

2007/2008 User Survey results

Francesca Verdier



NERSC User Group Meeting
October 3, 2008



Response Profile

457 respondents

- 70% “big user” response rate
- 43% “medium user” response rate
- 16.3% overall response rate



- ASCR 11.3%
- BER 11.8%
- BES 28.5%
- FES 13.7%
- HEP 12.4%
- NP 22.3%

- Pls 15.2%
- Proj Mgrs 15 %
- Users 69.8%

2007/2008 Survey Questions

- 128 satisfaction questions scored on a 7-point scale
- average score: 6.07

Satisfaction score	meaning	Num times selected
7	Very satisfied	4,985
6	Mostly satisfied	3,748
5	Somewhat satisfied	832
4	Neutral	584
3	Somewhat dissatisfied	251
2	Mostly dissatisfied	75
1	Very dissatisfied	51



Areas of Highest Satisfaction

Area	Num responses	Scores (7 = very satisfied)
Account Support and Consulting	241 - 310	6.55 - 6.71
NGF Reliability and Uptime	65 - 66	6.67 - 6.68
HPSS Reliability and Uptime	96 - 111	6.54 - 6.66
Network Performance within NERSC	111	6.59
Jacquard, Seaborg and Bassi Uptimes	82 - 112	6.54

Areas of Lowest Satisfaction

Area	Num Responses	Scores 4 = neutral 5 = somewhat satisfied
Bassi Batch Wait Time	183	4.46
Franklin Uptime	257	5.04
Franklin I/O Performance	233	5.15
Training: web and onsite classes	60 / 51	5.30 - 5.39
Visualization Software on Franklin, Bassi, Seaborg	37 - 65	5.34 - 5.37
Jacquard Batch Wait Time	126	5.47
Data Analysis and Visualization	231	5.48



Areas of Increased Satisfaction from 2006

Area	Num Responses	Scores 7 = very satisfied 6 = mostly satisfied	Score Change
Seaborg Batch Wait Time	138	5.53	+0.59
24 x 7 Support	161	6.35	+0.31
Seaborg Uptime	138	6.54	+0.30
Seaborg: Overall	143	6.34	+0.25
Available Software	398	6.22	+0.24



Areas of Decreased Satisfaction from 2006

Area	Num Responses	Scores 6 = mostly satisfied 5 = somewhat satisfied 4 = neutral	Score Change
Bassi Batch Wait Time	183	4.46	-1.39
Jacquard Visualization Software	38	5.58	-0.54
Jacquard Batch Wait Time	126	5.47	-0.40
Bassi Queue Structure	176	5.57	-0.35
Bassi: Overall	187	5.96	-0.30



Where do you do Vis and Data Analysis?

location	2005	2006	2007/2008
All at NERSC	6.7%	3.9%	7.2%
Most at NERSC	9.7%	13.8%	17.4%
Half at NERSC	20.8%	15.7%	17.4%
Most elsewhere	26.2%	35.4%	27.1%
All elsewhere	33.3%	28.0%	28.2%
Don't need	4/1%	3.1%	8.8%



What does NERSC do well?

150 responses, from which:

- **68.7% - provides powerful resources for their science**
- **39.3% - excellent support services, responsive staff**
- **16% - good software support, easy to use environment**

The NERSC facility is fantastic. I'm very pleased with the hardware available, the people, the help, and the queues.

Good computing. Good storage. We always need more.

What NERSC is best at is the combination of large-scale computing facilities with more flexible queuing policies than in other comparable facilities. Also the existence of "small-scale supercomputers" (Jacquard) is very useful to make tests.



What should NERSC do differently?

108 responses, from which:

- **22.2% - improve Franklin services**
- **16.7% - improve job scheduling, resource allocation**
- **14.8% - provide more or different computing resources**

It would be great if NERSC could magically improve the stability of Franklin... Unfortunately, hardware failures increase with the size and complexity of the system.

Bassi has become so busy as to be almost useless to me.

More disk space to users. The whole point of having a LARGE cluster is to do LARGE simulations. That means LARGE amounts of data. We should get more storage space (individually).

Another very useful improvement would be to have more memory per core.



How does NERSC compare to other centers?

104 responses, of which:

- 59.6% - NERSC is excellent or better**
- 23.1% - NERSC is comparable or mixed review**
- 10.6% - NERSC is not as good**
- 6.7% - no comparison made**

NERSC is the best supercomputer user facility I have worked with. It provides the best user services and has an enormous software repository.

Comparison with ORNL Jaguar

Pro: More user-friendly queue policies, in particular for small-scale jobs

Con: Slower execution time

I was surprised to see that Franklin is about 4 times slower than EMSL's MPP2 for computational chemistry runs.