u>Actual</u>
FY2006	2.0 <i>1.8 Tflops at FNAL 0.2 Tflops at Jlab</i>	2.6 <i>FNAL Kaon: 2.3 JLab 6N: 0.3</i>
FY2007	2.9	2.98 <i>JLab 7N</i>
FY2008	4.2	5.75 <i>FNAL JPsi</i>
FY2009	2.0	2.65 <i>FNAL JPsi</i>
<b>Total</b>	<b>11.1</b>	<b>13.98</b>

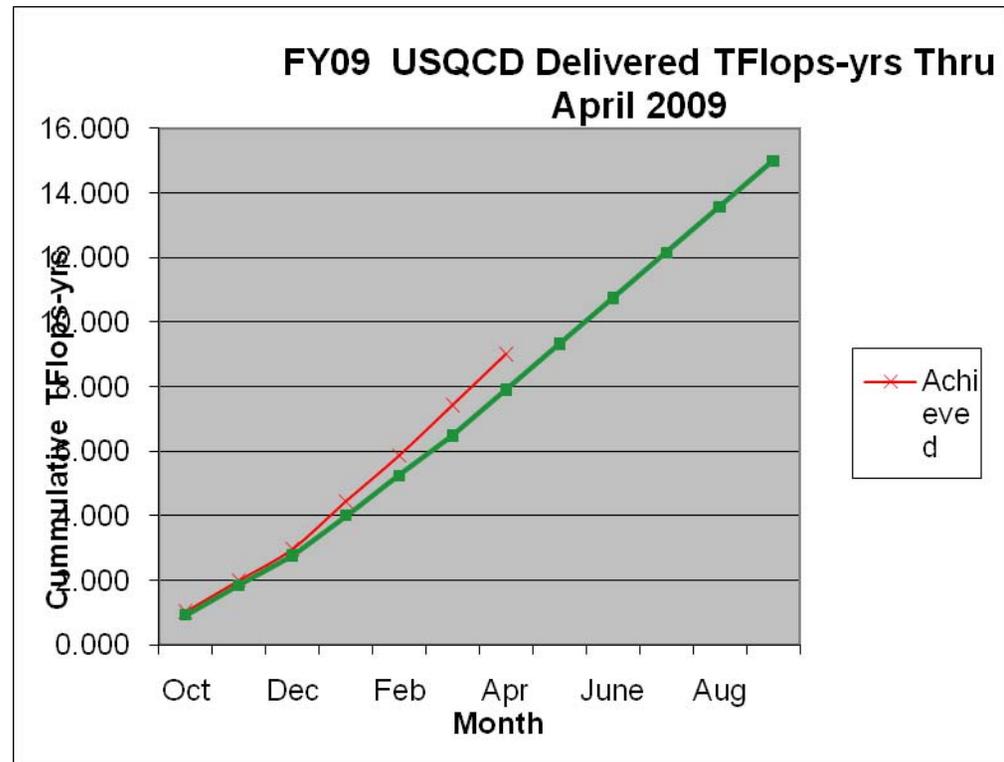
# Milestone Performance (*Tflops-yrs delivered*)

## ■ FY08

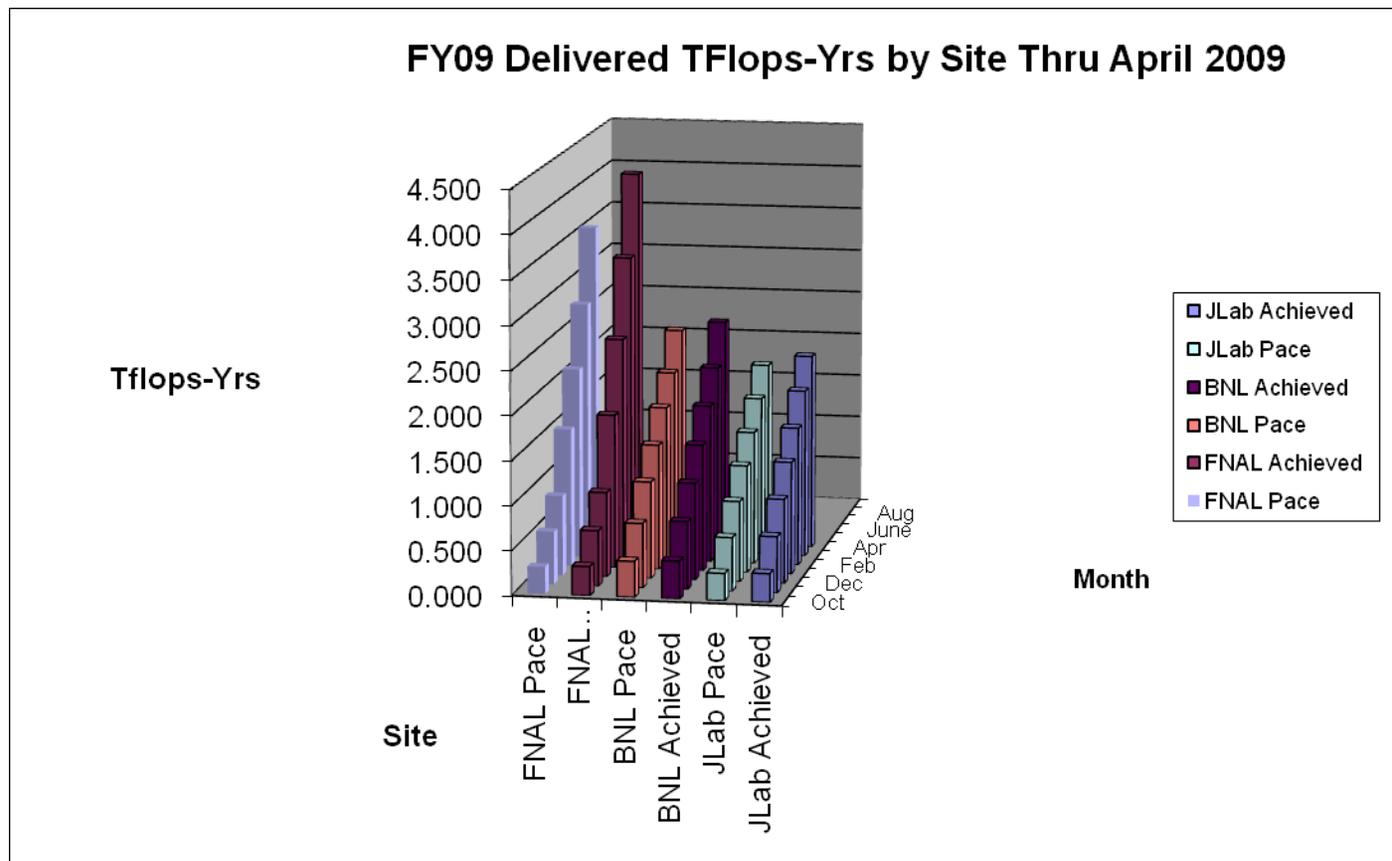
- FY08 performance goal = 12.0 Tflops-yrs delivered
- Total delivered = 12.07 Tflops-yrs (*100.6% of goal*)

## ■ FY09

- FY09 performance goal is 15 Tflops-yrs
- Planned pace goal through April is 7.90 Tflops-yrs
  - Planned pace takes into account JPsi coming online in two pieces (Jan and Apr '09)
- Through April, SC LQCD has delivered 9.01 Tflops-yrs (114% of goal)
- Actual performance data through April 2009 are shown to the right



# Delivered Tflops-Yrs by Site – FY09 Performance



## FY08 Performance Goals and Milestones

- Annual performance goals & milestones defined in OMB Exhibit 300 document include:

<i>Item</i>	<i>FY08 Goal</i>	<i>Actual</i>
Deployed Tflops	4.1	5.8*
Delivered Tflops-yrs	12.0	12.1
% machine uptime (weighted average by capacity)	93%	96%
% helpdesk tickets closed within 2 business days	92%	96%
Frequency of cyber security vulnerability scans	Monthly	Daily / wkly
Number of distinct users	30	66
Customer satisfaction rating	87%	91%

*\* FY08 deployment actually occurred in early FY09, due to planned deployment across FY08/09 boundary*

- We exceeded all performance goals in FY08
- Performance is monitored through monthly stakeholder calls, quarterly DOE OCIO progress reports, and annual progress reviews
- LQCD Project continues to receive “green” scores on quarterly reports

# Status of Progress towards FY09 Technical Performance Indicator Metrics

<i>Measurement Indicator</i>	<i>Performance Goal</i>	<i>Performance Results (through Apr 2009)</i>	
% of improvement in customer satisfaction rating (on a scale of 1 to 10)	Increase from 91% to 96% <i>(Additional 5% improvement over FY08 survey rating)</i>	TBD (Survey will be conducted in Aug/Sep 2009)	(TBD)
Number of distinct users of the facility (includes DOE labs, LQCD and academic communities)	Increase to 70	Total = 79 <i>(Through Apr '09: FNAL=44; JLab = 23; BNL=12)</i>	(+)
% of helpdesk tickets closed within 2 business days	Increase from 92% to 95%	94%	(+)
% of average machine uptime at the Metafacility	Increase from 93% to 95%	95% <i>Through Apr '09: BNL: 94.1%; FNAL: 95.7% JLab: 96.3% Capacity-weighted average = 95.4%</i>	(+)
TF-yrs delivered towards completion of the FY09 scientific program	Increase to 15.0 TF-yrs	Actual thru Apr = 9.01 TF-yrs <i>(FY09 forecast = 18.02 TF-yrs if pace continues)</i>	(+)

# Status of Progress towards FY09 Technical Performance Indicator Metrics (2)

<i>Measurement Indicator</i>	<i>Performance Goal</i>	<i>Performance Results (through Mar '09)</i>	
% of delivered node hours consumed by jobs with an exit error status	8% (Additional 10% reduction from FY08 result of 9%)	11% BNL: QCDOC does not permit the use of exit codes for this purpose	(-)
Aggregate computing resources provided by the project	Increase from 11.9 Tflops to 15.7 Tflops	17.55 (Includes the deployment of JPsi and retirement of 4G)	(+)
Increase frequency of vulnerability scans on nodes visible from Internet	Increase from monthly to bi-weekly	<u>FNAL</u> : Daily automated scans by site security staff. Log files read daily by LQCD staff. <u>JLab</u> : Daily vulnerability scans on all externally-facing systems. Access Control Lists (ACLs) have been tightened on the various enclaves. <u>BNL</u> : Daily vulnerability scans on all externally-facing systems. Log files collected daily and read by cyber-security staff.	(+)

- Performance against metrics is monitored through monthly stakeholder calls, quarterly DOE OCIO progress reports, and annual progress reviews



# Financial Performance

## FY2008 Cost Performance *(Final)*

- Period of Performance *(Oct-07 through Sep-08)*

	<u>Operating Funds</u>	<u>Equipment Funds</u>	<u>Total</u>
Budget			
FY07 Carry-Forward	\$ 34K	\$ 274K	\$ 308K
FY08 Budget	<u>\$ 930K</u>	<u>\$ 1,570K</u>	<u>\$ 2,500K</u>
Total Avail. Funds	\$ 964K	\$ 1,844K	\$ 2,808K
Actual Final Costs			
	\$ 827K	\$ 244K	\$ 1,071K
% of budget	86%	13%	38%
% of yr complete	100%	100%	100%

- Operating fund expenditures below budget because effort required to support and maintain QCDOC was much less than anticipated.

-Equipment costs below budget because FY08 cluster procurement was obligated in late FY08 but not costed until early FY09. Actual cluster cost was within planned budget.

-All unspent FY08 funds were carried forward into FY09.

## FY2009 YTD Cost Performance *(through Apr 2009)*

- Period of Performance *(Oct-08 through Apr-09)*

	<u>Operating Funds</u>	<u>Equipment Funds</u>	<u>Total</u>
Budget			
FY08 Carry-Forward	\$ 138K	\$ 1,599K	\$ 1,737K
FY09 Budget	<u>\$ 1,022K</u>	<u>\$ 678K</u>	<u>\$ 1,700K</u>
Total Avail. Funds	\$ 1,160K	\$ 2,277K	\$ 3,437K
Actual Costs			
	\$ 618K	\$2,091K	\$ 2,709K
% of budget	53%	92%	79%
% of yr complete	58%	58%	58%

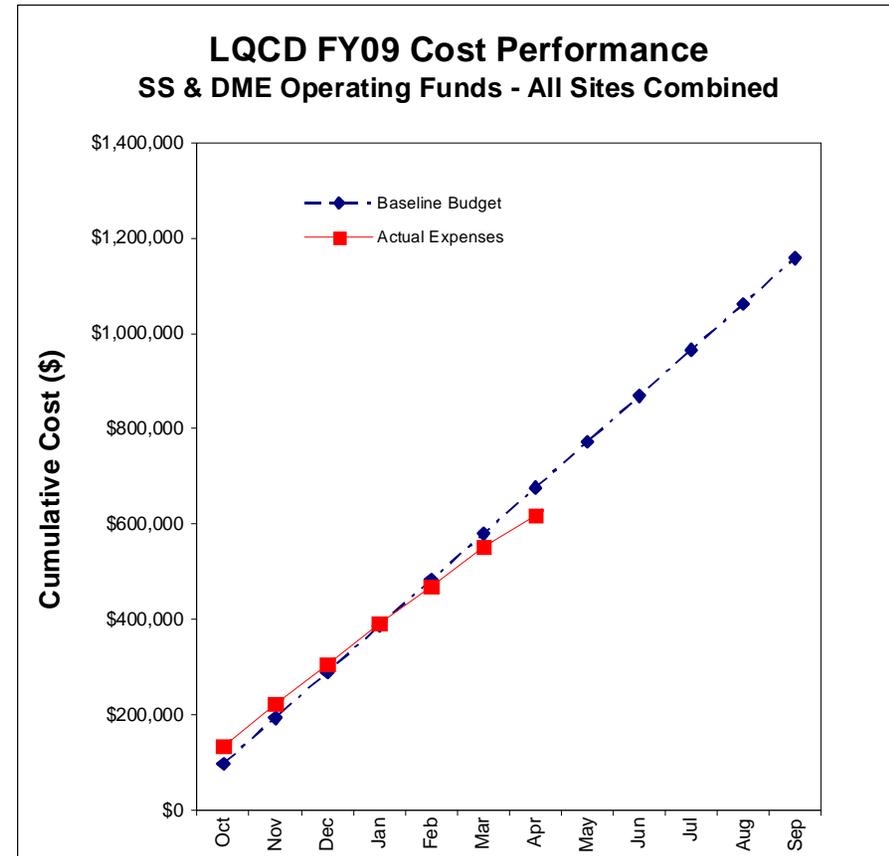
-Equipment costs to date are associated with FY08/09 JPsi procurements. Remaining EQ funds will be used to purchase additional storage hardware (e.g., disk arrays and associated servers).

- Since EQ funds are spent on large lump-sum expenditures, we expect a very non-linear EQ spend rate as seen in the above table. We do not anticipate exceeding the EQ budget.

# FY09 Spend Rate

## Operating Funds – All sites combined

- Operating funds support salary costs associated with project management; procurement and deployment of new hardware; and operation and maintenance of existing systems.
- Spend rate is tracked through monthly site accounting reports
  - Spend rate for steady-state operations support is fairly constant
  - Site-specific spend rates increase during acquisition & deployment activities
- Combined spend rate through April 2009, for all three sites, is trending slightly downward compared to a linear baseline forecast.
  - Level of support required at BNL to support QCDOC has been significantly less than anticipated
  - Salary costs for deployment activities at FNAL are slightly above budget due to carry-over of FY08 deployment activities. SS spend rate in line with budget.
  - Spend rate at JLab approaching linear baseline forecast. Higher at start of year due to one-time server purchase..

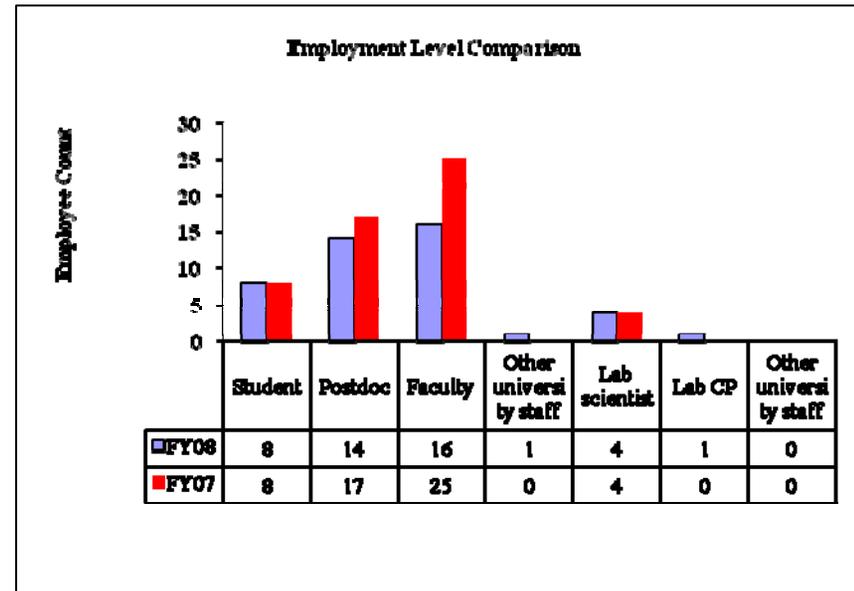
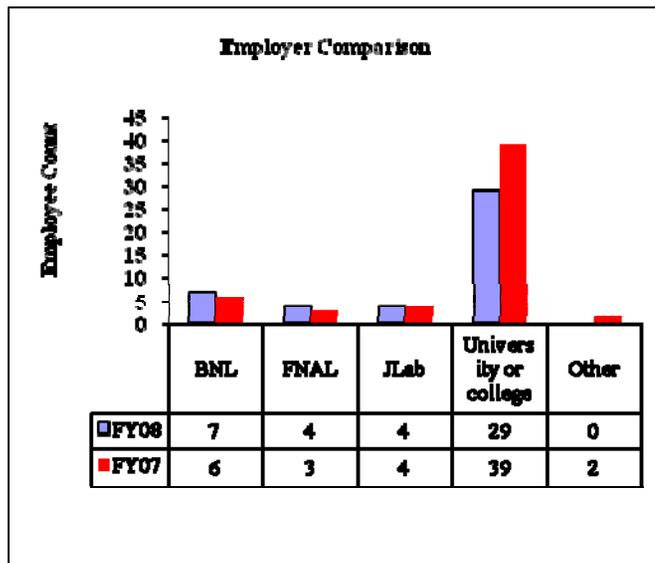




# User Survey Results

# FY08 User Survey: General Profile

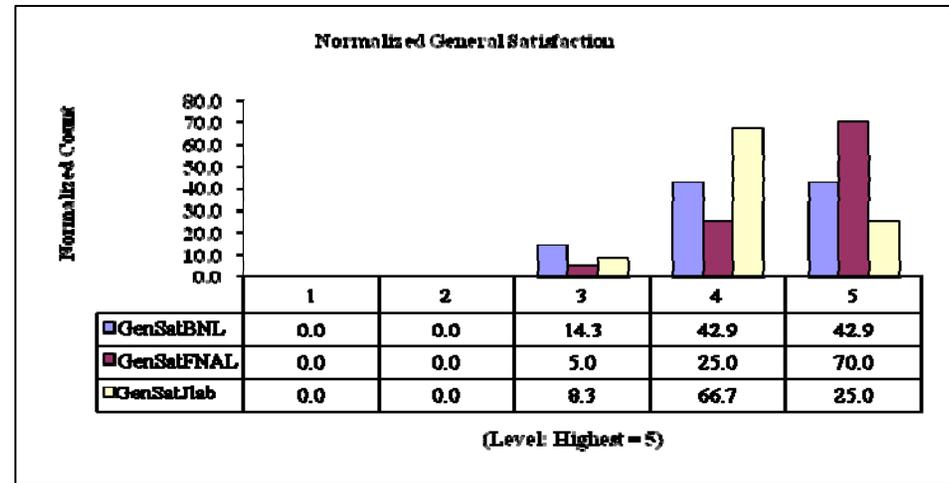
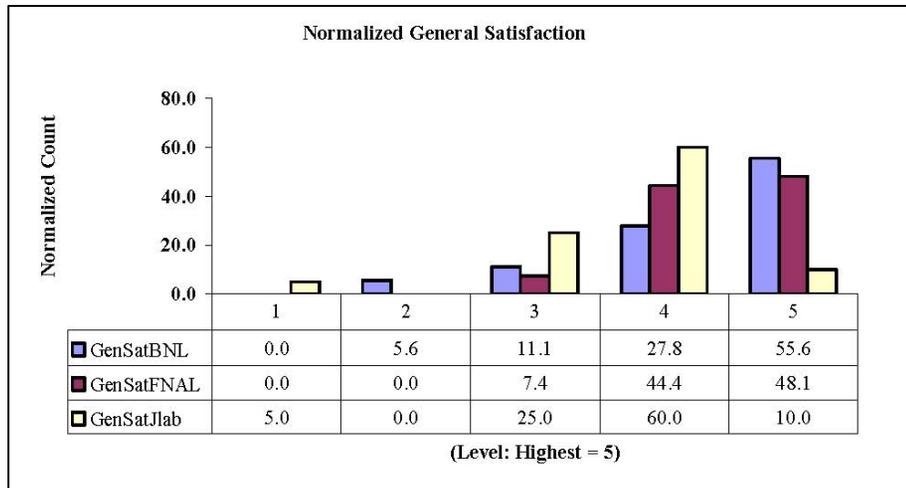
- The second LQCD user survey was conducted in the fall of 2008. Covered a total of 24 areas, including overall satisfaction with computing facilities, documentation, user support, and the proposal / resource allocation process.
- Received responses from 44 users, less than the 56 responses from the FY07 survey.



# User Survey: FY08 General Satisfaction - 91%

**FY07 Result = 82%**

**FY08 Result = 91%**



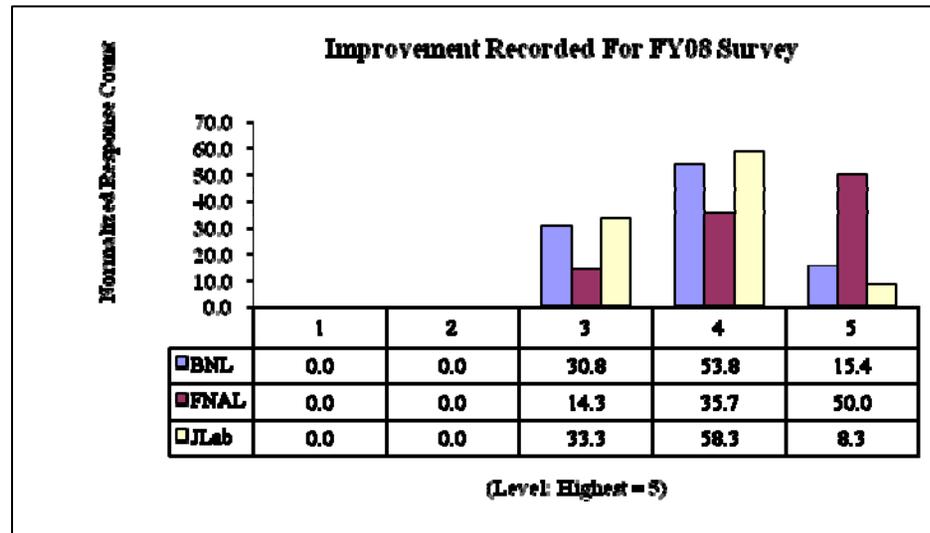
## Positive Comments

- Support staff is very responsive when there is a problem – even during off-hours.
- A few times, admins helped me improve my utilization without even being asked; much appreciated.
- I have used QCDOC for several years and am very pleased with administration of said computers.
- I only have the highest praise for the FNAL LQCD operation. Outstanding.
- The staff at JLab are helpful in attending to our needs.

## Constructive Criticism

- The only thing which is very annoying is the cyber-security related trainings. It is really too much for the users.
- I find the security at BNL and JLab inconvenient and no longer use them.
- Sometimes, my program takes more time than usual to finish running; suspects a problem with the nodes.

# User Survey: FY08 Overall Improvement



## Positive Comments

- Things improved as machines stabilized.
- Not much has changed at FNAL over the last year. Things are still running smoothly.
- Things ran well last year and continue to do so this year.
- My satisfaction for the FNAL operation remains very high.

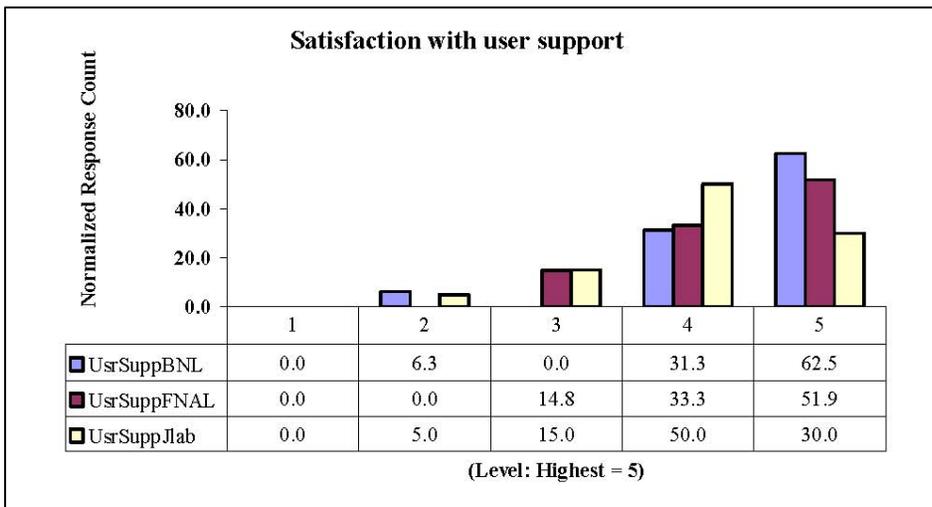
## Constructive Criticism

- The choices don't seem to match the question.

# User Survey: FY08 User Support

Satisfaction with User Support increased from 85% in FY07 to 100% in FY08

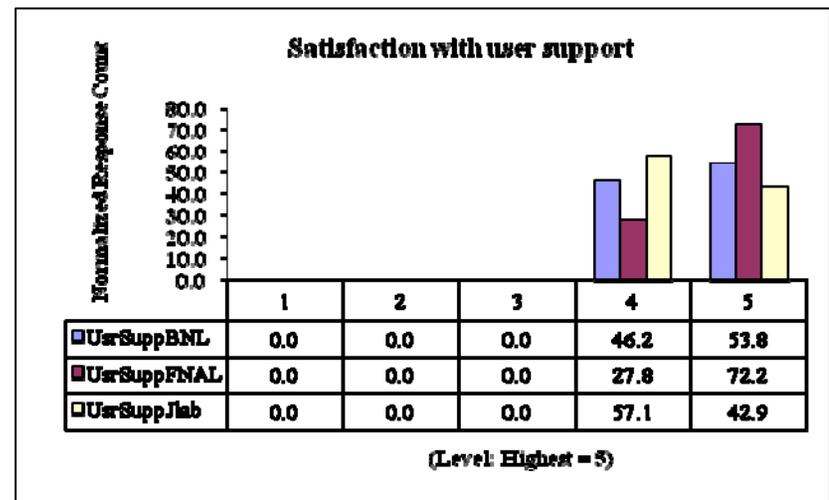
## FY07 Result



### Positive Comments

- Very good help.
- Support personnel at FNAL have been very responsive.
- Very satisfied with the admins at JLab.

## FY08 Result



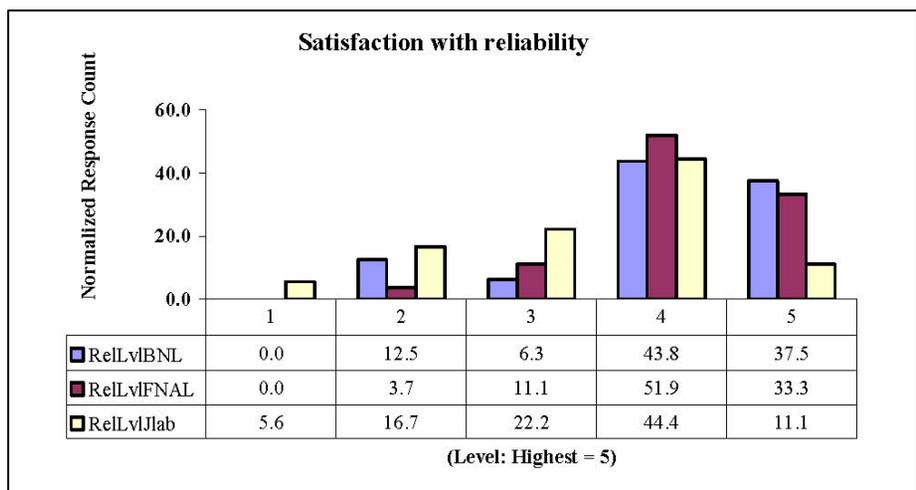
### Constructive Criticism

- No comments entered by survey respondents

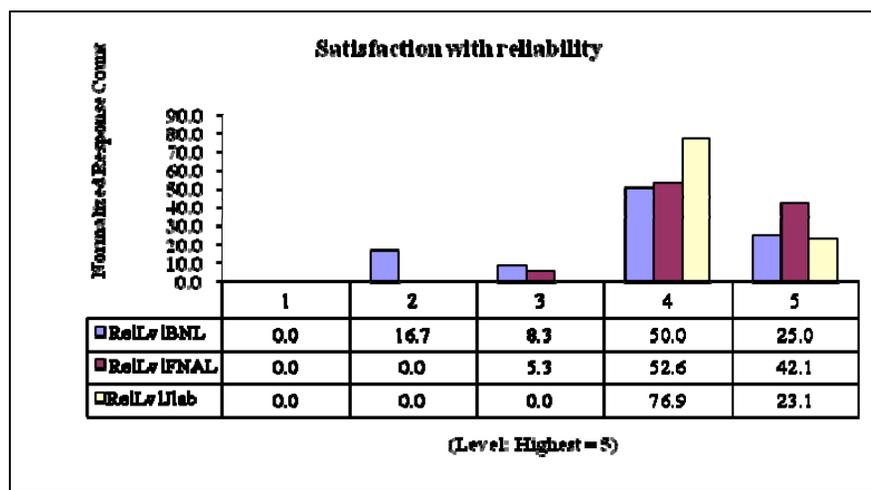
# User Survey: FY08 Reliability

Overall machine reliability increased from 75% in FY07 to 91% in FY08

## FY07 Result



## FY08 Result



### Positive Comments

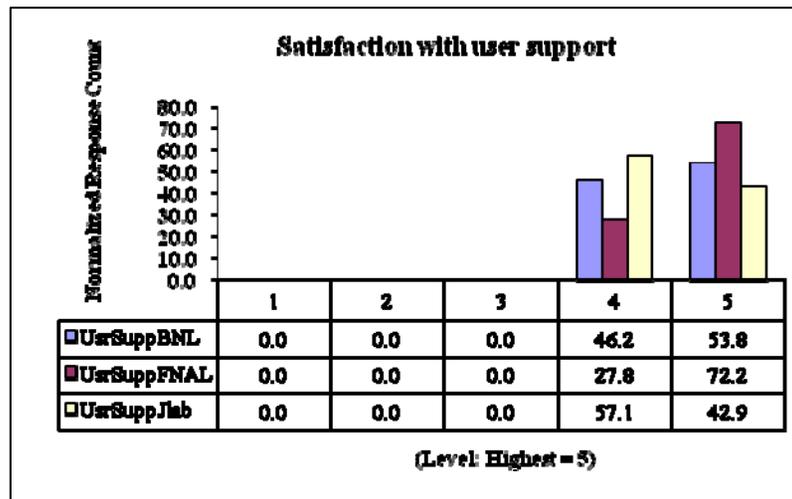
- No free-form comments were submitted by survey respondents

### Constructive Criticism

- I have a few more failed jobs that I would like, which seem to stem from bad nodes that are still being used. Have been told that nodes are often oversubscribed for memory, which can cause problems. Nodes stay in compute queue and don't automatically get rebooted.
- QCDOC machines are unreliable and prone to frequent hardware failure.

# User Survey: FY08 Responsiveness

Staff responsiveness increased from 89% in FY07 to 98% in FY08



## Positive Comments

- The JLab staff is very quick to respond as well as very helpful in explaining the current issues and giving timelines for getting up to speed again.
- Exceptionally nice and helpful staff.
- Stratos has done a great job. Bob Mawhinney has also done very well managing the different users and allocations.
- Outstanding support from staff at FNAL.
- The staff are stellar.

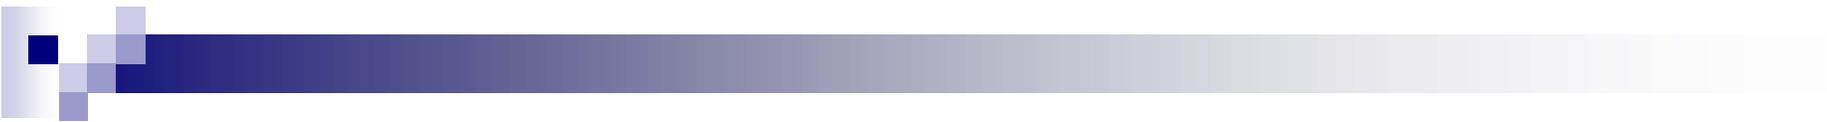
## Constructive Criticism

- Occasionally we're left a bit in the dark (or sent jargon-full e-mails that I don't understand) when volatile dCache or tape storage is down.



# FY08 User Survey Summary

- Overall, the FY08 year-end user survey resulted in a number of positive statements and constructive suggestions.
- General satisfaction level has increased from 82% in FY07 to 91% in FY08
- Users remain satisfied with helpdesk response (96% in FY07; 97% in FY08)
- Users agreed that several improvements occurred during FY08
- User satisfaction with the “call for proposals” and allocation process has improved
- Cyber-security issues continue to frustrate users.
- There may be some opportunities for improving communication within the collaboration
  - Improved communication of committee activities and important decisions



## Project Summary

- LQCD computing project continues to run smoothly
- Our site managers continue to do a very good job of operating their respective systems to minimize downtime and maximize output.
- We have been successful in meeting our key performance goals and milestones.
- We have been successful in deploying new systems and operating our facilities within budget.
  - FY08/09 procurement was successfully executed according to the approved acquisition plan.
  - Acknowledge that the host laboratories provide significant resources (e.g., space, power, cooling, networking, etc.)
- Results from the FY08 User Survey indicate an increase in overall satisfaction level. The user survey continues to be a useful tool by providing insight into where to focus efforts to improve operational effectiveness and the user experience.



# Questions?