2012 NERSC User Survey Results

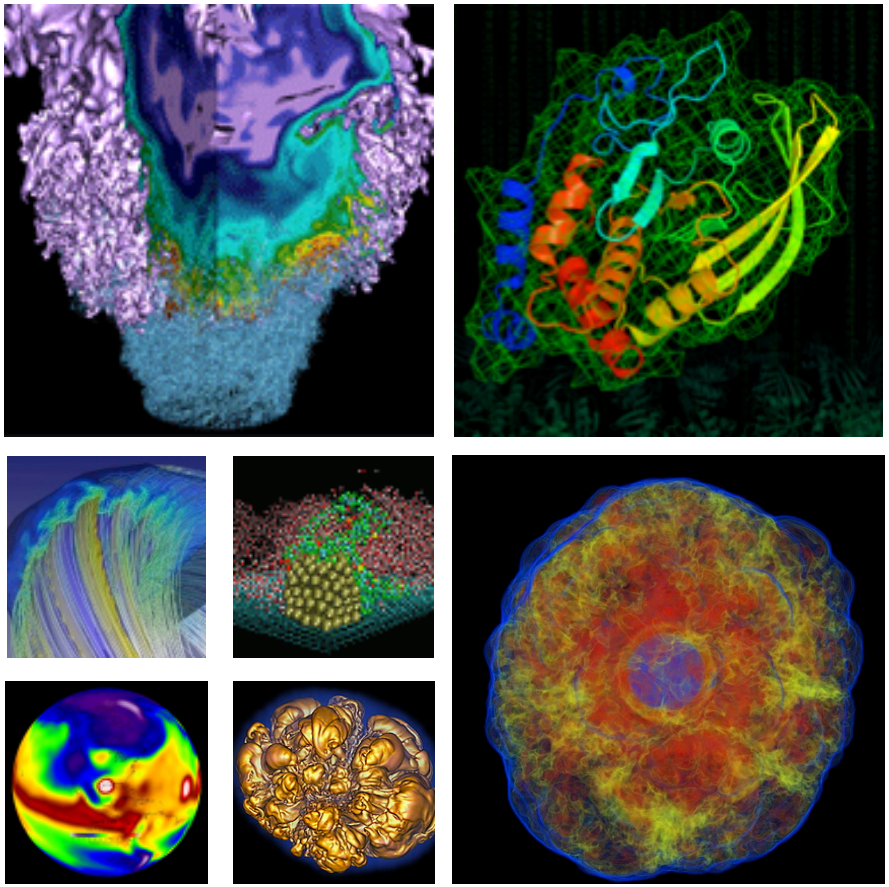

Francesca Verdier
fverdier@lbl.gov

Response Profile

481 respondents (+ 71 JGI only)

- **67.6% "big user" response rate**
- **36.2% "medium user" response rate**
- **11.8% overall response rate**

Respondants by Office

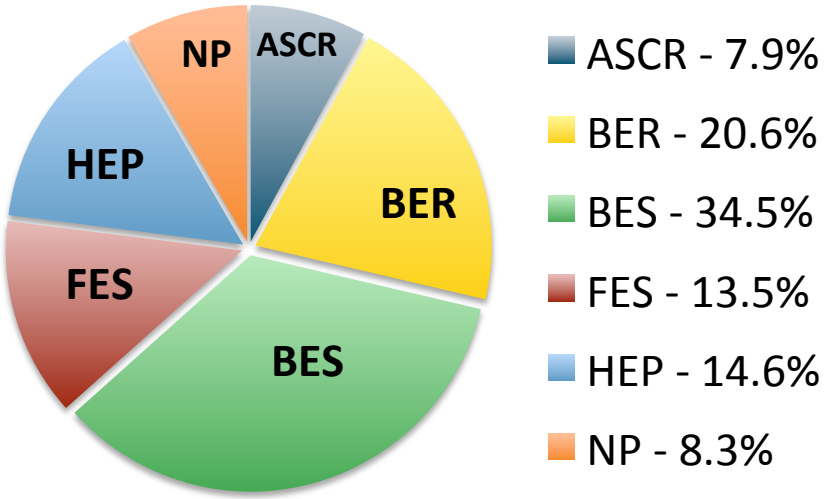

Respondants by Role

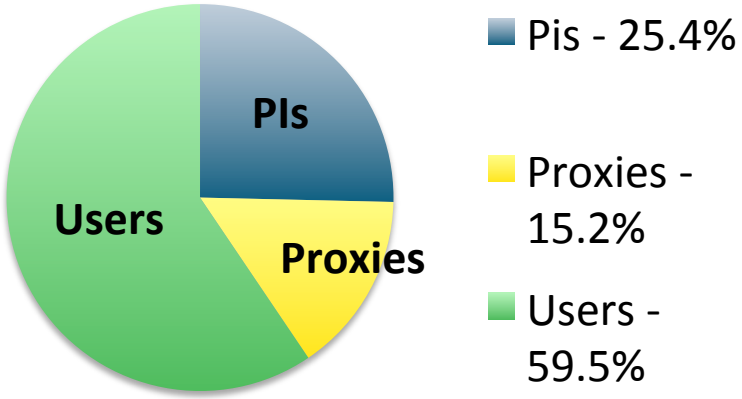

2012 Survey Question & Scores

- 97 satisfaction questions scored on a 7-point scale
- Average score: 6.32 (excludes JGI only)
- Minimum satisfactory score 5.25

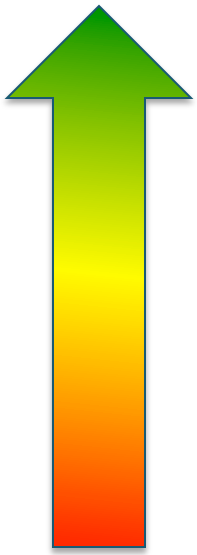

Satisfaction score	meaning	Num times selected
7	Very satisfied	10,843 (57.1%)
6	Mostly satisfied	5,477 (28.8%)
5	Somewhat satisfied	1,264 (6.7%)
4	Neutral	898 (4.7%)
3	Somewhat dissatisfied	353 (1.9%)
2	Mostly dissatisfied	95 (0.5%)
1	Very dissatisfied	58 (0.3%)

Response Profile

	Num	Resp. Rate	Avg. Score	% MPP hrs	< 1 year	1-3 years	> 3 years
Big MPP	98	67.6	6.32	47.3%	7.4%	42.1%	50.5%
Medium MPP	197	36.3%	6.29	13.0%	10.8%	45.6%	43.6%
Small MPP	149	8.1%	6.35	0.6%	30.4%	40.5%	29.1%
All (inc. JGI)	552	11.8%	6.22	60.8%	20.1%	44.4%	35.5%

Key Ratings	Big MPP	Med MPP	Small MPP	All
Satisfaction with NERSC	6.53 (-0.01)	6.53 (-0.01)	6.40 **(-0.14)**	6.32 **(-0.21)**
Available Hardware	6.16 **(-0.32)**	6.25 **(-0.23)**	6.25**(-0.23)**	6.15 **(-0.33)**
Services	6.67 (0.07)	6.60 (-0.00)	6.47 (-0.13)	6.42 **(-0.18)**
Available Software	6.31 **(0.21)**	6.22 (0.11)	6.12 (0.02)	6.08 (-0.03)
Mass Storage Facilities	6.12 (-0.06)	6.20 (0.02)	6.09 (-0.09)	6.32 **(-0.21)**
Consulting Overall	6.64 (0.05)	6.69 (0.10)	6.65 (0.05)	6.61 (0.03)
Web www.nersc.gov	6.48 (0.14)	6.57 **(0.23)**	6.42 (0.07)	6.42 (0.08)
NIM nim.nersc.gov	6.51 (0.16)	6.59 **(0.25)**	6.53 **(0.18)**	6.47 **(0.13)**

Overall User Satisfaction with NERSC Continues to be High (excludes JGI only)

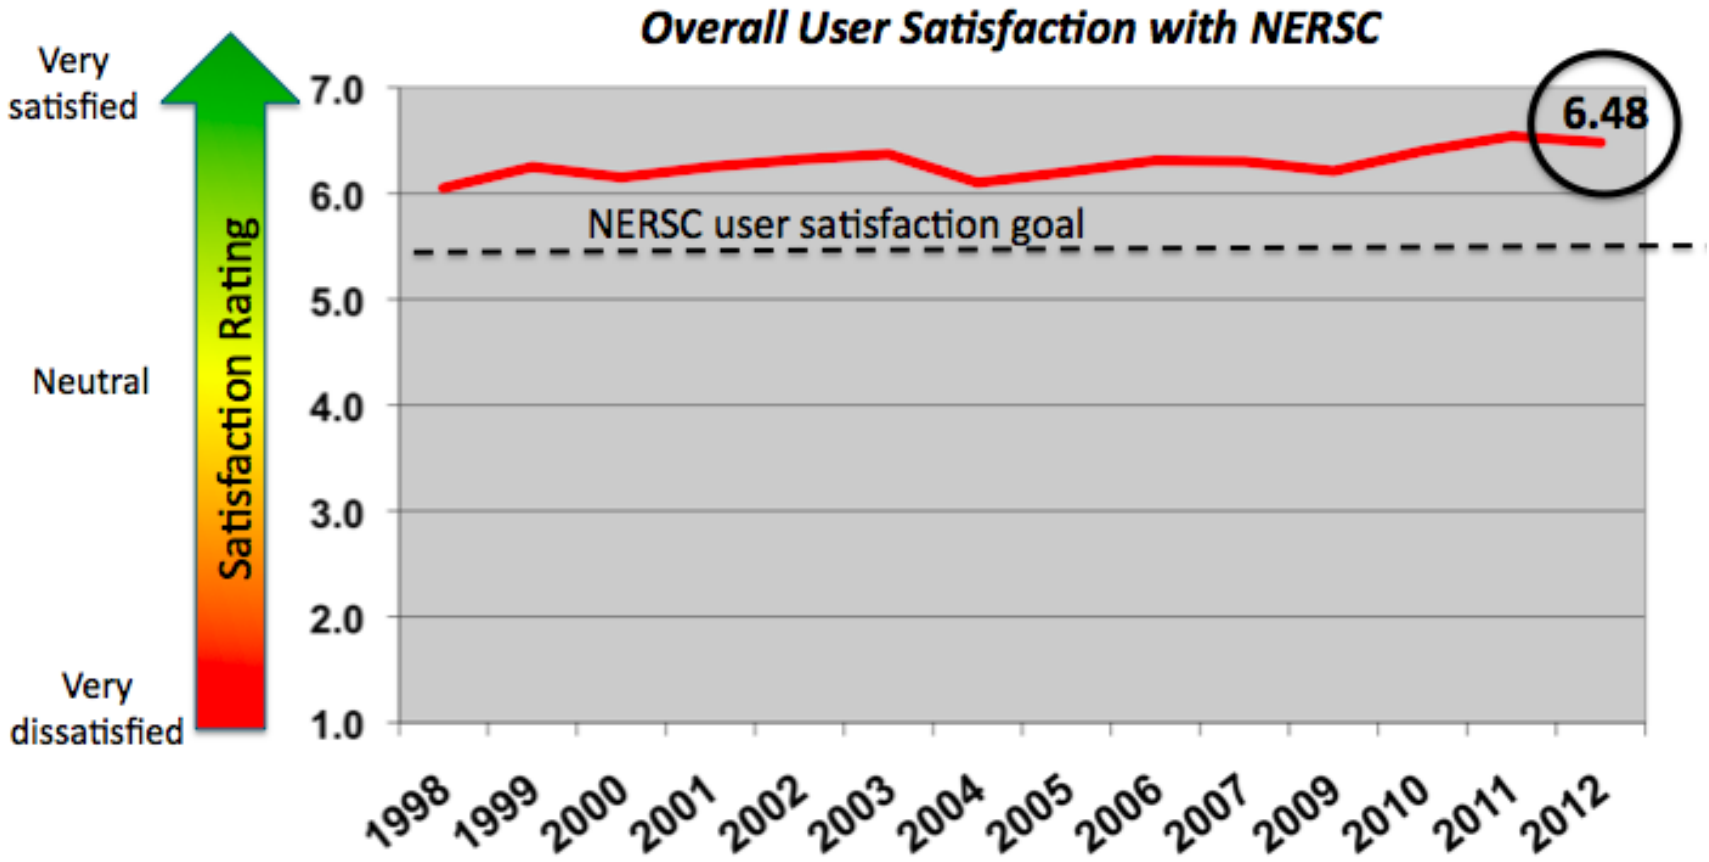

Areas of Highest User Satisfaction

(Scores > 6.5; out of 481 respondents)

Topic	Num Resp.	Score
HPSS uptime, reliability	198, 196	6.71, 6.69
Project reliability, overall, uptime	179, 185, 180	6.69, 6.59, 6.58
Global Homes reliability, uptime	280, 283	6.68, 6.62
Account Support	383	6.67
Security	324	6.66
Consulting overall, response time, technical advice, special requests	350, 349, 339, 215	6.65, 6.60, 6.57, 6.52
Web – accuracy of information, My NERSC, System Status, NIM	347, 299, 313, 371	6.56, 6.55, 6.53, 6.52
Global Scratch uptime, reliability	232, 228	6.56, 6.53
Training – New Users Guide	222	6.55
Services overall	450	6.54
Hopper uptime	375	6.54
Network performance within NERSC	260	6.53

Areas of Lowest User Satisfaction

(Scores < 6; out of 481 respondents)

Topic	Num Resp.	Score
Hopper batch wait time; batch queue structure	369, 367	4.90, 5.86
Web – mobile site; ease of use with mobile devices	49, 54	5.45, 5.63
Carver batch wait time; batch queue structure; ability to run interactively	153, 152, 114	5.64, 5.80, 5.95
NX overall	90	5.69
Visualization software	165	5.81
Euclid overall	49	5.82
HPSS user interface	198	5.85
Dirac GPU testbed – ability to run interactively	25	5.92
Training presentations; video tutorials	94, 52	5.95, 5.98

Areas of Increased Satisfaction 2011 →2012

(Scores with significant increases; out of 481 respondents)

Topic	Num Resp.	Score	Change
Carver batch wait time	153	5.64	+ 0.48
Training – New Users Guide	222	6.55	+ 0.32
Training – NERSC classes	83	6.19	+ 0.29
Data analysis software	164	6.01	+ 0.26
Web - searching	253	6.02	+ 0.24
Web – ease of finding information	366	6.29	+ 0.22
HPSS – overall satisfaction	209	6.49	+ 0.19
NIM (NERSC Information Management)	371	6.52	+ 0.17
Training – web tutorials	150	6.39	+ 0.16
Web – overall (www.nersc.gov)	385	6.49	+ 0.15

Note: The score for Carver batch wait time was 5.16 in 2011 – the only score for which we needed to show improvement ✔

Areas of Decreased Satisfaction 2011 →2012

(Scores with significant decreases; out of 481 respondents)

Topic	Num Resp.	Score	Change
Hopper batch wait time	369	4.90	- 0.96
Available computing hardware	471	6.28	- 0.20
Hopper overall	375	6.29	- 0.18
Hopper batch queue structure	367	5.86	- 0.17
Global Homes overall	291	6.47	- 0.14
Global Homes reliability	280	6.68	- 0.12

Areas of Most Importance to Users

(as shown by number of responses; out of 481 respondents)

Topic	Num Resp.	Score	Change
Overall satisfaction with NERSC	478	6.48	-
Available computing hardware	471	6.28	- 0.20
Services overall	450	6.54	-
Mass storage facilities overall	397	6.18	-
Web www.nersc.gov overall, ease of finding information	385, 366	6.49, 6.29	+ 0.15, + 0.22
Software environment, available software, programming libraries	384, 382, 363	6.37, 6.22, 6.36	-
Account support	383	6.67	
Hopper up time, overall, batch wait time, batch queue structure	375, 375, 369, 367	6.54, 6.29, 4.90, 5.86	- , - 0.20, - 0.96, -0.17
NIM – nim.nersc.gov	371	6.52	+ 0.17
Consulting overall	350	6.55	-

Advanced Architectures & Programming Models

Architecture	GPUs	Multi Threaded	MIC	IBM Cell
Big MPP	22.4%	9.2%	5.1%	3.1%
Medium MPP	20.3%	14.2%	3.6%	2.5%
Small MPP	22.8%	17.4%	6.0%	2.7%

Program. Model	OpenMP	CUDA	Pthreads	CUDA Fortran	OpenCL	OpenACC	UPC
Big MPP	49.0%	12.2%	3.1%	7.1%	1.0%	5.1%	0.0%
Medium MPP	42.1%	16.8%	10.7%	8.1%	5.6%	1.5%	1.0%
Small MPP	38.3%	19.5%	14.1%	3.4%	6.7%	3.4%	4.7%

What Does NERSC Do Well?

	Big MPP	Medium MPP	Small MPP	Total
User Support, staff	28	46	33	107
Hardware, HPC resources	19	32	20	71
Well managed center, allows science, all	11	20	13	44
Uptime, reliability	17	14	11	42
Software support	9	16	8	33
Web, documentation, training	6	14	10	30
Data, I/O, networking	6	9	7	22
Batch structure, policies	6	8	6	20
Communications to users	4	10	2	16
Security, ease of use, account mgt	4	5	2	11
Allocations	3	3	3	9

What Does NERSC Do Well – Sample Comments

NERSC provides reliable massively parallel resources and storage, supported with high quality consulting and software, and provides users the year to year continuity they need to stay productive. The staff accommodates special needs, and I am delighted overall. I only wish more time were available, and queue wait times were shorter. **This continuity is very important -- having a safe place to archive data, and not having to recode for a new system every year is essential.**
Support. I love having **24 hour support by people who actually know what they are talking about**.
NERSC provides highly professional and reliable computing with a broad array of capabilities. This to me is its primary strength. **The ability to have many members of a single, international, collaboration work in a common environment and common space and to apply a wide range of computing to shared data has been key** to the surveys I have been a member of. The access to high performance, massively parallel computing for simulations and simultaneously to smaller machines for analysis has been extremely helpful.
NERSC is the flagship of the DOE. **Best people, facilities, policies of fairness**.

Does NERSC Provide What You Need? What Else Do You Need?

	Big MPP	Medium MPP	Small MPP	Total
Yes	25	22	31	78
Need more cycles, HPC resources	9	16	7	32
Need different architectures (Vis cluster, more memory per core, large shared memory, interactive with no limits. mid range, testbeds	6	11	8	25
Different batch policies (HTC, longer wall, quicker turnaround for low, higher job limits, better support for mid range, more transparent, more fair)	7	6	4	17
More software support (e.g. python)	1	7	5	13
Need more data (local disk, automated data imports, larger quotas, faster file access	5	4	1	10
More documentation, training	1	1	3	5

What Else Do You Need? – Sample Comments

Hopper is a capability system optimized for *huge* jobs. It would be nice to **have a capacity system that would allow high-throughput of medium size jobs** (e.g. of order 1000 cores instead of order 20000 cores). It is incredibly difficult to push jobs with ~3000 cores. Jobs with < 512 cores and much larger than 10000 cores seem to start just fine. Our jobs are mid-size and take more time in the queue than they actually request in walltime. This is very frustrating.
I prefer a **larger memory on carver to run matlab for visualization**. The speed to access files is sometimes very slow so it would be nice to speed up it. I would prefer a larger home directory disk memory.
The biggest issues I have are **queue waits (way too long)** and **variation in wallclock run time**. It takes about 2 days to run a 3 hour job, so my throughput is not very good. I think variability in job wallclock is mostly due to I/O, esp run during day vs night, making it harder to keep my est time for batch jobs accurate. Also have had issues with **slow filesystem response** from hopper frontend nodes, both on GSCRATCH, SCRATCH, and home areas. Often compile times are very slow (not always) and things like ls, cd, or just getting your login prompt can take quite a while. These are not on directories with thousands of files. My guess is all of this is **due to lack of physical resources**. All my experiences with nersc help and staff have been excellent.
No, we are never satisfied. We need **more compute power and more on-line disk storage**, and we need it yesterday.