

Lattice QCD Computing Project

Project Management Overview

Bill Boroski
LQCD Contractor Project Manager

2008 Annual Review
Brookhaven National Laboratory
May 13-14, 2008



Outline

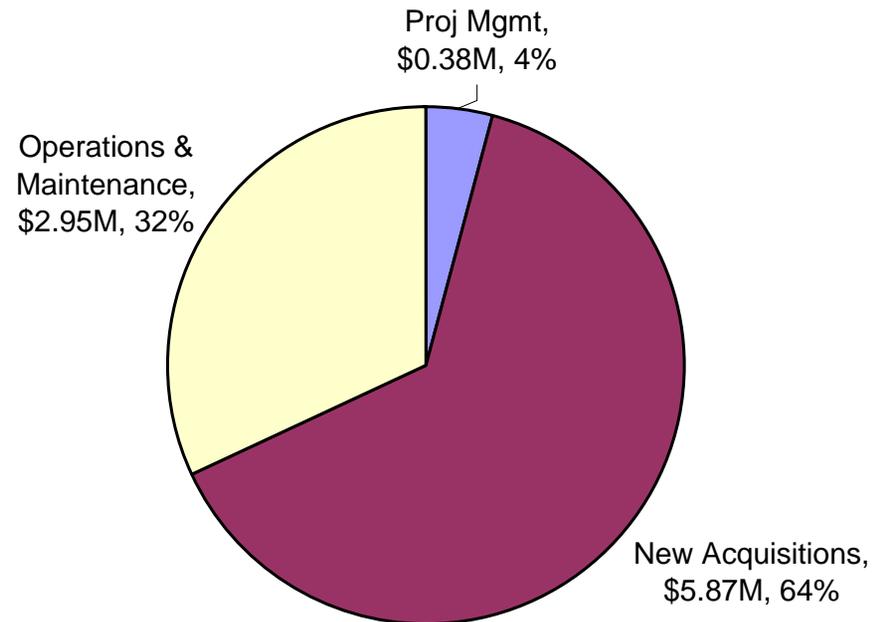
- Project Scope
- Management & Oversight Structure
- Work Organization and Planning
- System Deployments
 - FY07 Deployment Summary (7n at JLab)
 - FY08/09 Procurement Plan
- Performance Measures and Metrics
 - Technical, Scientific, Cost and Schedule
- FY07 User Survey Results

Project Scope and Budget

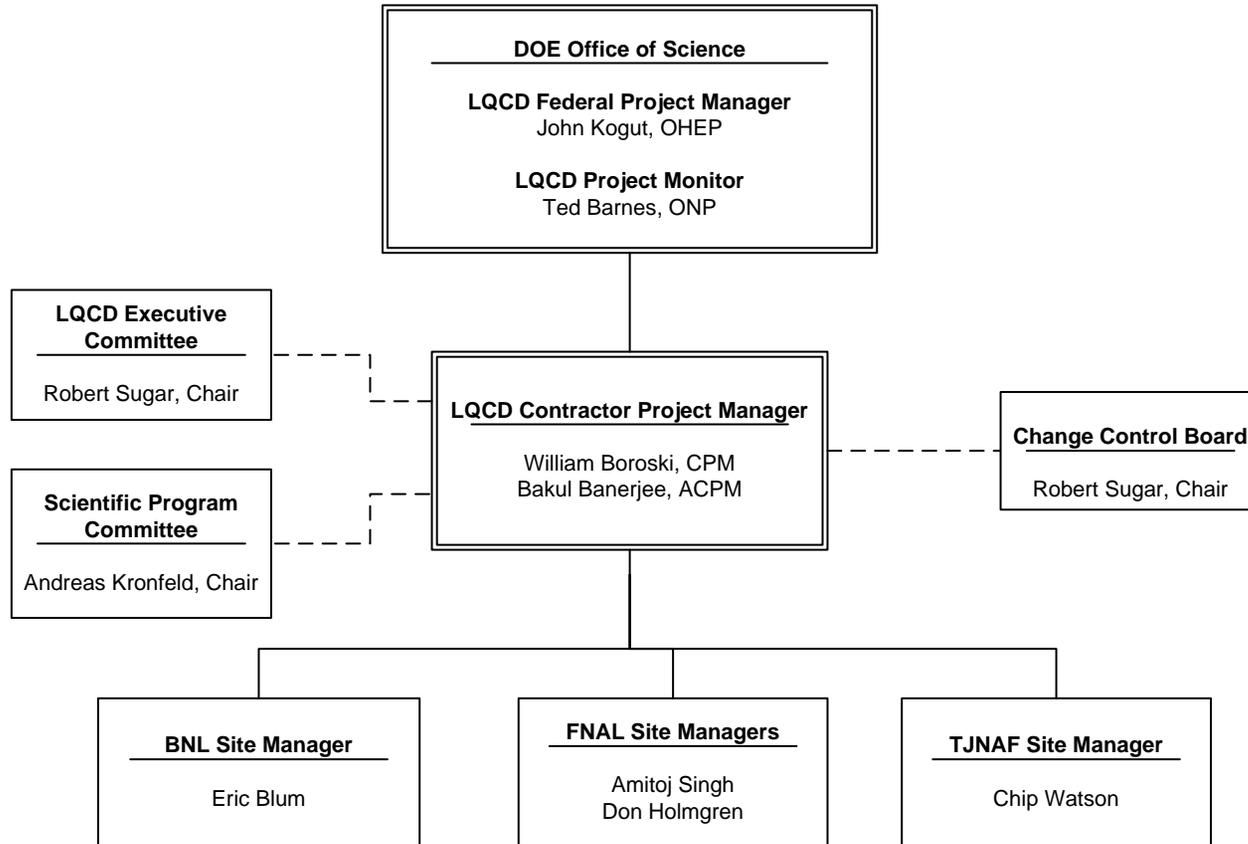
- Acquire and operate dedicated hardware at BNL, TJNAF, and FNAL for the study of quantum chromodynamics
 - Budget: \$9.2 million (*provided jointly by OHEP and ONP*)
 - Period of performance: FY06 through FY09

- Project funding covers:
 - Project management
 - Operations and maintenance of existing systems
 - Acquisition and deployment of new hardware

- Not in scope
 - Software development
 - Scientific software support



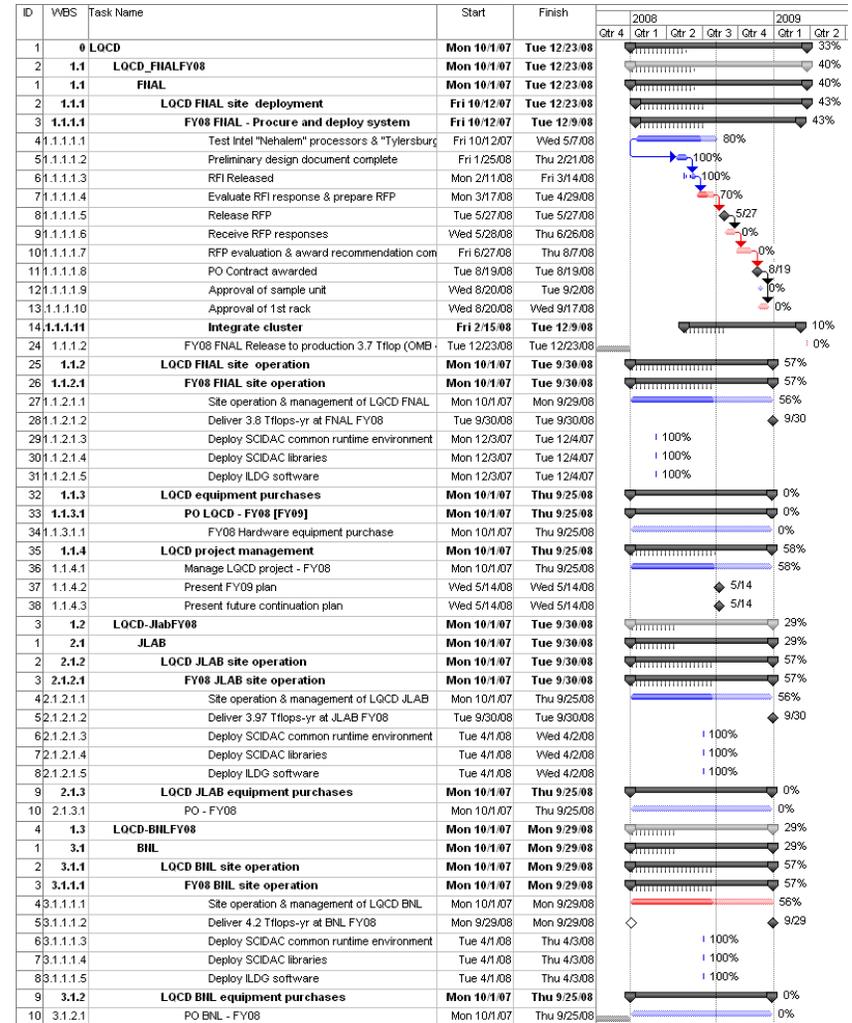
Management Organization



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

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;
 - 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 call for a major new acquisition in each year of the project.

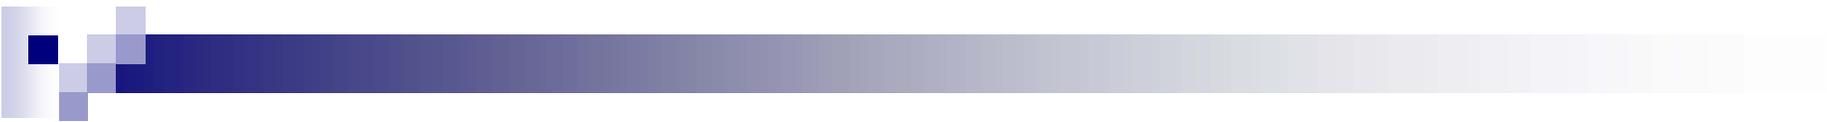
- Procurements treated as sub-projects

- Procurement and deployment plans, with timeline and milestones, are developed as part of the annual planning and budgeting process.
 - Planning takes into account performance requirements and goals, existing facility capabilities and required facility upgrades, technical advances, etc.
 - Current activities are focused on the selection, procurement and deployment of the FY08/09 cluster at FNAL



Project-funded Workforce Staffing Model

- In general, the project-funded staffing model has been:
 - Site management support: 0.25 FTE/yr per site
 - Steady-state sys admin at BNL: 0.75 FTE/yr
 - Steady-state sys admin support at JLab: ~1.0 FTE/yr
 - Sys admin support at Jlab increased to 1.65 FTE to support 7n acquisition/deployment in FY07
 - Steady-state sys admin support at FNAL: ~1.65 FTE/yr
 - Project management support at FNAL: 0.5 FTE/yr



Communications and Reporting

- **Bi-weekly 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**



Configuration Management and Change Control

- Change control process defined in the PEP
- CCB chaired by Bob Sugar, Executive Committee chair
- Membership includes members of the LQCD Executive Committee and senior management at the three labs.
- Change Control Board:
 - Bob Sugar, Chair (LQCD Executive Committee Chair)
 - Bill Boroski, LQCD Contractor Project Manager
 - Steve Gottlieb, USQCD scientific representative
 - Tom Schlagel, Head, Information Tech. Division, BNL
 - Vicky White, Head, Computing Division, FNAL
 - Roy Whitney, CIO, JLab

Configuration Management and Change Control (2)

- Change control thresholds:

Level	Cost	Schedule	Technical Scope
LQCD Federal Program Manager (Level 0)	Any increase in total project cost	3-month or more delay in Level-1 milestone date	Change of any WBS element that could adversely affect performance specifications
LQCD CCB (Level 1)	Cumulative increase of more than \$125K in WBS Level 2	> 1-month delay of a Level-1 milestone date or >3-month delay of Level-2 milestone.	Any deviation from technical deliverables that does not affect expected performance specifications
LQCD Contractor Project Manager (Level 2)	Any increase of >\$25K in WBS Level 2	>1-month delay of Level 2 milestone date	Technical design changes that do not impact technical deliverables.

- Associate Contractor Project Manager maintains change control log and records.
 - One CR processed in FY2007 – Change in FY07 cluster deployment schedule milestone.



System Deployments



FY07 Cluster Deployment Summary (7N)

- FY07 deployment goal was to procure and deploy a 2.9 Tflops system at JLab by June 30, 2007.
- The winning proposal for the 7N cluster consisted of purchasing 396 dual-core nodes, with an option to upgrade to 2.1 GHz quad-core AMD processors when they became available.
 - Cost of the upgrade was estimated at 20% per node, with a likely performance boost of 40-60%.
 - Testing showed that implementation of the quad-core upgrade was necessary to achieve our performance goal
 - With quad-core processors, we anticipated deploying a minimum of 2.9 Tflops
 - Without the quad-cores, we anticipated deploying 2.2 Tflops.
 - Early estimated ship date for quad-core production chips was Aug/Sep 2007, beyond the performance goal milestone date.
- A plan for procuring dual-core processors, and exercising the upgrade option if appropriate, was presented at the 2007 review and accepted by the review committee as reasonable.

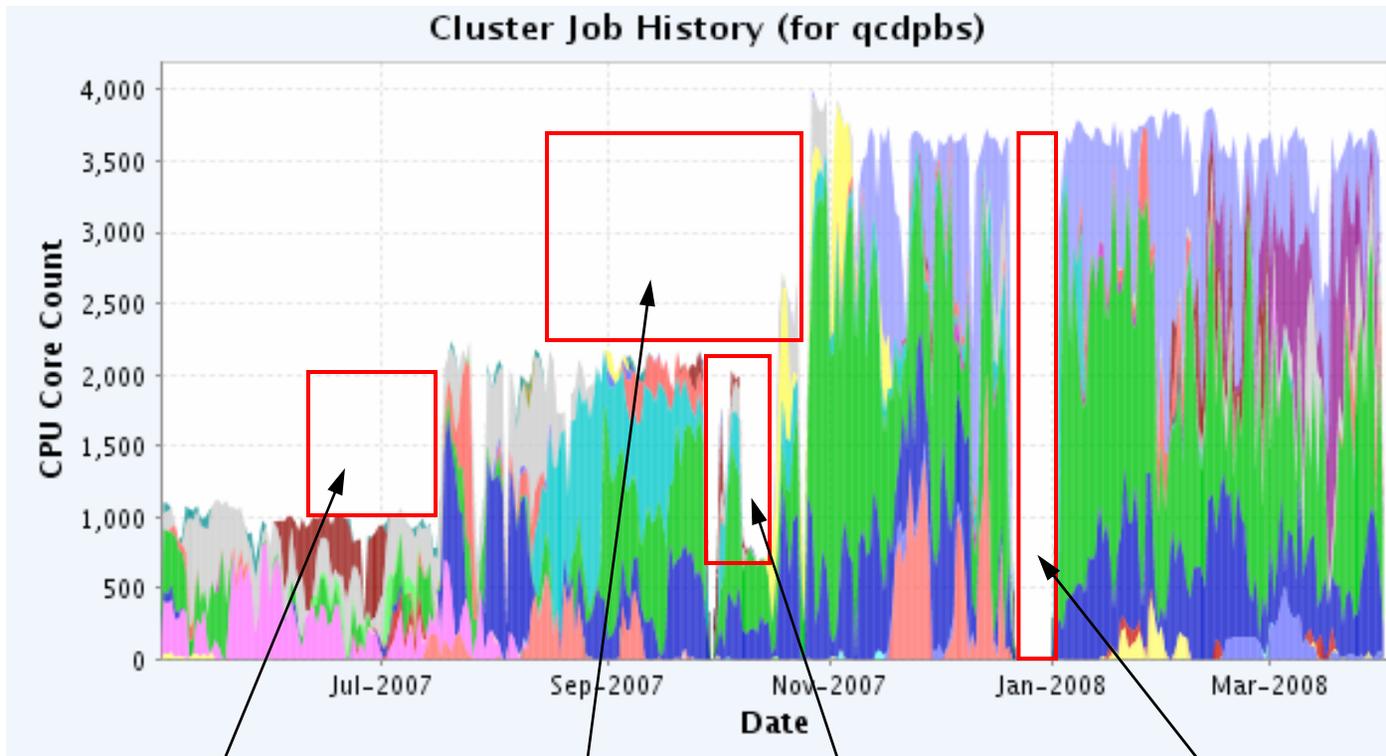


7N Deployment Summary (cont'd)

- A formal change request was processed to move back the deployment milestone date by 6 months to provide time for the quad-core upgrade.
 - The CR conditions required approval by the CCB and Federal Project Manager. Approval was received on 21-Jun-2007.
- 7N Time Line
 - July 15: Initially online with dual dual-core AMD Opterons
 - Memory bandwidth 8.5 GB/s for 4 threads (Triad); 5.17 GFlops/node
 - Total deployment: 2.04 TFlops
 - October 8-17, 2007: Upgrade to quad-core Opterons
 - Upgrade delayed by AMD release of production quad-core chips
 - Took delivery of 1.9 GHz processors, instead of planned 2.1 GHz
 - Memory bandwidth 9.1 GB/s for 8 threads (Triad); 7.95 GFlops/node
 - Total deployment: 3.15 Tflops (vs. goal of 2.9 Tflops)
 - Dec 20 to Jan 3, 2008: 20-year JLab site power maintenance
 - Feb 15 to Mar 14, 2008: Rolling memory upgrade to 8 GB
 - Memory bandwidth 10.8 GB/s

JLab Cluster Usage (July 07 -> March 08)

(4G + 6N + 7N)



Friendly user period
(dual-duals)

7N quad-core
installation

Cycles lost to AMD
delays
(0.7M 6N node-hrs)

Cycles lost to 20-year
maintenance work
(0.5m – 6n node-hrs)

FY08/09 Procurement Planning

- FY08/09 procurement planning is now underway (*Details in Don Holmgren's talk*)
- Performance goals from the OMB Exhibit 300:

FY08 Goals:	Planned Completion Date
Procure and deploy 4.2 Tflops at FNAL	12/30/08
12 Tflops-yrs aggregate computing delivered	09/30/08
FY09 Goals:	
Procure and deploy 2.0 Tflops at FNAL	06/30/09
15 Tflops-yrs aggregate computing delivered	09/30/09

- We intend to combine the FY08/09 procurements with an option clause, as presented during the 2007 review
 - Cost savings associated with reduced labor costs
 - Approach endorsed by the 2006 and 2007 Review committees



Performance Measures and Metrics



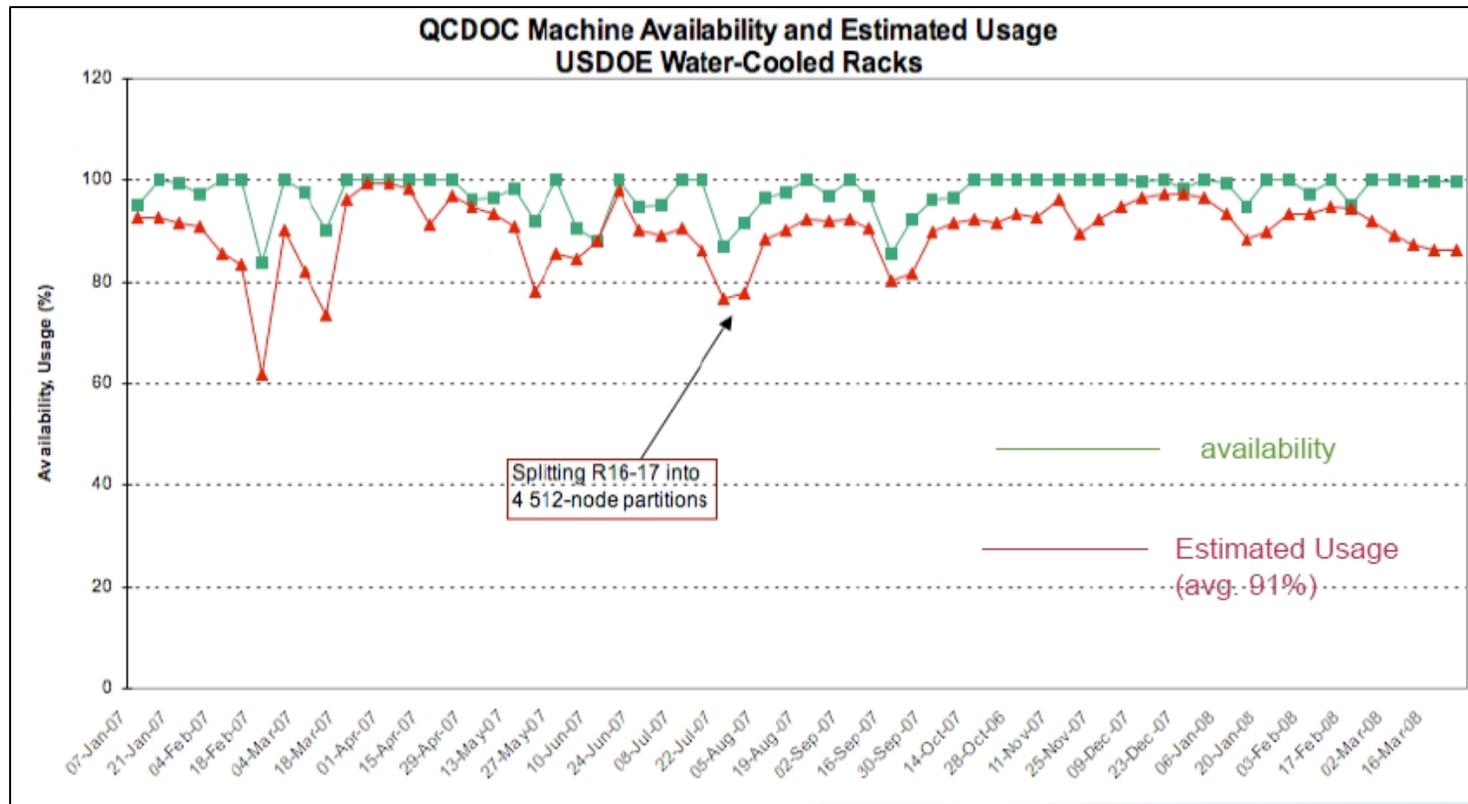
LQCD Hardware Performance Data

- 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/>
 - 4G, 6N, and 7N at JLab: <http://lqcd.jlab.org/>
 - Kaon, Pion, QCD at FNAL: <http://kaon2.fnal.gov/cluster/usage.html>

- 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.

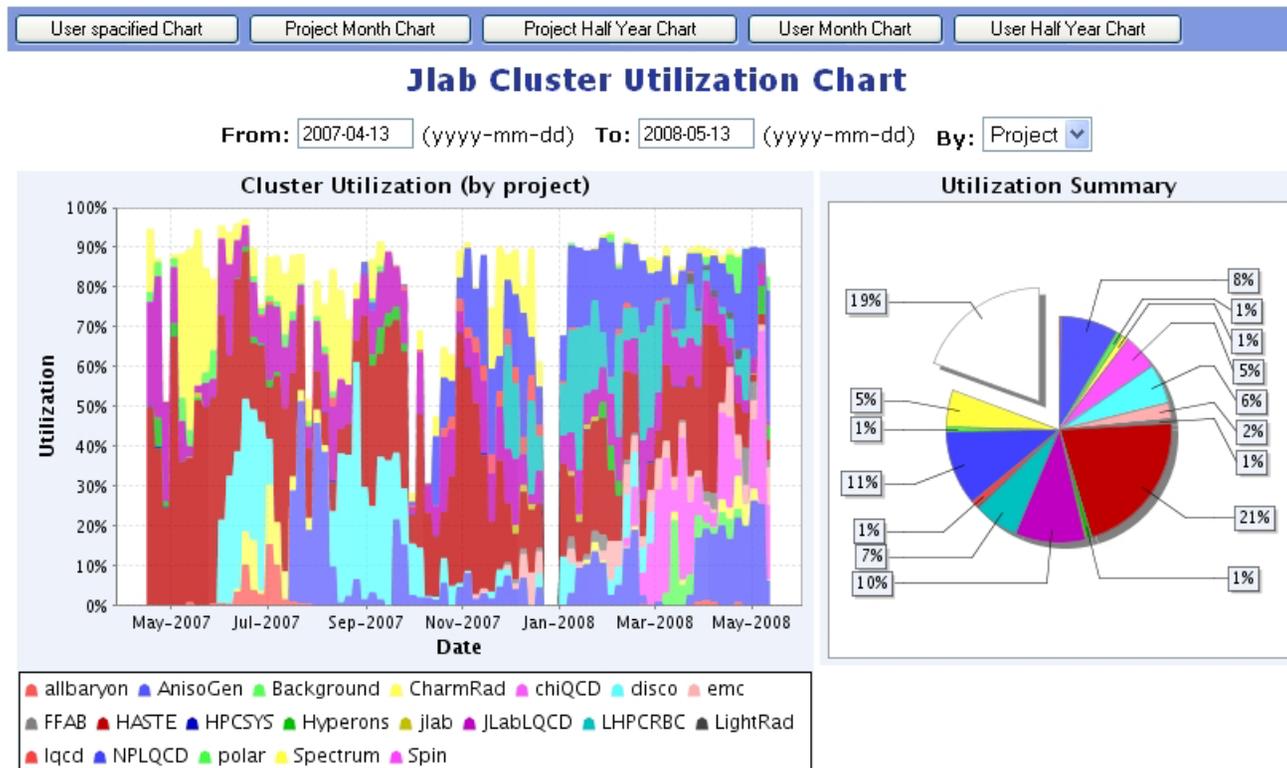
BNL QCDOC Utilization

- Period of performance: January 2007 through March 2008



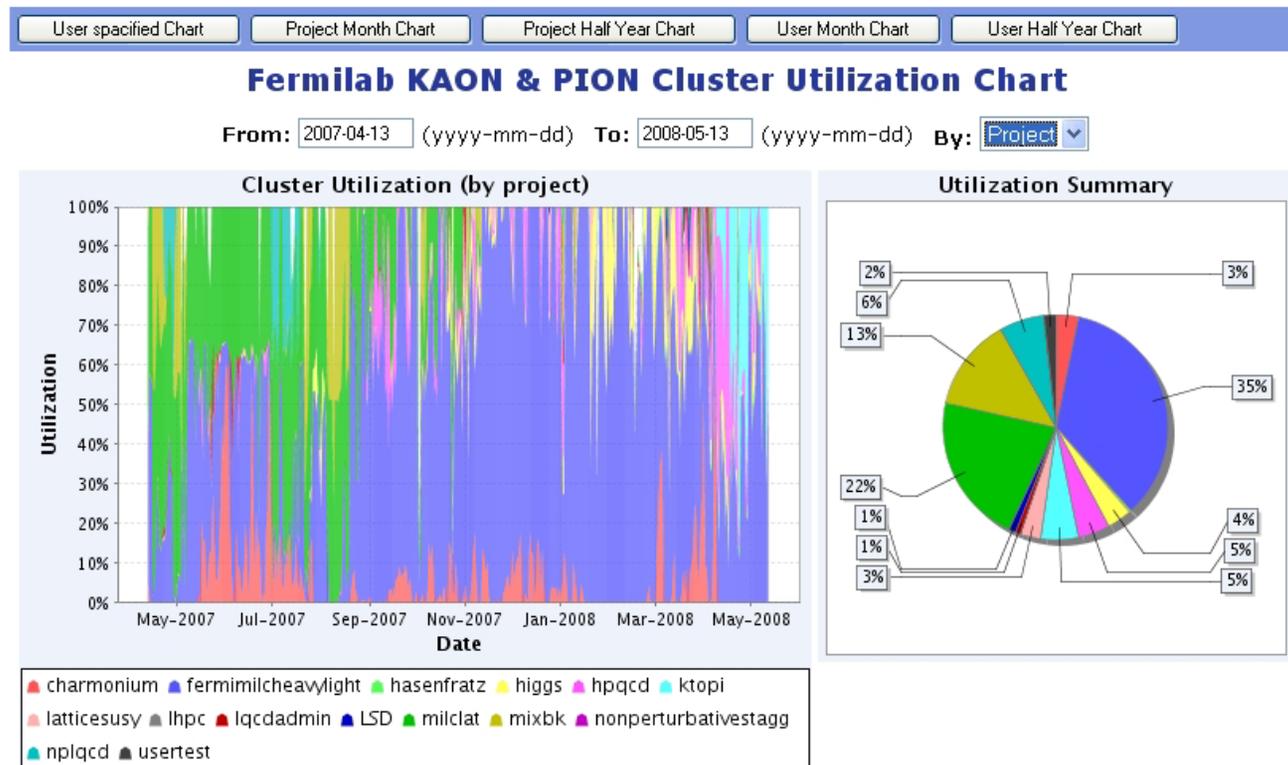
JLab Cluster Utilization (4G, 6N, 7N)

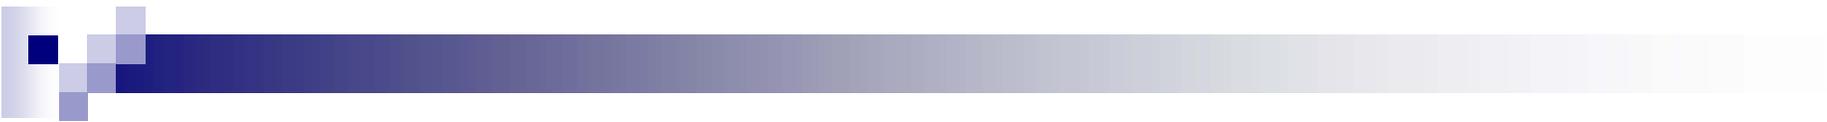
- Period of performance: 4/13/2007 through 5/13/2008



FNAL Cluster Utilization (Kaon & Pion)

- Period of performance: 4/13/2007 through 5/13/2008





FNAL Cluster Utilization (corrected)

- A caveat: the java software used to generate our utilization plots has problems when the project and user counts get too high.
- The JLab data appear to be fine, but the FNAL data are likely missing the “idle” time and perhaps some other projects.
- Corrected data obtained by dividing the number of normalized hours from Oct 1 through April 30 into total billed core hours. Normalized hours correspond to a “snowmass” year.
 - We use “snowmass years” for delivered Tflops-yrs; they correspond to a time available of approximately 91% (8000 hours / 8800 hours).
- “Corrected” FNAL utilization data
 - QCD: 97.6%
 - Pion: 88.1%
 - Kaon: 96.1%
 - Weighted average by capacity (Tflops) = 94.0%



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

 - 36 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:
 - “Delivered Tflops–yrs”
 - Defined as available capacity expressed as average of DWF and asqtad inverter performance
 - “1 year” = 8000 hours
 - “Deployed Tflops”
 - Defined as incremental capacity brought on-line, expressed as average of DWF and asqtad inverter performance

Milestone Performance (*Tflops deployed*)

	Tflops Deployed	
<u>Year</u>	<u>Current Baseline</u>	<u>Actual</u>
FY2006	2.0 <i>(FNAL: 1.8 Tflops)</i> <i>(JLab: 0.2 Tflops)</i>	2.6 <i>(FNAL Kaon: 2.3 Tflops)</i> <i>(JLab 6N: 0.3 Tflops)</i>
FY2007	2.9	2.98 <i>(JLab 7N: 2.98 Tflops)</i>
FY2008	4.2	tbd
FY2009	2.0	tbd

Cumulative FY06/07 milestone = 4.9 Tflops

Total FY06/07 actual = 5.6 Tflops

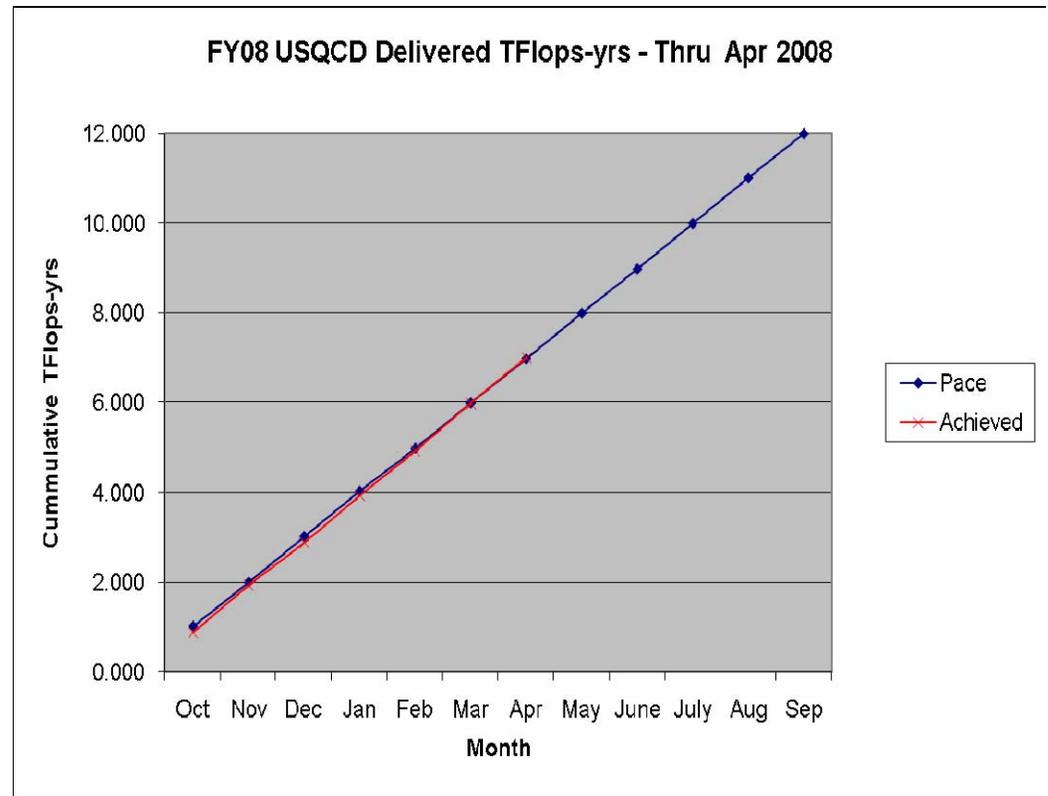
Milestone Performance (Tflops-yrs delivered)

■ FY07

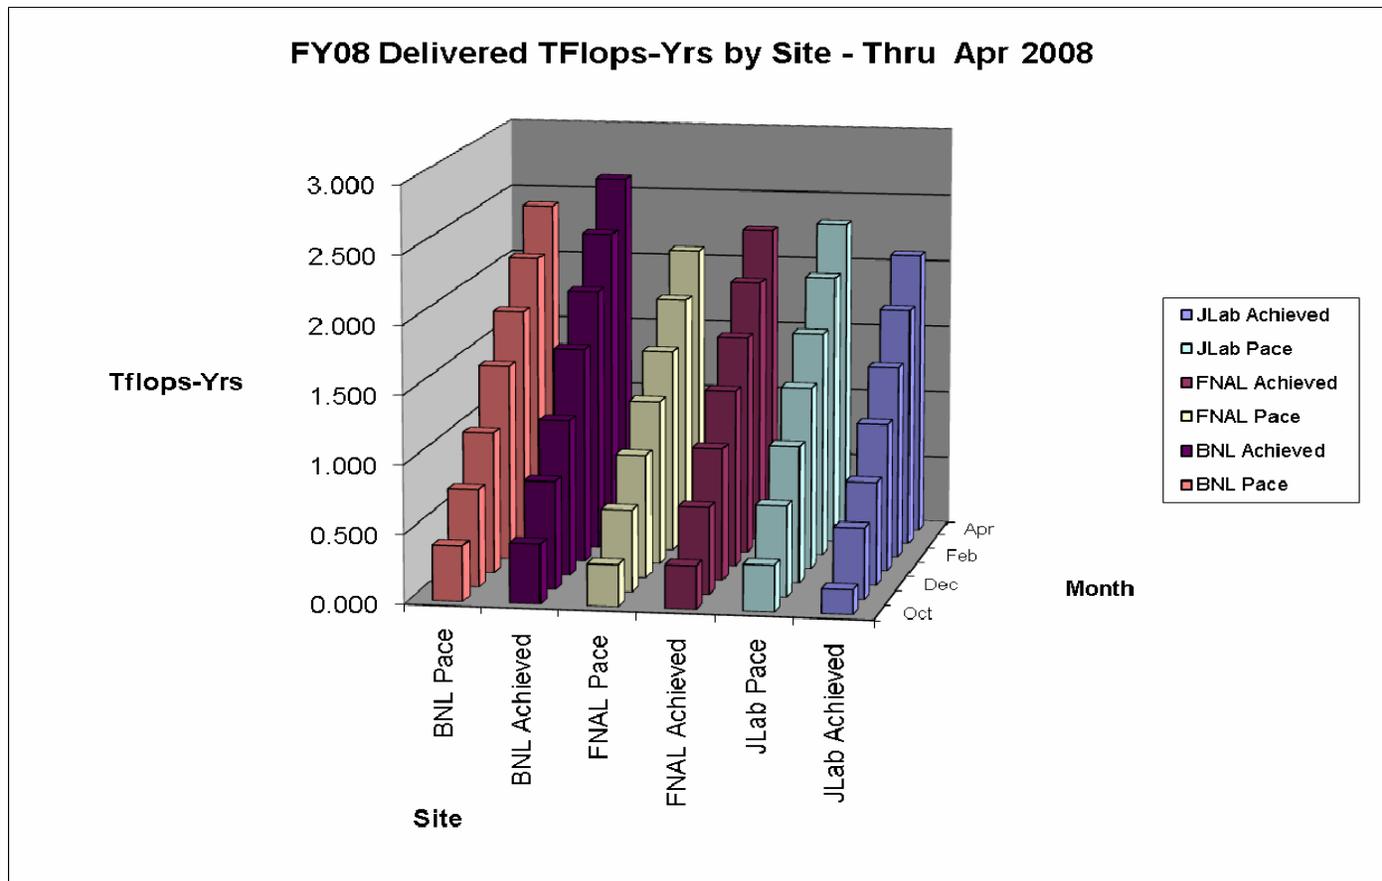
- FY07 performance goal = 9.0 Tflops-yrs delivered
- Total delivered = 9.674 Tflops-yrs (107.5% of goal)

■ FY08

- FY08 performance goal = 12 Tflops-yrs
- Linear pace goal through April is 6.98 Tflops-yrs
- Through April, LQCD has delivered 7.0 Tflops-yrs (100.3% of goal)



Delivered Tflops-Yrs by Site – FY08



Status of Progress towards FY08 Technical Performance Indicator Metrics

<i>Measurement Indicator</i>	<i>Performance Goal</i>	<i>Performance Results (through Apr 2008)</i>	
% of improvement in customer satisfaction rating (on a scale of 1 to 10)	Increase from 82% to 87% <i>(Additional 5% improvement over FY07 survey rating)</i>	TBD	(TBD)
Number of distinct users of the facility (includes DOE labs, LQCD and academic communities)	Increase from 25 to 30	Total = 82 (Thru Apr '08: FNAL=42; JLab = 25; BNL=15)	(+)
% of helpdesk tickets closed within 2 business days	Increase from 90% to 92%	96% (Ave for three sites thru Apr '08)	(+)
% of average machine uptime at the Metafacility	Increase from 92% to 93%	93.2% (BNL: 98.9%; FNAL: 98.2% JLab: 82.6%) Capacity-weighted average = 93.2% JLab's shortfall is due to the quad-core upgrade in October and the site-wide power outage in Dec/Jan. The trend indicates that we will exceed this metric (JLab was 95% in Nov and 89% in Dec, despite the power outage.	(+)

Status of Progress towards FY08 Technical Performance Indicator Metrics (2)

<i>Measurement Indicator</i>	<i>Performance Goal</i>	<i>Performance Results (through Apr 2008)</i>	
% of delivered node hours consumed by jobs with an exit error status	Additional 10% reduction from FY07 baseline	FNAL: 6.9% JLab: TBD BNL: QCDOC does not permit the use of exit codes for this purpose	(-)
Aggregate computing resources deployed by the project	Increase from 11.5 Tflops to 15.6 Tflops (Additional 4.1 Tflops)	To date, we have deployed 11.75 TFlops. Will meet goal if we deploy 4.2 Tflops as planned, but will not likely be in production by year end.	(tbd)
Increase frequency of vulnerability scans on nodes visible from Internet	Increase from monthly to bi-weekly	FNAL: Daily automated scans by site security staff. Log files read daily by LQCD staff. JLab: Daily vulnerability scans on all externally-facing systems. ACLs have been tightened on the various enclaves. BNL: Daily vulnerability scans on all externally-facing systems. Log files collected daily and read by cyber-security staff.	(+)

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



Financial Performance

FY2007 Final Cost Performance (Actual)

- Period of Performance (*Oct-06 through Sep-07*)

	<u>Personnel</u>	<u>Equipment</u>	<u>Total</u>
Budget			
FY06 Carry-Forward	\$ 20K	\$ 52K	\$ 72K
FY07 Budget	<u>\$ 980K</u>	<u>\$ 1,520K</u>	<u>\$ 2,500K</u>
Total Avail. Funds	\$ 1,000K	\$ 1,572K	\$ 2,572K
Actual Costs to Date			
	\$ 966K	\$ 1,298K	\$ 2,264K
% of budget	97%	83%	88%
% of yr complete	100%	100%	100%

Personnel costs in reasonable agreement with budget.

Equipment expenses under spent in large part due to cost savings associated with 7n upgrade.

FY2008 Year-to-Date Cost Performance

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

	<u>Personnel</u>	<u>Equipment</u>	<u>Total</u>
Budget			
FY07 Carry-Forward	\$ 34K	\$ 243K	\$ 277K
FY08 Budget	<u>\$ 930K</u>	<u>\$ 1,570K</u>	<u>\$ 2,500K</u>
Total Avail. Funds	\$ 964K	\$ 1,813K	\$ 2,777K
Actual Costs to Date			
	\$ 372K	\$ 205K	\$ 520K
% of budget	39%	11%	19%
% of yr complete	50%	50%	50%

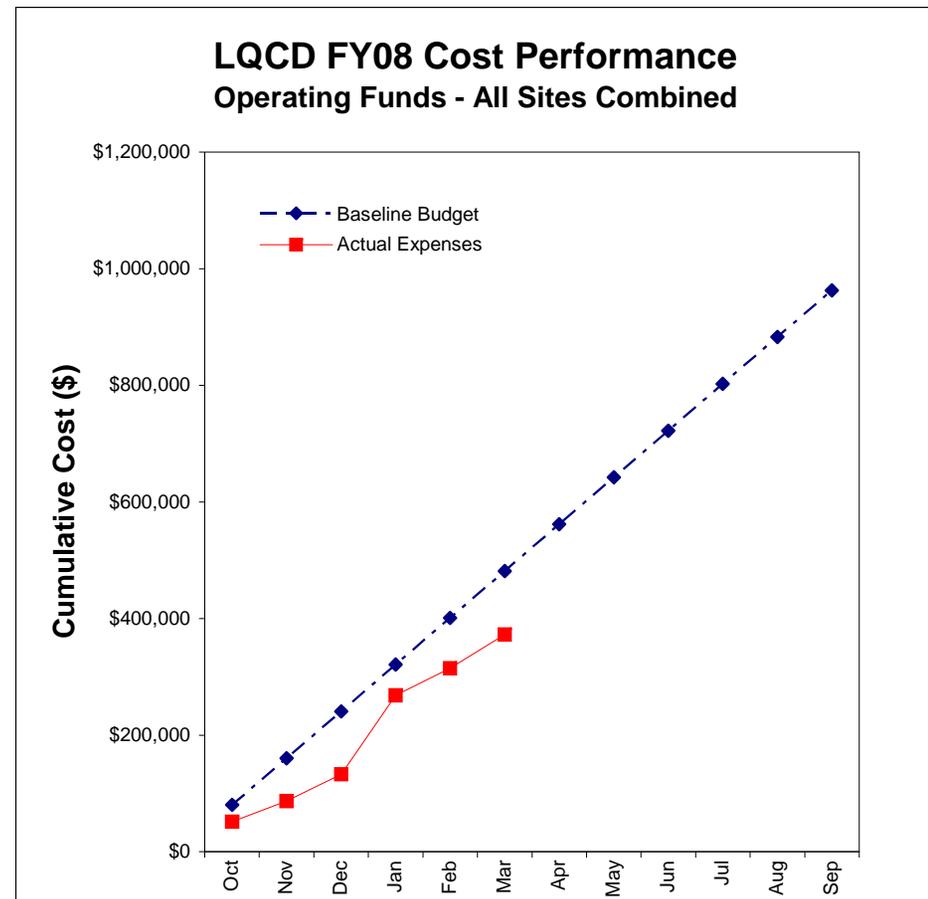
Personnel costs YTD costs appear under spent; expect ramp-up in late FY08 to support new cluster deployment (*more on the next slide*).

Equipment expenses to date related largely to 7n upgrade; large expenditure will occur late in FY08

FY08 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
 - Spend rate increases during acquisition & deployment activities
- Combined spend rate through March 2008, for all three sites, is trending slightly downward compared to a linear baseline forecast.
- YTD costs through March are below linear forecast:
 - Level of support required at BNL has been less than anticipated
 - FY08 deployment activities at FNAL are not scheduled to begin until FY08-Q4; spend rate will increase to support additional activity.



FTE Performance

<i>Project-funded Personnel Support (FTE-yrs)</i>							
Institution	FY06		FY07		FY08		FY09
	Plan	Actual	Plan	Actual	Plan	Actual	Plan
BNL	1.00	0.77	1.00	1.05	1.00	0.28	1.00
JLab	0.90	0.88	1.90	1.44	1.35	1.21	1.35
FNAL (technical)	2.00	1.82	1.85	1.88	2.10	2.08	1.87
FNAL (proj mgt)	0.50	0.19	0.50	0.16	0.5	0.37	0.5
Total	3.90	3.66	5.25	4.53	4.95	3.94	4.72

FY08 actual reflects normalized average through Mar-2008

- Level of effort is reported and tracked on a monthly basis.
- Personnel support levels are adjusted to support new cluster procurement and deployment activities
- Level of project-funded technical support at JLab and FNAL may be on the low side. In FY09, we may consider re-allocating funds if level of support required at BNL remains low.



Total Project Cost Performance Summary

- Period of Performance (*Oct-05 through Mar-08*)
- Project duration complete: 63%

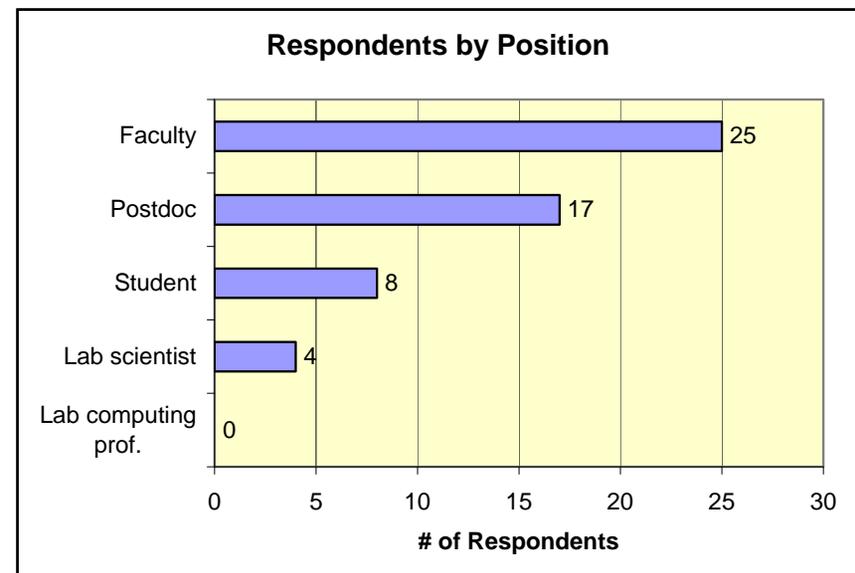
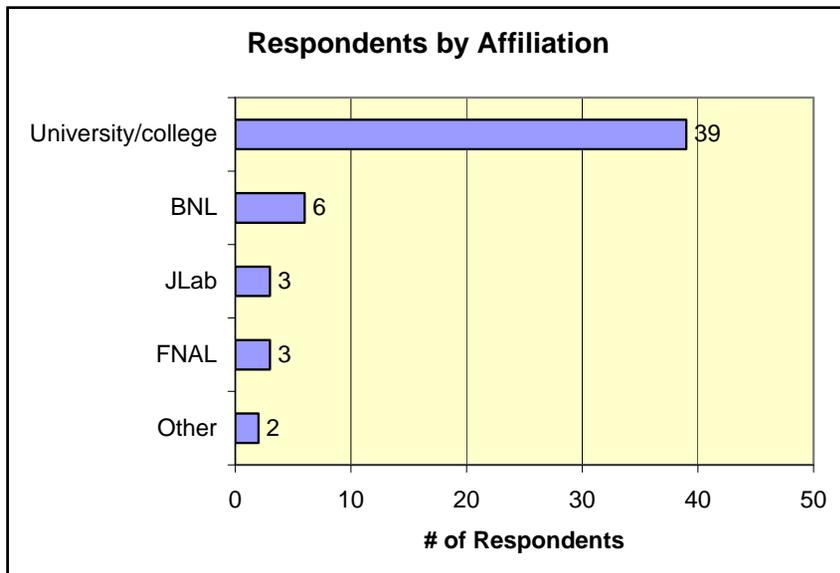
	<u>Personnel</u>	<u>Equipment</u>	<u>Total</u>
Total Project Budget	\$ 3,330K	\$ 5,870K	\$ 9,200K
Less			
Actual Costs to Date	<u>(\$ 2,060K)</u>	<u>(\$ 3,144K)</u>	<u>(\$ 5,204K)</u>
Budgeted Funds Remaining	\$ 1,270K	\$ 2,726K	\$ 3,996K
% of budget spent	62%	54%	57%



FY07 User Survey Results

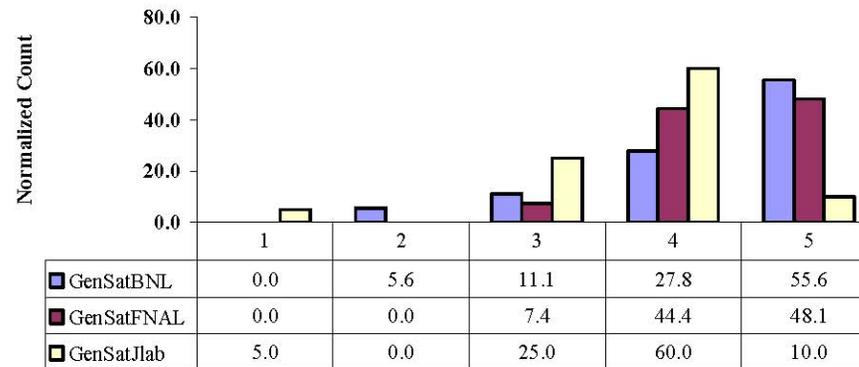
User Survey

- On-line user survey was conducted last fall. Covered a total of 23 areas, including overall user satisfaction with computing facilities, documentation, user support (helpdesk), success of job submission, and the proposal / resource allocation process.
- Included pull-down menus (ratings from 1 to 5) and free-form text entries.
- Received responses from 54 users out of an estimated user base of 60.



User Survey: General Satisfaction - 82%

Normalized General Satisfaction



(Level: Highest = 5)

Positive Comments

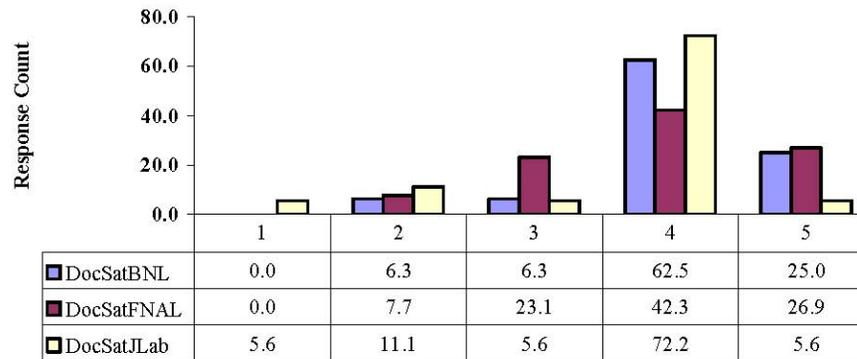
- Very effective.
- Outstanding service at FNAL in every respect of cluster management
- Increased support staff at JLab has resulted in increased level of satisfaction.

Constructive Criticism

- Long queues; high failure rates.
- No stable environment; unclear responsibilities of staff.
- Three different security systems is annoying.
- Low IO rate on QCDOC makes it unsuitable for some valence calculations.

User Survey: Documentation - 78%

Satisfaction with documentation



(Level: Highest = 5)

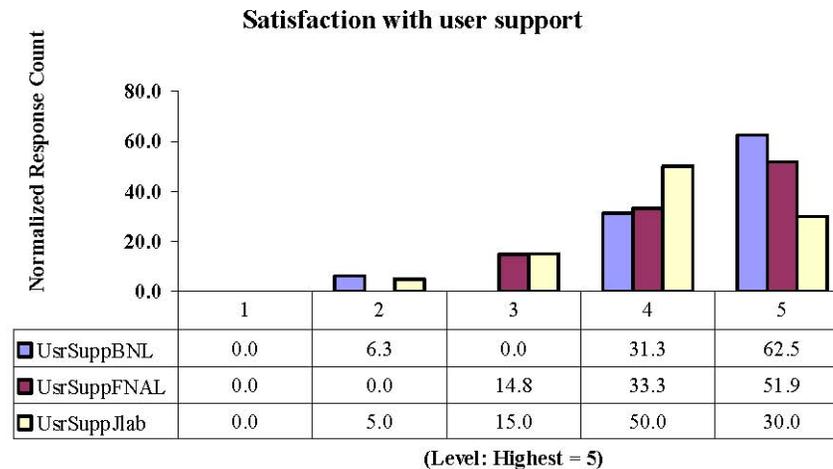
Positive Comments

- All sites do a reasonable job at site level documentation.

Constructive Criticism

- Difficult to find information needed. Best if user experience doesn't change over long periods of time.
- Documentation on PBS queuing was great, but needed to be told where to find it (location on web not obvious).
- Much easier to get information by word-of-mouth than looking online.

User Survey: User Support - 86%



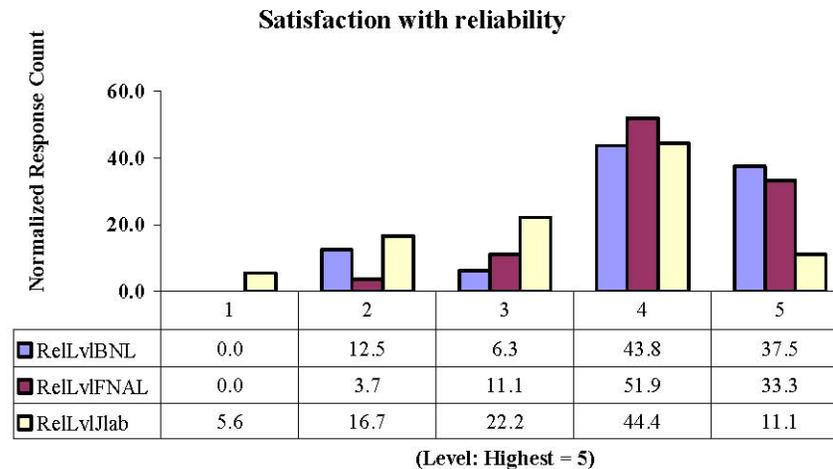
Positive Comments

- Good and responsive service. Queries are resolved by e-mail and problems are generally solved.
- User support staff at all sites are helpful and quite motivated.
- Always ready to provide uncomplicated, immediate solutions.
- Great job!
- Don Holmgren and Amitoj Singh deserve medals.

Constructive Criticism

- Took weeks to get QCDOC account at BNL.

User Survey: Reliability - 74%



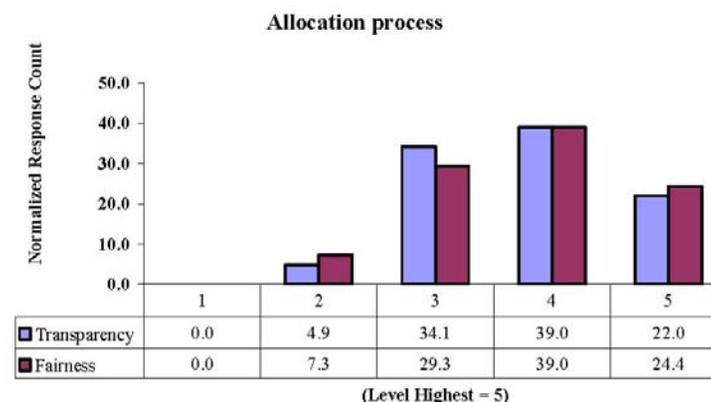
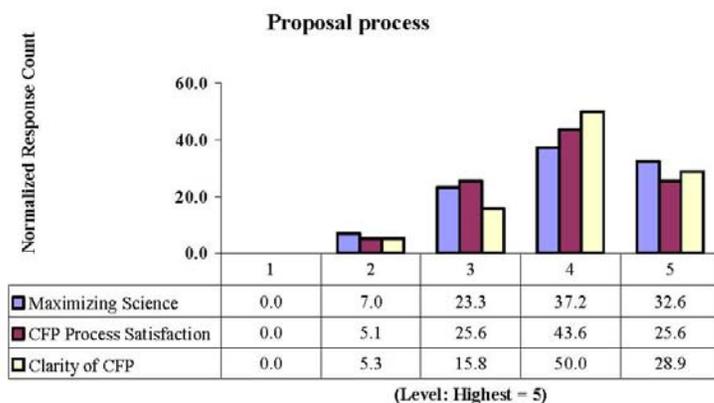
Positive Comments

- Have never had a problem with job failure or unreasonable downtime.
- Failure rate has decreased to the point where only a few percent of time is lost.

Constructive Criticism

- Significant number of sporadic, inexplicable job failures result in wasted CPU time.
- Some problems with jobs hanging, thereby allowing time to effectively be wasted.
- Irregular service interruptions are bit of an inconvenience.
- Getting jobs to successfully start can be a chore on QCDOC; once running, is quite stable.

User Survey: Proposal/Allocation Process - 69%



Positive Comments

- 70% of respondents felt that allocation process helps maximize scientific output.
- 98% of respondents felt that the time allowed for proposal preparation was adequate.
- 83% found the Call for Proposals to be adequately clear (i.e., no additional clarification needed).
- Different groups adopted different strategies in submitting proposals (e.g., several small projects encompassed in single large proposal. Committee fairly recognized this in allocations.

Constructive Criticism

- 2007 proposals were not easy to find on the web.
- Process works best for established big user groups with long-term computational programs.
- Not enough allocation for applications.
- Communication process less effective for individuals who do not belong to large collaborations.



User Survey Summary

- Overall, the survey resulted in a number of positive statements and constructive suggestions.
- We have used the data to generate an internal scorecard to help us quickly assess areas requiring attention and quantitatively measure improvements.
 - Quantifies user satisfaction in 23 distinct areas, overall and by site
- We plan to use the survey results and analysis to develop an action plan to address shortcomings and (hopefully) improve the user experience.
- We will conduct another survey in late summer/early fall
 - Assess impact of changes made.
 - Provide quantitative measure of performance to satisfy performance metric (5% improvement in customer satisfaction rating).



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.
 - Acknowledging that the host laboratories also provide significant resources (e.g., space, power, cooling, networking, etc.)
- The user survey has provided us with a good idea of where to focus efforts to improve operational effectiveness and the user experience.



Questions?