

Results of the FY10 User Survey for Lattice QCD (LQCD) Computing Facility

Unique Project (Investment) Identifier: 019-20-01-21-02-1032-00.

Operated at
Brookhaven National Laboratory
Fermi National Accelerator Laboratory
Thomas Jefferson National Accelerator Facility

for the
U.S. Department of Energy
Office of Science
Offices of High Energy and Nuclear Physics

Version 0.2

Revision Date
January 20, 2011

PREPARED BY:
Bakul Banerjee, FNAL

Change Log

Revision No.	Pages Affected	Effective Date
Revision 1.0	Entire Document	

Table of Contents

Results of the FY10 User Survey for Lattice QCD (LQCD) Computing Facility.....	1
1 Objective.....	4
2 Summary of Results.....	4
3 Comments.....	6
3.1 General.....	6
3.2 Improvements over the past year.....	6
3.3 Support.....	7
4 General comments from the survey analyst.....	7
5 General requirements.....	7
6 Deployment of the survey.....	8
7 Methodology for the analysis of the survey.....	8
7.1 General employment information.....	8
7.1.1 Employment level.....	8
7.1.2 Usage of LQCD computers.....	8
7.1.3 Job submission.....	9
7.1.4 Facility usage.....	9
7.2 User satisfaction.....	9
7.2.1 Overall user satisfaction with facilities.....	9
7.2.2 Overall improvement.....	10
7.2.3 Documentation.....	10
7.2.4 User support.....	11
7.2.5 Reliability.....	12
7.2.6 Responsiveness.....	13
7.2.7 Ease of access.....	14
7.3 Support and communication.....	14
7.3.1 E-mail.....	14
7.3.2 Web support.....	15
7.3.3 Other tool support.....	15
7.4 Helpdesk (analysis of the last problem solved):.....	16
7.4.1 Knowledge of how to ask help.....	16
7.4.2 Help requested by facility.....	16
7.4.3 Time to initial response (working hours).....	16
7.4.4 Closeout from initial response.....	17
7.4.5 Working days needed to solve the problem.....	17

- 7.5 Proposal and allocations 17
 - 7.5.1 Satisfaction with the proposal process 17
 - 7.5.2 CFP clarification and time needed 17
 - 7.5.3 Allocation process 18
- 7.6 Running jobs 18
 - 7.6.1 Overall job submitted 18
 - 7.6.2 Overall job failure rate 18
 - 7.6.3 Job failure rate due to hardware 19
 - 7.6.4 Job failure rate due to access failure 19
 - 7.6.5 Job failure rate due to user error 19
- 7.7 Mass storage 19
 - 7.7.1 Adequacy of disk storage 19
 - 7.7.2 Tape library use and quality 19

1 Objective

To fulfill the goal of the LQCD facility to continuously serve our user community in the best possible manner, an initial user survey was conducted during August and September of 2007. Follow-up surveys were conducted at the end of FY2008 and FY2009. The objective of these surveys is to assess the level of satisfaction experienced by the users of the Lattice QCD Computational Facility. Results of the FY10 survey, presented in this document, indicate how satisfied or dissatisfied users were. Using the results of these surveys, the Integrated Project Team (IPT) considers ways to improve and optimize the services using the limited resources available to the project.

2 Summary of Results

Although the LQCD project team always strives to provide complete satisfaction to its customer community, the recently completed user survey will allow the team to fine-tune the services to improve user satisfactions. A comprehensive set of questions for the FY10 survey was defined by the project team in collaboration with the LQCD Executive Committee and the Scientific Program Committee. The questions were designed to identify performances of individual facilities, namely, the two clusters at Fermilab (FNAL) and Thomas Jefferson Lab (Jlab) and the QCDOC machine at Brookhaven National Laboratory (BNL). However, it should be noted that the project management team that includes site managers at all three facilities work together closely toward to fulfill the common goal of the project.

The survey, targeted toward users of the LQCD Computing Facility, was executed using the Zoomerang Survey Tool already in place at Fermilab. Similar to the surveys done during 2007 and 2008, survey questions were grouped in eight categories. A total of 44 questions were posing, often including questions specific to the three laboratories. Answers to some of the questions had alphanumeric values. For subjective questions, we asked users to choose from 1 to 5 satisfaction ratings with 5 being highly satisfied. Graphical views of the data collected and tabulated are given in the section titled "Detailed Results". The statistical data presented are normalized for each laboratory as needed to remove any bias.

For subjective rankings, rankings of 4 and 5 were used to infer satisfaction. Users were asked to provide short comments in several categories. Comments are included in this report verbatim. These comments often reveal underlying issues and may be helpful to the site managers. If for an area of the survey, more than one laboratory has a percentage rating below 80%, considered with associated comments, that particular area may be an area of concern. These items are bolded in the summary section

It is important to put forth a word of caution regarding the survey. Since the total population of users is relatively small, the outliers may affect the results of the survey significantly. A single unsatisfied customer may affect the satisfaction ranking for an area.

Descriptions of each category and the summary of the survey results associated with it are given below:

1. General: Questions under this category are designed to collect demographic data of the user community.
 - a. Among the total of 39 respondents, 26 users are employed by a university or a college, the rest are mostly employed by the participating laboratories.
 - b. 14 users are faculty members. Post docs make up a significant portion of the rest.

- c. 21 users submit jobs daily. Only 8 users submit jobs occasionally.
- d. Most active users submit an average of less than 10 jobs per week.
- e. User satisfaction: These questions assessed the overall user satisfaction and related satisfaction levels related to documentation, user support, reliability of the machines, responsiveness and accessibility.

	BNL	FNAL	Jlab
Overall satisfaction	89%	88%	67%
Documentation	75%	65%	83%
User support	86%	92%	83%
Reliability	78%	87%	61%
Responsiveness	88%	92%	89%
Ease of access	89%	63%	89%
Improvement made during the past year	43%	89%	73%

2. Communication: The topics covered were various modes of communications including e-mails, web communications, and other tools.
 - a. Email: 88%, 93%, and 94% of BNL, FNAL, and Jlab users found email related communications to be satisfactory
 - b. Web support: BNL, FNAL, and Jlab received satisfactory ratings of 89%, **77%, and 70%** of users respectively. Although overall satisfaction ratings for documentation have improved over the past three years, the web support seems to be an area of concern for FNAL and Jlab.
 - c. Other tool support: Users were asked about their satisfaction level regarding various general purpose user tools, for example, various possible command line tools. 100%, 83%, and 83% of users indicated satisfaction.
3. Helpdesk: All three LQCD facilities operate site specific helpdesks. An extensive set of questions were posed to determine the usage and efficacy of the Helpdesk at each site. After determining the awareness of the existence of the helpdesk, users were asked to rate their satisfaction regarding the last helpdesk request they submitted in terms of time to initial response and close out of the helpdesk ticket, and the level of satisfaction with the helpdesk request. 92% of users knew how to ask for help. They were asked to consider the last problem report they submitted. The response to the evaluation of the last problem report is given below:
 - a. The normalized spread of the helpdesk request submission among BNL, FNAL, and Jlab is 14%, 47%, and 39%
 - b. Time to initial response: 32 out of 39 helpdesk requestors received initial response within 6 hours.
 - c. Problem solved by initial response: 69% of problems were solved using the initial response. About 95% of the problems were solved within 3 days. It is likely that a small fraction of problems may require modification of the system and may not be solved for months.
 - d. Satisfaction with the helpdesk: 83% of users found the help received satisfactory.
4. Proposal and allocation: These questions were designed to understand the satisfaction level related to proposal and allocation processes.

- a. Scientific output: In FY10, 85% of users believe that the allocation process helps maximizing the scientific output.
 - b. User meetings: The need for user meetings was assessed in this question. 21% users indicated that additional user meetings may be beneficial. It indicates a higher demand for additional meetings from previous years.
 - c. Satisfaction with the process: 86% of users found the allocation process satisfactory.
 - d. Clarity of the Call for Proposal(CFP): 93% of users thought that the CFP was clear
 - e. Transparency: 86% of users found the proposal process adequately transparent. This may be an area of concern.
 - f. Fairness: **78%** of users found the allocation process to be fair.
5. Running jobs: The objective was to assess the success of job submissions. The percentage of users less than or equal to 10 job failures over the past year was 85% at BNL, 42% at FNAL and 60% at Jlab. A successful running of jobs on the LQCD machines depends on multiple factors including hardware, software, user configuration, experience and preferential styles.
 6. Mass storage: 86% of users are satisfied with the storage solutions provided. At Jlab and FNAL, 44% and 33% of users found the tape libraries to be satisfactory.
 7. General comments: Users provided an extensive set of comments, both general and specific.

3 Comments

3.1 General

1. BNL: After Stratos left no one responded to my e-mails about not being able to access the qcdochosta and qcdochostb servers. Luckily I didn't end up needing to.
2. BNL: Don't know if there was a help desk specifically, no one replied to me at all about inability to login into qcdochosta, qcdochostb.
3. FNAL: It seems storage space is often limited, which sometimes makes certain running inconvenient or causes delays.
4. FNAL: The FNAL staff is very responsive and helpful.
5. Jlab: It is very difficult to run jobs with 128 cores or more
6. Jlab: I am generally very happy with the JLAB facility. This year, I am still very happy with the staff. However, I believe JLAB has grown too fast to keep up with its hardware (the ARA growth). Since early 2010, I have had a large number of job failures for many reasons, many of them not satisfactorily resolved. I believe the size of the support staff is no longer big enough to maintain the factor of 10 growths in the JLAB cluster/gpu machine. At facilities where the machines work smoothly, they have routine maintenance, in the most extreme cases, taking the entire compute machine off-line for 12 hours per week to perform these tasks. I think JLAB needs to implement such measures.
7. Jlab: 7n cluster has been very unreliable

3.2 Improvements over the past year

1. BNL: What improvements?

2. FNAL: I'm optimistic about D_s, but until it's online and available I have fewer resources at my disposal (with pion gone) than I did this time last year
3. FNAL: Only used FNAL resources for part of year -- can't compare improvements.
4. Jlab: The web site needs to be updated.

3.3 Support

1. user installation request... eg, "Can I have this python package installed in the main /usr/ section?"
2. Better documentation of Chroma
3. Nothing in my mind at this moment.
4. The (old) tape system would constantly timeout. I have not extensively used the new LUSTRE system, so I can't speak to its efficiency.
5. FNAL: The support team at FNAL does an excellent job. Very responsive and helpful.
6. Jlab: I had some difficulties with documentation and file system management at JLAB early in the year, but all is well now.
7. Jlab: I have routine job failures at JLAB, which they are currently not able to resolve. These range from I/O problems to the file server, to mpi jobs failing to launch. They need to figure out how to perform scheduled routine maintenance on the cluster to trouble shoot problematic nodes. They need to address the I/O communications problems with the new /lustre file system.
8. Jlab: Documentation often seems to be hard to find or outdated (or both). I've found it necessary to request help rather than consult documentation.
9. Jlab: Some of the user documentation on the JLab website is outdated.
10. FNAL: With future upgrade to lustre at FNAL, disk space issues may be met, but at the moment it's borderline...

In summary, results of the user survey indicate that the LQCD facility did not performed as well as it was expected.

4 General comments from the survey analyst

1. The survey was an exhausting one with a total of 78 fields to be filled by the survey takers. Perhaps this may be the cause of low response rates. Statistically, low response rates usually results in poor overall ratings.
2. Compared to other years, most laboratories did not do well in most categories.
3. There is an indication that users are interested in larger and newer facilities, yet newer environments lead to user dissatisfaction. BNL did well, but Jlab with the new machines did not.

5 General requirements

The target audience of the survey includes USQCD collaborators, Principle Investigators, faculty members, researchers, students and post-docs who submit jobs to the LQCD Computing Facility at any of the three sites, BNL, FNAL, and Jlab.

The on-line survey should be easily accessible from various outside organizations for a limited time. The survey should be anonymous. Principle Investigators for LQCD projects submitting jobs to the LQCD Computing Facility are requested to assure that the students and post-docs are included in the “sdac” mailing list.

6 Deployment of the survey

The request to participate in the survey was sent to the USQCD mailing list on October 28, 2010, followed by reminders and encouragements. The survey closed on December 12, 2010.

7 Methodology for the analysis of the survey

After the closing, the survey data was retrieved in the tabulated format from the Fermilab’s survey website. 39 users responded to the survey. Detailed results

7.1 General employment information

FY10

Employed by	Count
BNL	4
FNAL	1
Jlab	3
University or college	26
Other	2

7.1.1 Employment level

FY10

Type	Count
Student	4
Postdoc - University	10
Faculty - University (tenured or tenure track)	14
Research Scientist - University (non-tenured)	1
Research Scientists – Laboratory	4
Lab computing professional	0
Other university staff	1

7.1.2 Usage of LQCD computers

FY10

Usage	Freq.
Daily	21
Weekly	9
Monthly	3
Occasionally	8

7.1.3 Job submission

FY10

Avg. Jobs	Freq.
10	11
20	6
50	6
100	6
200	4
500	5
1000	1
5000	0

7.1.4 Facility usage

FY10

Facility	Users
BNL	7
FNAL	22
Jlab	18

7.2 User satisfaction

7.2.1 Overall user satisfaction with facilities

FY10

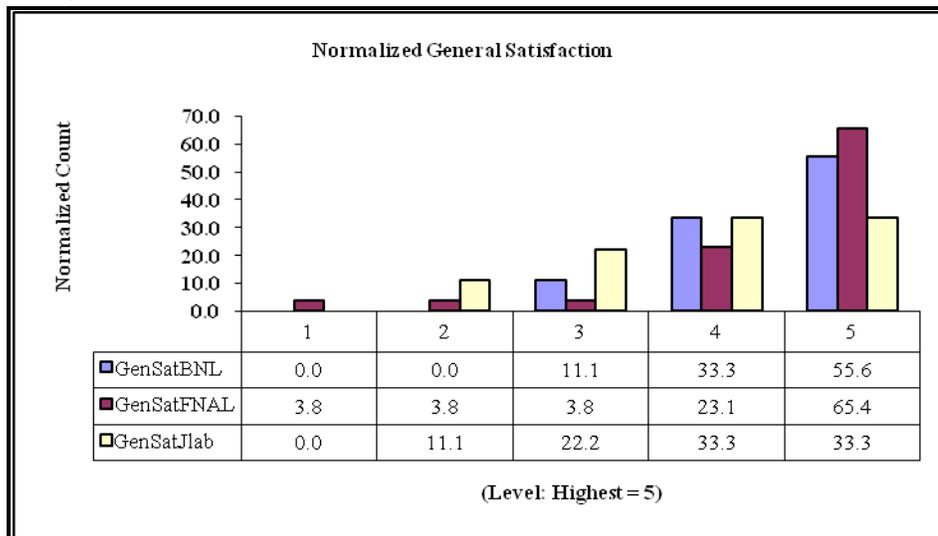


Figure 1: General satisfaction

Overall Satisfaction			
	BNL	FNAL	Jlab
Median	5	5	4
Mode	5	5	5
Average	4.44	4.42	3.89
% 4 or 5	88.9%	88.5%	66.7%
Skew	-1.18	-2.24	-0.53
# Responses	9	26	18

Figure 2: Statistics for general satisfaction

7.2.2 Overall improvement

FY10

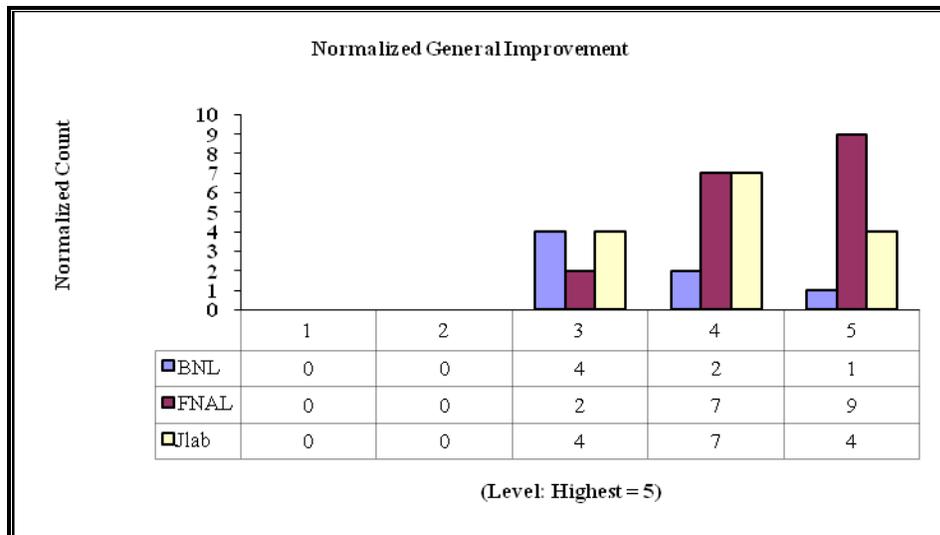


Figure 3: Normalized general improvement

Improvements			
	BNL	FNAL	Jlab
Median	3	4	4
Mode	3	5	4
Average	3.57	4.21	4.00
% 4 or 5	42.9%	84.2%	73.3%
Skew	1.32	-1.82	0.00
# Responses	7	19	15

Figure 4: Statistics for normalized general improvement

7.2.3 Documentation

FY10

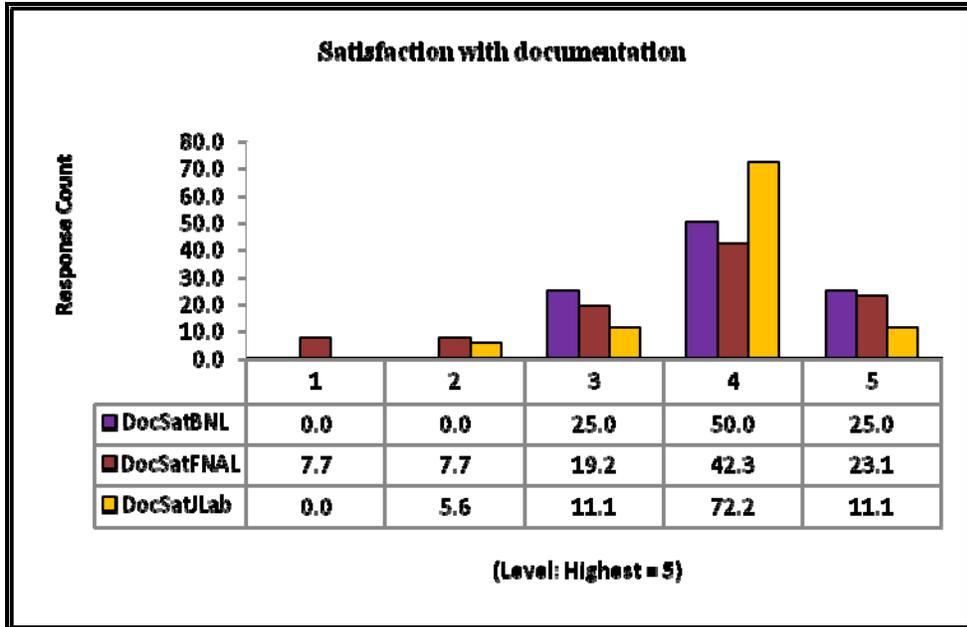


Figure 5: Normalized satisfaction with documentation

User Documentation	BNL	FNAL	Jlab
	Median	4	4
Mode	4	4	4
Average	4.00	3.65	3.89
% 4 or 5	75.0%	65.4%	83.3%
Skew	0.00	-0.94	-1.20
# Responses	8	26	18

Figure 6: Statistics for satisfaction with documentation

7.2.4 User support

FY10

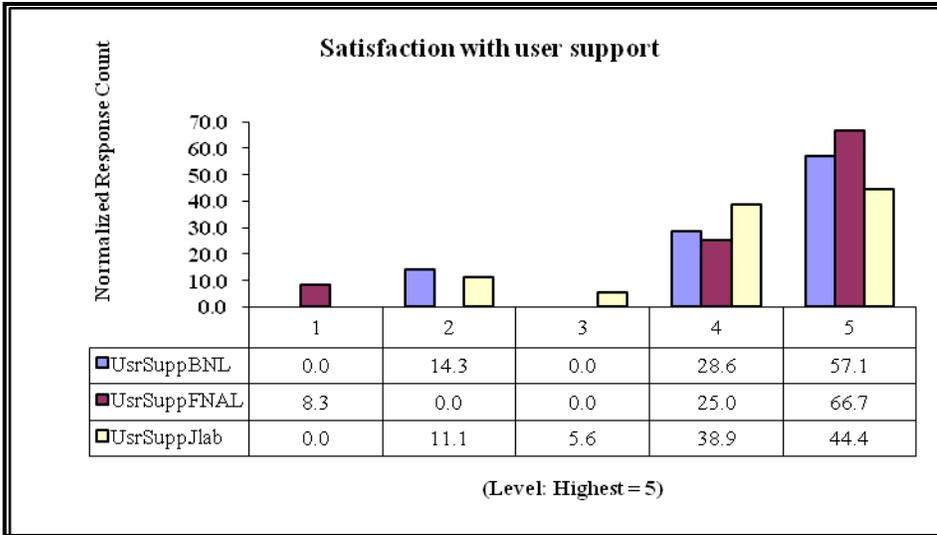


Figure 7: Normalized user support

User Support	BNL	FNAL	Jlab
Median	5	5	4
Mode	5	5	5
Average	4.29	4.42	4.17
% 4 or 5	85.7%	91.7%	83.3%
Skew	-1.96	-2.54	-1.19
# Responses	7	24	18

Figure 8: Statistics for user support

7.2.5 Reliability

FY10

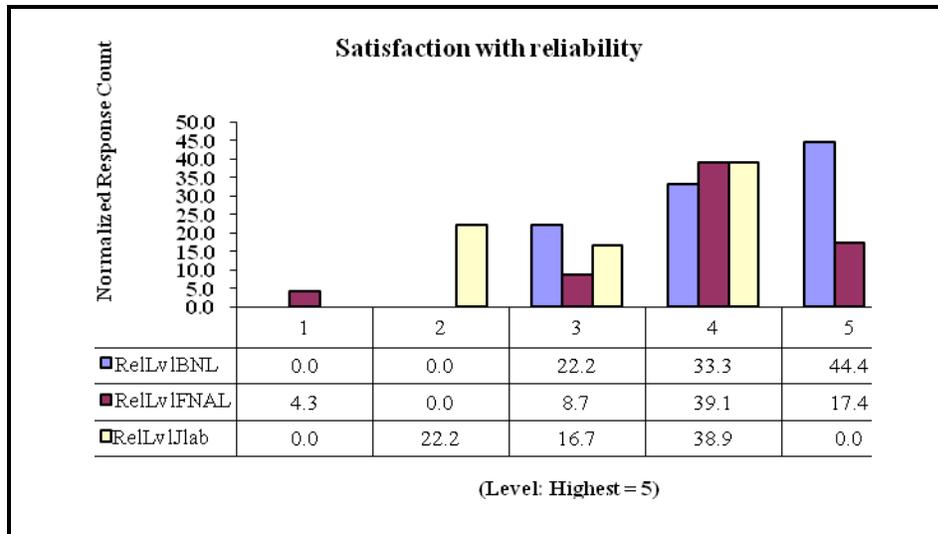


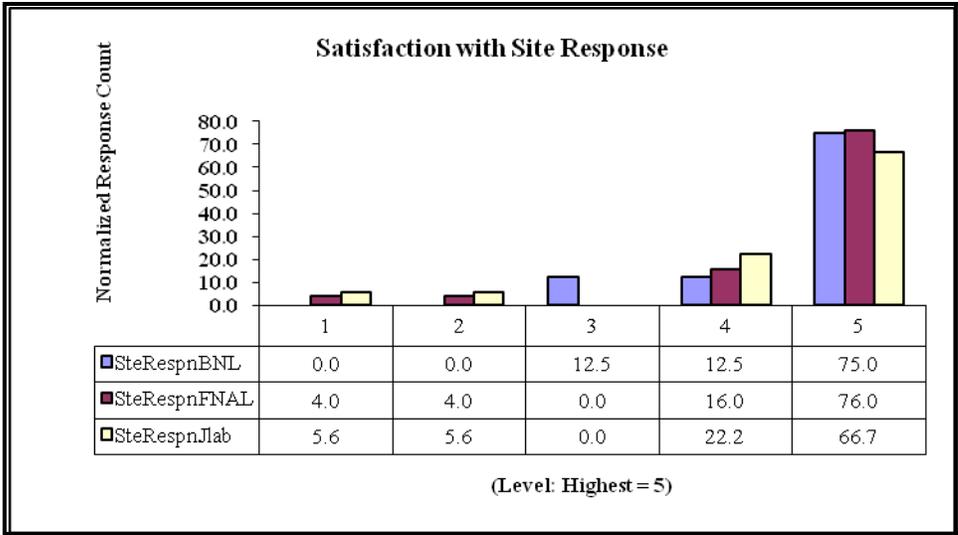
Figure 9: Normalized satisfaction with reliability

Reliability			
	BNL	FNAL	Jlab
Median	4	4	4
Mode	5	5	4
Average	4.22	4.26	3.61
% 4 or 5	77.8%	87.0%	61.1%
Skew	-0.41	-1.90	-0.38
#			
Responses	9	23	18

Figure 10: Statistics for reliability

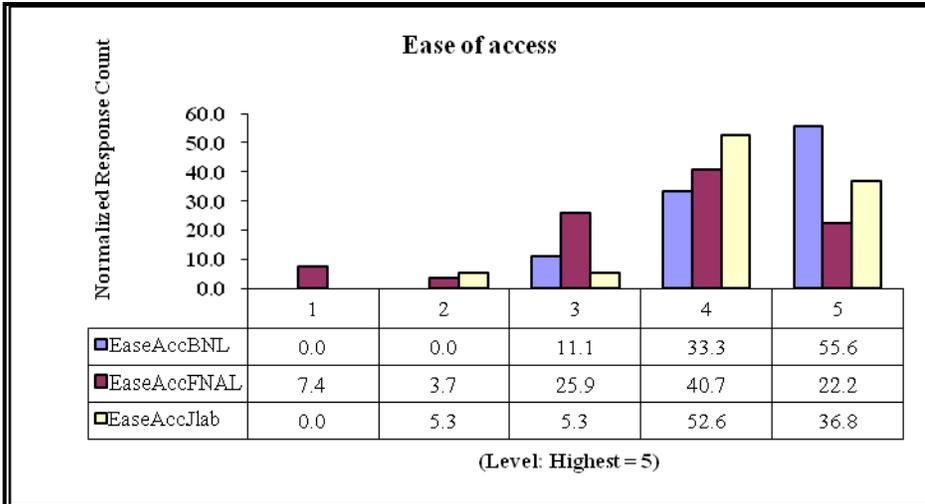
7.2.6 Responsiveness

FY10



7.2.7 Ease of access

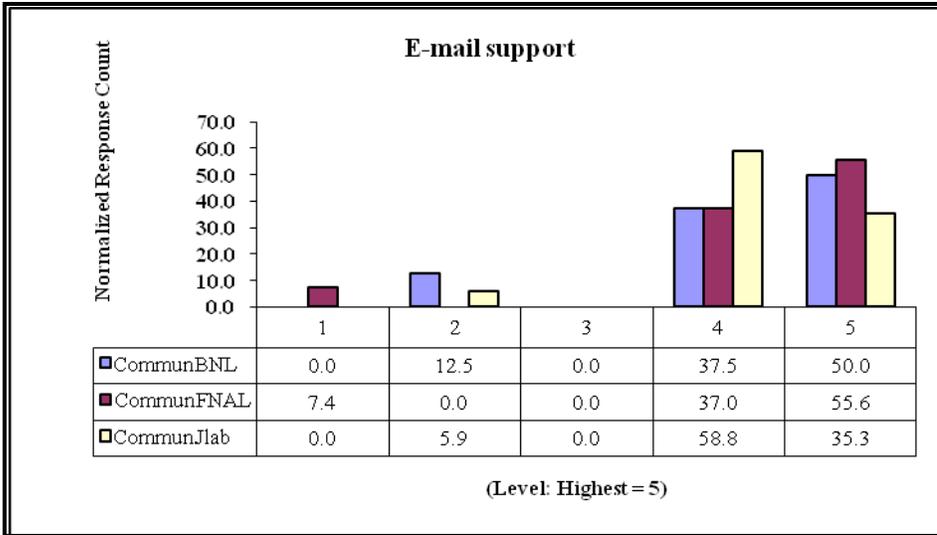
FY10



7.3 Support and communication

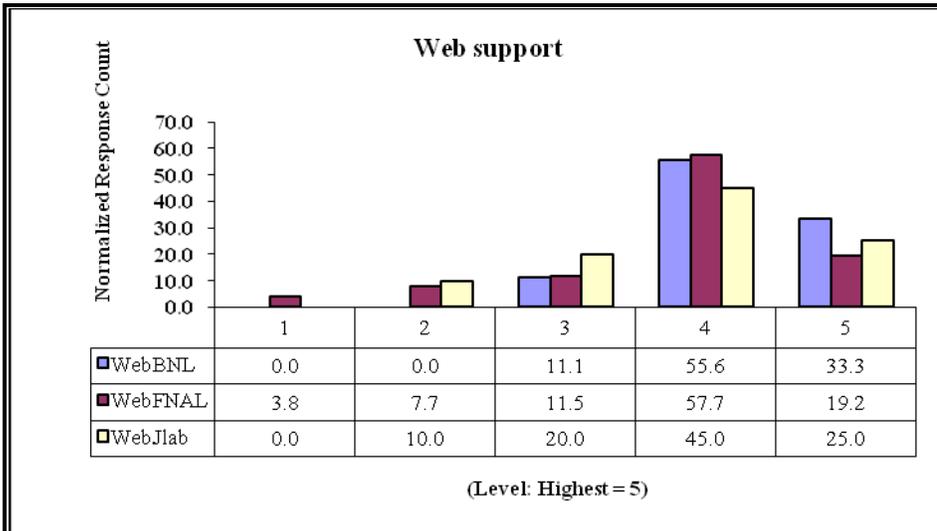
7.3.1 E-mail

FY10



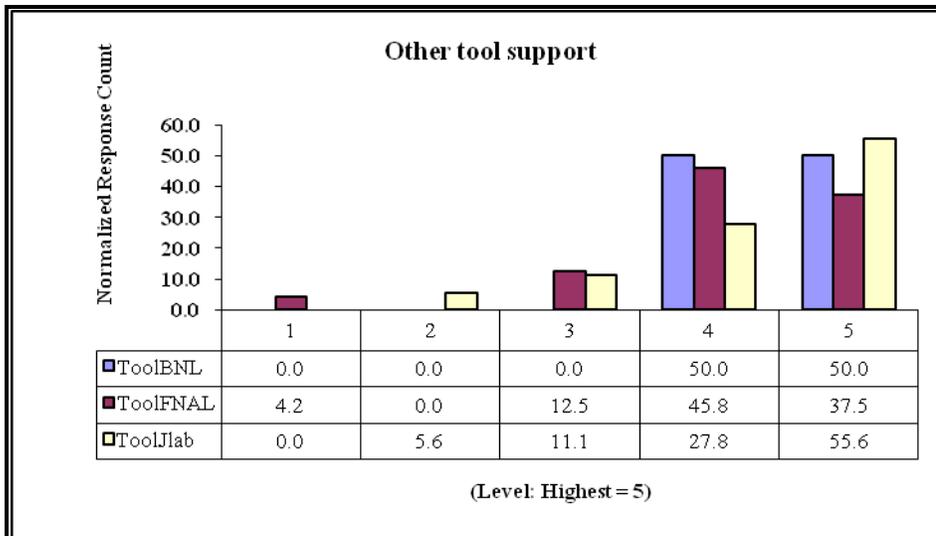
7.3.2 Web support

FY10



7.3.3 Other tool support

FY10



7.4 Helpdesk (analysis of the last problem solved):

7.4.1 Knowledge of how to ask help

FY10

Knows	Count
Y	36
N	3

7.4.2 Help requested by facility

FY10

Help asked	BNL	FNAL	Jlab
Y	5	17	14

7.4.3 Time to initial response (working hours)

FY10

<=Hours	Freq.
6	32
12	2
24	3
>=24	2

7.4.4 Closeout from initial response

FY10

Closed?	Count
Y	25
N	11

7.4.5 Working days needed to solve the problem

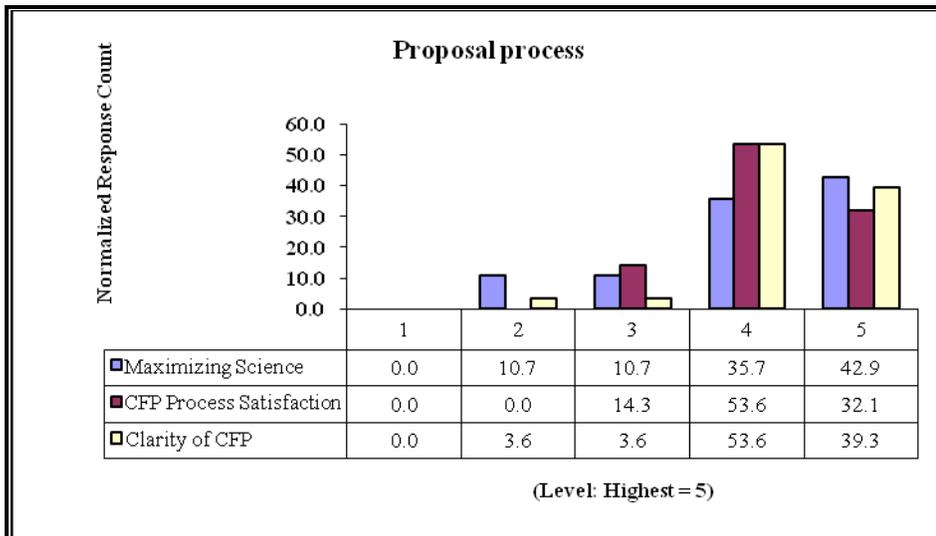
FY10

<=Days	Freq.
1	28
3	7
5	2
>5	2
>5	3

7.5 Proposal and allocations

7.5.1 Satisfaction with the proposal process

FY10



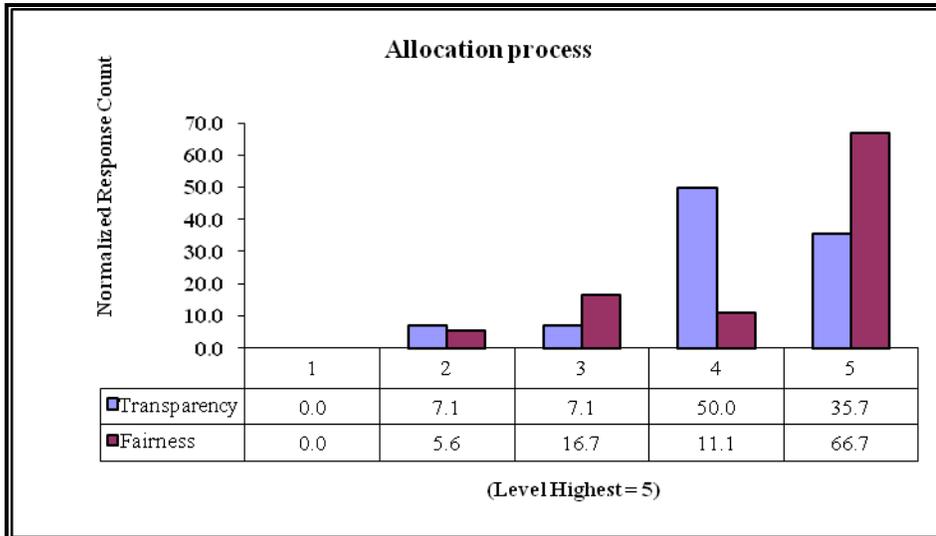
7.5.2 CFP clarification and time needed

FY10

Rating	Clarification needed?	Time adequate?
Y	3	25
N	24	2

7.5.3 Allocation process

FY10



7.6 Running jobs

7.6.1 Overall job submitted

FY10

< =Job Submitted	BNL Frequency	FNAL Frequency	Jlab Frequency
1	34	21	21
10	0	0	0
500	5	8	8
1000	1	2	2
5000	0	4	4
9999	0	5	5

7.6.2 Overall job failure rate

FY10

# of Failed job is <=	BNL Freq.	FNAL Freq.	Jlab Freq.
1	35	24	28
10	2	10	4
100	3	3	1
500	0	1	5
1000	0	0	1
2000	0	1	1

7.6.3 Job failure rate due to hardware

FY10

# of Failed job is <=	BNL Freq.	FNAL Freq.	Jlab Freq.
1	36	29	33
10	0	6	1
100	4	5	6
500	0	0	0
1000	0	0	0
2000	0	0	0

7.6.4 Job failure rate due to access failure

FY10

# of Failed job is <=	BNL Freq.	FNAL Freq.	Jlab Freq.
1	39	27	32
10	1	3	4
100	0	9	4
500	0	1	0
1000	0	0	0
2000	0	0	0

7.6.5 Job failure rate due to user error

FY10

# of Failed job is <=	BNL Freq.	FNAL Freq.	Jlab Freq.
1	36	25	30
10	3	7	4
100	1	7	6
500	0	1	0
1000	0	0	0
2000	0	0	0

7.7 Mass storage

7.7.1 Adequacy of disk storage

FY10

Y	31
N	5

7.7.2 Tape library use and quality

FY10

Tape library OK	# of users of tapes at Jlab	# of users of tapes at FNAL
Y	15	13
N	19	21
N/A	3	3

8 Additional analysis of the survey data

In addition to using the standard percentage of responders who gave a score of either 4 or 5, mean, median and skew functions were also calculated.

	Overall Satisfaction			Improvements		
	BNL	FNAL	Jlab	BNL	FNAL	Jlab
Median	5	5	4	3	4	4
Mode	5	5	5	3	5	4
Average	4.44	4.42	3.89	3.57	4.21	4.00
% 4 or 5	0.89	0.88	0.67	0.43	0.84	0.73
Skew	-1.18	-2.24	-0.53	1.32	-1.82	0.00
# Responses	9	26	18	7	19	15

	User Documentation			Web Site			Ease of Access			Reliability		
	BNL	FNAL	Jlab	BNL	FNAL	Jlab	BNL	FNAL	Jlab	BNL	FNAL	Jlab
Median	4	4	4	4	4	4	5	4	4	4	4	4
Mode	4	4	4	4	4	4	5	4	4	5	5	4
Average	4.00	3.65	3.89	4.22	3.81	3.85	4.44	3.67	4.21	4.22	4.26	3.61
% 4 or 5	75.0%	65.4%	83.3%	88.9%	76.9%	70.0%	88.9%	63.0%	89.5%	77.8%	87.0%	61.1%
Skew	0.00	-0.94	-1.20	-0.13	-1.28	-0.57	-1.18	-0.96	-1.15	-0.41	-1.90	-0.38
# Responses	8	26	18	9	26	20	9	27	19	9	23	18

According to this analysis, more than half of our users ranked the three sites in all categories but "improvements" as 4 or better (median is 4+, average is 4.0+ except in the following categories:

- Overall satisfaction: JLab = 3.89
- Improvements: BNL = 3.57
- User documentation: FNAL = 3.65, JLab = 3.89
- Reliability: JLab = 3.61
- Ease of access: FNAL = 3.67
- Web site: FNAL = 3.81, JLab = 3.85

From the above data, we can draw the following conclusions:

- a) Web site and user documentation are correlated, as one would expect. FNAL needs to improve on both of these categories.

- b) Ease of access points to user difficulties with FNAL kerberos access. Under the current situation with Kerberos, not much improvement can be done regarding this issue.
- c) Reliability at JLab is associated with the resources devoted to the new ARRA machines which affected JLab support.