2012 XSEDE User Satisfaction Survey – Final Report

Prepared by Julie Wernert, Indiana University

16 July 2012

version 1.0

Table of Contents

A. Document History iii
B. Document Scope iv
C. Document Body – XSEDE 2012 User Satisfaction Survey Results 1
C.1. Survey Process Overview and Methodology 1
C.1.1. Survey implementation 1
C.1.2. Information Regarding Sources of Survey Error 2
C.2. Respondent Demographics 3
C.3. Adequacy, Utilization, and Knowledge of Cyberinfrastructure and XSEDE 5
C.4. Satisfaction with XSEDE Resources 15
C.4.1. Satisfaction with Computation and Data Resources 15
C.4.2. Satisfaction with XSEDE Software 18
C.4.3. Computational Resources - Science Gateways 19
C.4.4. Future Computational, Data, and Gateway Utilization Plans 21
C.4.5. Accelerators and Co-Processors 22
C.5. Satisfaction with XSEDE Support Services 23
C.5.1. User Training and Development 23
C.5.2. User Support Staff 26
C.5.3. XSEDE User Portal 27
C.5.4. User Support Documentation 28
C.5.5. Allocations 31
C.5.6. Extended Collaborative Support Services 32
C.5.7. Extended Collaborative Support Services – Future Needs 33
C.6. Final User Comments 35
D. Appendices D-1
D.1. Appendix 1 – Recruitment and Reminder Letters D-2
D.1.1. Appendix 1a – Initial Recruitment Letter D-2
D.1.2. Appendix 1b – First Reminder Message D-3
D.1.3. Appendix 1c – Second Reminder Message D-4
D.1.4. Appendix 1d – Final Reminder Message D-5
D.2. Appendix 2 – Survey Interface (Introductory and Conclusion Pages, screenshots) D-6
D.3. Appendix 3 – Study Information Sheet (screenshot) D-7
D.4. Appendix 4 – Representative Question Interfaces (screenshots) D-8
D.5. Appendix 5 - Complete Survey Text D-10

A. Document History

Relevant Sections	Version	Date	Changes	Author
Entire Document	0.9	07/12/2012	Baseline document for internal review	J. Wernert
Section C, Appendices	0.91	07/15/2012	Text comments moved from appendices to inline tables; Full survey text, correspondence, and screenshots added to appendices	J. Wernert, C. Stewart
Entire Document	1.0	07/16/2012	Applied official XSEDE document formatting	J. Wernert
Entire Document	1.1	8/13/2012	Minor edits	J. Wernert, Glenn Brook

B. Document Scope

This document summarizes the responses to the 2012 XSEDE User Satisfaction Survey planned throughout the spring of 2012 and conducted with XSEDE users and NSF grant awardees in May-June, 2012. This was the first annual survey of XSEDE users, and while it borrowed from the methodology and types of questions from TeraGrid User Surveys from previous years, it cannot be directly compared to those prior results. Successive years of this survey can be compared and analyzed with the results in this document to legitimately identify important trends, emerging needs, and notable changes in user satisfaction.

As with any survey instrument or resulting report, one should exercise caution in reading too much into specific results, either positive or negative. The authors made every effort to accurately summarize and convey the survey results they received so as to not introduce any bias. Readers should pay specific attention to the survey methodology detailed in Section C.1, especially sources of survey error described in Section C.1.2. Moreover, readers should frame their interpretation of responses in the context of the respondent demographics detailed in Section C.2.

Please direct any questions regarding the methods used in the administration of this survey or the summarization of responses provided in this report to Julie Wernert at Indiana University, jwernert@iu.edu .

C. Document Body – XSEDE 2012 User Satisfaction Survey Results

C.1. Survey Process Overview and Methodology

Work began on the 2012 XSEDE User Satisfaction Survey in mid-February 2012, with the aim of defining a new process for evaluating how users perceive and use XSEDE services (systems, software, support, etc.), what they recommend for changes and improvements to XSEDE services, and how those perceptions and inputs change over time.

Drawing on lessons learned and data gathered from previous TeraGrid surveys, a new, comprehensive survey instrument was developed and sent two distinct populations. The survey instrument is similar, in parts, to past surveys, but divergent based on the population being surveyed.

The two populations surveyed are:

1. XSEDE users: categorized as PIs, postdoctoral fellows, research staff, graduate students, undergraduate students, and XSEDE staff conducting work on XSEDE resources using XSEDE allocations.
2. Researchers who are federally funded PIs but not XSEDE users: Researchers funded as a PI by the NSF between 2007-2011, inclusive, but who do not have accounts or allocations on XSEDE.

Survey results from the first population—XSEDE users—aim to provide information about user satisfaction with XSEDE systems and services, specifically:

- To what extent users of XSEDE find the services offered by XSEDE satisfactory in terms of capability and availability to support their research needs.
- What capabilities do users wish XSEDE would provide, and how would these services expand their research capabilities and scientific impact?

Data from the second population aims to in some ways determine if XSEDE users are better provisioned with cyberinfrastructure resources and support than researchers who are not users of XSEDE. PIs who are XSEDE users could provide insight regarding the value of XSEDE services and unmet cyberinfrastructure needs within the US research community.

Working with members of the broader XSEDE User Engagement team, including Vincent Betro, Jay Boisseau, Glenn, Brook, Katie Cohen, Maytal Dahan, Chris Hempel, Therese Miller, Craig Stewart, and Justin Whitt, and various service owners, and beginning with the questions used in past TeraGrid surveys, a new survey instrument was developed.

C.1.1. Survey implementation

The web survey was launched on May 1, 2012, after receiving approval from XSEDE leadership and the IU Institutional Review Board, and closed on June 11, 2012. The Indiana University Center for Survey Research (CSR) administered this survey for Principal Investigator Craig Stewart, Indiana University Office of the Vice President for Information Technology and CIO. The study was funded by the National Science Foundation (NSF). It was programmed using Qualtrics Web Survey Software. The sample and email messages (initial invitation and subsequent reminders) were managed within the CSR SQL database.

The sample population included: (1) 5,000 XSEDE users selected from a list of some 13,000+ current XSEDE users provided by the XSEDE organization, and (2) 5,000 NSF principal or co-principal investigators funded between 2007 and 2011, inclusive, and who were not current XSEDE users; this portion of the population was generated from the publicly-available NSF award database and compared to the list of XSEDE users to eliminate duplication.

The survey sample population was generated randomly and all responses were voluntary. Responses from those completing at least 50% of the survey were included in the final data. As of June 11, 2012 there were 734 responses, representing just over 7% percent of the survey population. This is lower than the anticipated 10% response rate. It should be noted that a similar percentage of survey respondents started the survey but did not finish; this may be indicative that the survey is too long or its scope too broad. It should also be noted that respondents were not required to answer any question; that is, respondents could skip questions if they did not wish to answer for any reason without voiding other responses.

C.1.2. Information Regarding Sources of Survey Error

Surveys of this kind are sometimes subject to types of inaccuracies for which precise estimates cannot be calculated. For example, findings may be influenced by events that take place while the survey is in the field. Events occurring since the time the surveys were completed could have changed the opinions reported here. Sometimes questions are inadvertently biased or misleading. The views of people who responded to the survey may not necessarily replicate the views of those who refused to respond to the survey.

C.2. Respondent Demographics

Respondents were overwhelmingly white, non-Hispanic, males representing the engineering, chemistry, computer and information sciences, biology, and physics.

What is your race?	Frequency	Percent
American Indian or Alaska Native	5	0.7%
Asian	164	22.3%
Black of African-American	10	1.4%
Native Hawaiian or other Pacific Islander	3	0.4%
White	484	65.9%
No Answer	68	9.3%
Total	734	

What is your ethnicity?	Frequency	Percent
Hispanic or Latino	40	5.4%
Not Hispanic or Latin	627	85.4%
No Answer	67	9.1%
Total	734	

What is your gender	Frequency	Percent
Male	528	71.9%
Female	145	19.8%
No Answer	61	8.3%
Total	734	

,

Respondents' primary research field	Frequency	Percent
Engineering	117	16.3%
Chemistry	106	14.7%
Computer and Information Science	93	12.9%
Biology	85	11.8%
Physics	79	11.0%
Astronomy	47	6.5%
Mathematics	47	6.5%
Earth Science	41	5.7%
Atmospheric Sciences	21	2.9%
Psychology	11	1.5%
Education	7	1.0%
Medicine	6	0.8%
Politcal Science	6	0.8%
Sociology	6	0.8%
Economics	5	0.7%
Health and Wellness	3	0.4%
Diseases	1	0.1%
Other	38	5.3%
Total	719	

Of the respondents who identified their primary role within their current organization, the vast majority were faculty, research scientists/post-docs, or graduate students. Of these, over 50% identified themselves as current principal investigators or co-investigators on an NSF-funded award.

Respondents' primary organizational role	Frequency	Percent
Faculty	353	49.1%
Research scientist/Postdoctoral fellow	175	24.3%
Graduate student	113	15.7%
Executive director/Administrator	30	4.2%
Analyst/Programmer (staff)	16	2.2%
Project manager	10	1.4%
User support (staff)	7	1.0%
Undergraduate student	5	0.7%
System administrator (staff)	4	0.6%
Other	6	0.8%
Total	719	

What is your status relative to funding support from the National Institutes of Health (NIH)? (Check all that apply)	Frequency
I am currently funded by an NIH award	73
I am currently a Principal Investigator on an NIH award	20
I am currently a Co-Investigator on an NIH award	28
I am not currently funded by an NIH award, but I have been within the past five years	47
*Total # of selections / # of potential respondents	168 /734

What is your status relative to funding support from the Department of Energy (DoE)? (Check all that apply)	Frequency
I am currently funded by a DoE award	77
I am currently a Principal Investigator on a DoE award	24
I am currently a Co-Investigator on a DoE award	19
I am not currently funded by a DoE award, but I have been within the past five years	52
*Total # of selections / # of potential respondents	172 / 734

C.3. Adequacy, Utilization, and Knowledge of Cyberinfrastructure and XSEDE

Considering the entire survey population – both XSEDE and non-XSEDE users – over 70% indicated having adequate access to cyberinfrastructure resources either most or all of the time, and just over 50% indicated having adequate access to technical support most or all of the time. These results map closely to those reported in the 2011 XROADS Survey[1] of 5,000 NSF investigators who, at the time of the survey, were not using TeraGrid.

I have access to adequate…	Average (1=never, 4=all of the time)	Number of Responses	Distribution (1=never, 4=all of the time)				
			1	2	3	4	Histogram
...CI resources	2.923	711	6.5%	23.3%	41.6%	28.6%	
...techical support	2.603	713	9.0%	38.6%	35.6%	16.8%	

Of that same group, over 40% indicated that the 2012 User Satisfaction Survey was how they learned of XSEDE. Surprisingly, only some 15% reported being legacy TeraGrid users.

How did you first learn of XSEDE?	Frequency	Percent
This survey	291	40.3%
Advisor or principal investigator	128	17.7%
Legacy TeraGrid user	115	15.9%
Colleague	107	14.8%
Internet	28	3.9%
Conference or workshop	19	2.6%
NSF announcement/website	19	2.6%
Other	15	2.1%
Total	722	

Of current XSEDE users, 65% have used XSEDE (or TeraGrid) for more than one year, with nearly a quarter of respondents having used XSEDE for six months or less. Over 40% of respondents report using XSEDE at least once per week in the last year.

How long have you used XSEDE (or TeraGrid) resources, and/or overseen the use of XSEDE (or TeraGrid) resources by others?	Frequency	Percent
Less than 6 months	103	23.8%
6-11 months	48	11.1%
1-2 years	112	25.9%
3-5 years	99	22.9%
More than 5 years	70	16.2%
Total	432	

[1] Stewart, C.A., D.S. Katz, D.L. Hart, D. Lantrip, D.S. McCaulay and R.L. Moore. *Technical Report: Survey of cyberinfrastructure needs and interests of NSF-funded principal investigators*. 2011. Available from: http://hdl.handle.net/2022/9917

How frequently did you use XSEDE and/or TeraGrid services in the past year?	Frequency	Percent
More than 50 times per year (once per week or more)	177	40.4%
25-50 times per year	63	14.4%
10-24 times per year	55	12.6%
5-9 times per year	36	8.2%
1-4 times per year	43	9.8%
I have never used XSEDE	47	10.7%
I have in the past, but no longer use XSEDE	17	3.9%
Total	438	

Those who are aware of XSEDE but are not presently using its resources indicate a variety of reasons for not doing so. The complete list of reasons provided follows on the next pages.

Why are you not currently using XSEDE?

No time to learn a new system/my current research doesn't require massive HPC resources.

Administrative duties increase at home institution

All computational needs have been met for my research by local resources; if this does not continue to be the case I will definitely use XSEDE

Allocation lapsed

Both myself and my colleague [name deleted] are working on transition metal system so any of us can run the programme

Current research is more "analytic" than "numeric"; will plan to use in teaching though

Do not currently have the need

Don't even know what it does

Don't specifically need those resources right now, although may want them in near term future

FutureGrid provides great services for me

Have not yet requested services

I am funded to develop software (iRODS data grid), which is used as part of XSEDE infrastructure

I am just beginning to use them

I am not aware what services are available from XSEDE

I am still testing my code; I plan to use it in a very near future

I don't have a startup allocation/project that requires the use of these services. Previously, I received a Teragrid startup project in Aug 2010. The focus was for computational chemistry field problem using software at NCSA (called Materials Studi from the company Accelrys). During the course of the year I was transferred to several NCSA computing systems (Cobalt and then Ember) as some of these retired. Due to my busy schedule, I barely used my startup computation time. At the end of one year I requested an extension for 6 months and also came to know that the NCSA was not sure if they were licensing the computational chemistry software. I couldn't go about this, as the software of interest was not available.

I don't know

I have recently retired.

I need to be using XSEDE again! But there is a substantial burden associated with training graduate students and postdocs to make use of resources of this nature. Furthermore, our great challenge right now is large memory-space computing, not so much compute time, so parallel architectures don't help a lot.

I was conducting research supporting community engagement for the LEAD portal, which used TeraGrid resources. I had an allocation for the purposes of trying out the portal as part of our support for the project. I do not do research that uses high-end computational facilities.

I'm just starting to learn about the XSEDE systems through collaborators that are more involved. That is I provide the scientific content and they do the actual using of the services...

Just learning about it

Just signed up

Lack of time and no urgent need to pursue alternatives to existing local environment

My computations are in a testing/developing stage rather than a production stage.

My programmer carries out the production run, not me personally although I oversee the effort and make the needed scientific choices for the run.

My Teragrid grant expired

Need more computational resources

Never heard of it.

No current need for resources.

No current research need

No need for large simulations.

No SUs

Not clear to me what I would use it for.

Not easily available or well defined

Not enough time to pursue

Not related to the project and my new PI does not write proposals to XSEDE

Not sure how I would

Not sure how!

Not sure what it can do.

Not sure if I need them.

Only recently heard of it.

Presently research opportunities as an undergrad don't require me to

Startup allocation expired; now writing regular research allocation request

The nature of our work has changed to where the computational resources we utilize locally are sufficient; however, we are working to extend our work into an area that would require considerably larger computational resources, and will likely consider XSEE as a potential resource.

The startup allocation expired already.

Too difficult to access and use. Allocations process is burdensome and redundant with grant peer review.

What is XSEDE?

Over 88% of respondents reported that their use of XSEDE is primarily related to research, and over 81% indicated that XSEDE was either very helpful or essential to their research.

Is your use of XSEDE primarily related to research or education/outreach?	Frequency	Percent
Primarily research	325	88.3%
Primarily Education	14	3.8%
Equal parts research and education	29	7.9%
Total	368	

How important is XSEDE to your research?	Frequency	Percent
Essential; I would not be able to pursue my current research program without its use	174	46.6%
Very helpful; I would have difficulty pursuing my current research program without its use	129	34.6%
Helpful; It is useful, but I could pursue my research program without it	59	15.8%
Neutral; It is neither helpful nor unhelpful to my current research program	9	2.4%
Unhelpful; It is presently of negative net value to my research	2	0.5%
Total	373	

When asked "How can we make XSEDE more useful to your research program?" users reported a wide variety of suggestions which are listed on the next 4½ pages.

How can we make XSEDE more useful to your research program?

Coordination of queues at different sites/unified access to files

Provide computing resources that grow with the growth in community demand/maintain a stable mass storage facility -- not one that will pull the plug in a few years time when the centers are recompeted.

Offer secure archival storage (like the soon-to-retire gpfs-wan)

Longer run time limits on Lonestar (>24 hrs) for normal # cores (<4104 cores) OR Access to >4104 cores for 24 hours

Install open source codes like openFOAM on all machines/have better info about viz capabilities

Access to more FEA software.

Allow for longer job-running times on HPC machines.

Allow resource requests at anytime for various levels of reasons (with match levels of resource allocations and use priorities).

Any training or documentation related to advance MPI topics and considerations for codes that scale to thousands of processors.

Assign more SU allocation

Automatically terminate jobs when the allocation is consumed.

Being able to run QIIME successfully across the nodes would be very helpful.

Better storage services. Better links to DOE and NIH computing programs.

Better support for basic biology high throughput pipelines (e.g. blast, interpro, etc)

Better support for efficient use of software provided for computational chemistry: advice on parallelization of codes, more computer configurations that allow efficient use of the Gaussian09 program.

Better support for WRF model compilation

Bigger machines

Broaden range of software available. And, of course, more cycles are always better.

By continuing to provide more workshops and outreach activities such that there are XSEDE centers in each state. These centers could be a point of knowledge dissemination and could be a university or any research medium. That way, users can directly interact and also attend workshops (video-conferencing). And if each center has a few consultants during the workshops, it would help a lot.

By providing more help/support for students staring to use an XSEDE allocation (beyond the training sessions which are just general). For example, help with setting up interactive runs to find out why programs that run fine on other computational resources do not write out data on XSEDE clusters or crash with no error message and no obvious mistakes (upon repeated attempts).

By using uniform file system in different machines

Categorize, catalog and add more searching capability.

Clearer instructions and orientation as to which servers to sign in for service units; this was my first contact with this level of cyberinfrastructure and I was not sure how to go about, and I am sure I wasn't the first. Also, allow for changes in such allocations after granting allocations.

Coming from Germany it was quite hard to reopen my account when terra grid reached its end of life. Probably nothing that can be done about this . . .

Congressional research funding is inadequate. New machines rapidly become saturated with users, which slows the progress of research for everyone. XSEDE

Continue to support GPU-enabled nodes

Deploying more computing resources

Everything I need is available

Fully support R (http://cran.r-project.org/web/views/HighPerformanceComputing.html)

Further improve the data storage facilities and the data transfer rates between storage facilities and computing resources.

Give longer lasting grants of time on machines; make information about the characteristics of the systems easier to access; have more standardized methods of access to the computing machines themselves. Prospective new users are often discouraged from using XSEDE because of the time required to learn how to use the equipment. Another problems is that more common application software for particular fields needs to be available and the information as to where it is located needs to be more easily found. There is a lot of verbiage on the portal that could be reduced with better design.

Give more CPU time and storage space

Good question. I will know better in 3 weeks after we run our first major data production. So far so good.

I am a happy XSEDE user.

I am quite happy with it. It would be nice if it were easier to move and store data (Globus often falls short, for example, and mass storage systems are often woefully inadequate for the systems they're on), but it's not really a deal-breaker at this point.

I am satisfied :)

I do a lot of software development and compilation on Kraken is very slow because the front end machines that I log into have very limited resources (I think only two cores each).

I don't know. It seems to be working well for us.

I have not thought about this. So far I have been very pleased with what it is doing.

I have used XSEDE (originally Teragrid) for three years. In the past two years, we used Abe and Ranger. We got excellent technical and consulting services. This year, we switched to Trestles where the technical service is much poorer than the two previous clusters. The technical advisors respond our installation requests very slowly. Some problems are not even solved today. The mainly explains why our current usage is only 0.4% of the allocated quota. The technical and consulting service is of great importance to our use of these clusters. Improving the technical and consulting service will greatly help our use on Trestles and other XSEDE clusters in future.

I honestly can't think of anything - it seems pretty perfect at the moment.

I require a relatively small number of SUs (~300K) per year. A simplified proposal procedure, similar to a start-up allocation request, for allocations of this size would be much appreciated.

I run WRF simulations. Knowing more about which of the platforms are optimized for this type of research would be useful.

I understand that many other groups are also using the computer resources. However, if we can sometimes shorten our time for waiting for a job to be submitted, it will be very great.

I would appreciate faster and easier tape storage and retrieval. Tape systems such as Kraken's HPSS and "Ranch" are rather slow and clunky. Otherwise I have been very satisfied with resources.

I'd really like to see more educational opportunities available to people who are a little new to XSEDE but do have some background with computational science. I have found most seminars etc. are usually for beginners or more advanced users. I'm sort of intermediate, so I feel that I don't get much out of those resources.
Also, (this is a little off topic) I've posted questions in the past after searching the site for quite awhile to see if I can find the answer on my own and I've never gotten a response. I know that people usually want to work on answering "more important" questions, but if the information isn't on the web site, then what am I supposed to do other than ask?

If more resource time can be approved, that would be perfect

If the turnaround or waiting time can be shortened, it will be more useful to us.

If the XSEDE can give me more resource time, that would be very helpful to my research, since my research is related to super large scale simulations

If XSEDE can help in getting NSF funding

Improve file systems on clusters (lower latency)

Improve the reliability of Trestles.

In ongoing research projects it is difficult to predict required computer resources; an easy mechanism to switch resources or add resource would be appreciated.

Increase capacity to enable increased allocation. Enable individual allocations for graduate students.

Increase computing power of available resources, and give more startup allocation than now.

Increase grant request frequency (though this is pertaining to my particular current situation and might not be feasible).

Increase the computational time available.

Increase the computer resources by adding more fast nodes. The queue is quite long now.

Increase the processor and memory capacity of individual machines: Kraken, Ranger, etc.

Increase the variety of hardware setups. For example, it's very difficult to find a machine with significant local disk storage on each node, which is what we need for optimal performance.

Installation of more bioinformatics programs (genome assembly, RNA-seq pipelines, etc)

It already fully meets my needs.

It currently can be very difficult to get more allocation, especially as a student.

It is a bit time consuming to get a fob for new students or postdocs in my group.

It is great the way it is.

It is not clear to me how archival storage resources can be added to computing resource requests, especially when there is no archival storage system at the site of the compute allocation.

It is really good. I don't have any suggestion off the top of my head.

It is very good as it currently is

It should be more flexible, for example we are not allowed to run long queue jobs which would have been very helpful if we were allowed.

It will be better if ordinary researcher could apply XSEDE easier.

It works well for me, so I don't know you could make it more helpful

It would be good to have all the software updated quickly.

It's great as is for what I need.

It's great, except the queuing algorithm on Kraken is annoying. By having the largest job have the priority, sizable jobs requesting, for instance 2,000 cores for 8 hours, can sit in the queue forever.

Keep adding state of the art architectures and storage to complement our local resources.

Keep up to date software available. Make it possible for users to suggest upgrades to existing software (e.g., blender for visualization). Many GPU-containing machines have the GPUs switched off for months at a time making the machine essentially useless or its stated purpose. Machines are going down all the time -- terrible continuity.

Larger disk quotas

Less pressure to use up the allotted machine time while I am still investigating my problem.

Lonestar (TACC): It would be great if maximum runtime for normal jobs is 48 hr instead of 24 hr

Ranger (TACC): It seems to be much slower than Lonestar and Trestles

Trestles (SDSC): The system is very unstable, but it is very good when it's working well. So many small jobs are using just one node for long time with normal queue instead of shared queue, which I think is not the way cluster computers should be used.

Lower barriers to use: simplify grant procedures, shorten turnaround time, reduce technical obstacles (increase uniformity and improve documentation across systems or reduce "musical chairs" arising from system turnover; better and more responsive support for system problems).

Machines are extremely powerful and capable of running large jobs on large number of processors. Reducing the queue waiting time for jobs when submitting parallel jobs to machines can be of great help. Some machines are extremely easy to use. But many aren't. Either creating a common way of accessing all the machines or providing for less complex methods of usability will be of help too!

Make access easier to get

Make computational resources with more computational chemistry software available to the users, make more machines available.

Make it easier for new users to be added to a project.

Make it easier to add new users; have had delays in getting passwords set up.

Make it easier to obtain and maintain allocations.

Make it easier to submit proposals for SUs.

Make it easier/faster to extend the allocation time for a specific project if it runs out on Teragrid

Make more detailed documentation of resources

Make more types of software (in our case electronic structure calculations for molecules and for extended systems) available at every XSEDE site

Make queues, compilers, and long-term storage access more uniform across resources and sites so that the learning curve on going from machine to machine is smaller.

Make the website load faster

More allocations

More asynchronous training materials for scientist users

More complete documentation in the user guides and perhaps web/on-line versions of the training sessions that are offered that can be followed at an individual's own pace and time availability.

More CPU hours, longer wall-time runs.

More efficient technical support

More flexibility in the queue system. Sometimes I have high-priority jobs, but I may have to wait for a long time for them to start execution.

More flexible allocation -- increase as you go instead of having to write supplemental proposals

More GPUs + less huge machines + less e- bill = more $ to send interests to DC to make politicians care more about our/XSEDE mission. Just an idea...

More memory-intensive machines to access

More new computational resources, better and faster support on new resources.

More regular reporting of research being conducted on XSEDE resources with feedback from support staff or management regarding synergy with related software capabilities.

More representative in local school, perhaps

More services and solutions for high performance computing (or how to use supercomputers)

More SUs

More training for first time users with little or no computational experience. Also Gaussian should be available on all platforms.

More user home directory space on NICS resources (>2GB)

My biggest concern is my data. I would like a guarantee of long-term data storage, independent of my (current) allocation status and access to the data generated as part of my research.

My impression is that XSEDE provides ample resources and adequate opportunities to apply for these resources. My only complaint is that some of the resources to which we've had access have not scaled well (e.g. ranger), but this is likely just a matter of asking for access to other systems.

Great right now!

Need XSEDE to support persistent Science Gateways

No way to make it more useful.

Not much at this point

Of great use would be a brief consultation with a staff member with deep knowledge of XSEDE resources prior to submitting an allocation request, to better match my science needs with available resources.

Offer more nodes and reduce the queue waiting time

One of my students never could get our code running on Steele. The lack of a debug queue, and the lack of support or documentation, kept us from figuring out how to launch jobs and collect data at the end.

Our award is on Steele, and sometimes the XSEDE/TeraGrid queues are so clogged that it takes 1-2 weeks for my submitted jobs to run.

Other colleagues who use Steele also have the same complaint, and it just makes it difficult to use the resource when the wait time is often that long.

That is especially true for me when I have a very large number of jobs that take just 5-8 hours to run on a single node.

It would also be helpful to have a disk quota about twice as big on Steele.

With a larger disk quota, I could submit a larger number of jobs at a time before then compressing and moving the data off Steele, and that would speed up my research.

Please keep the NCSA unitree system working as smoothly as it has worked so far. Backups are essential and difficult to do elsewhere.

Possibly provide more on-campus "beginner" instruction in computational tools for biochemistry/chemistry - like the VMD/NAMD workshops. But what is done is good.

Probably making GPU SUs cheaper (they cost about 10 times the CPU SU)

Provide a far better on-demand experience and provide far better support for dynamic user definition of system images. I think the XSEDE organization tries t be too many things and as such fails to achieve success in any of them. The mission should be to deliver cycles and make it possible for users to do work.

Provide more computer cycles.

Provide more cycles

Provide more resources for visualization of data sets.

Provide more support to each specific project and flexibility on software installation

Provides a more uniform developing experience across all the architectures/systems available on Xsede

Provision of more computation units.

Queue waiting time too long.

Quit changing things (resources come and go; protocols for accessing resources change; job submission requirements change; etc.) It is not easy to learn how to deal with these things in the first place, and just when you get comfortable, it changes.

Relax wall-clock limit for a job if practical.

Resource user guides are limit to basic information and does not have placeholder for user to share their experience. A blog is must needed per resource or XSEDE wide to share user experience. New users can be benefited a lot from this service

Shorter queue waiting times, especially at Kraken

Shorter waiting time and more longer maximum wall time.

Simplify resource application/better document XSEDE user portal

Small interactive cluster for data analysis

So far it has responded well to all my needs, maybe more computational resources would make queues shorter and simulations faster

Spatial analysis tools, perhaps bring more of the CIGI, cyberGIS tools into XSEDE for immediate use by researchers. Plus make these known.

Streamlining the software installation across all the resources affiliated with XSEDE

The 24 hr wall clock limit on many of the resources makes using the resources unusually difficult compared to the resources of past. I would like to see at least three-day wall clock limits reinstituted, perhaps for lower core count jobs (maybe fewer than 56 cores). This would make XSEDE resources more useful to my research program.

The current state is fine to me.

The queuing system sometimes is bad. Maybe, optimizing the system such that small-scale jobs do not have to wait too long to be executed would be ideal.

The state of documentation is pitiful. Many links are outdated. User support ranks from adequate to totally unresponsive depending on the site.

The wait-time of a submitted job is sometimes very long.

The website is very convoluted and hard to navigate. It lacks certain useful functionality (e.g., transferring files, etc.).

Third party commercial application support needs to be strengthened. significant science is conducted by large number of users who use familiar commercial software and most resource providers(RP) are not enthusiastically supporting such users. XSEDE should strongly encourage RPs to widen such support.

To my knowledge, there is no tape system at the site (SDSC) where I have received my compute allocation.

I am able to use other tape resources provided by DOE funding, but the presence of a tape system at the XSEDE site where I am performing the bulk of y work would have been useful.

Undergraduate allocation option--I need less resources, use them less often; perhaps a streamlined application process for smaller "chunks" of time that can easily be renewed would be appropriate; expected outcome include undergraduate researcher coauthor and/or undergraduate presentations

Very often queues way too long, especially on Kraken

Wait times in the queue have often impeded our research, especially on Kraken.

We run simulation on a machine I have in my group for 30 days on 32 cores without interruption. We would like to have the facility to use machines for a longer time continuously. This will make a big difference to us.

We use the program Gaussian on PSC's Blacklight system. We have found that jobs run better when we use 16 cpus compared to 32 cpus, but jobs that request 32 cpus get preference in the queue. For example, if a 16 cpu job is submitted, you are lucky if the job starts within 24 hours. When a 32 cpu job is submitted, it usually starts within the hour. It would be helpful to us if 16 cpu jobs were not given such low priority.

We're trying to figure that out.

Getting data in the appropriate format for HPC national resources usage.

Would like the queue to go faster. some 3hr 800 core jobs take over 48 hours to even run.

XSEDE and Teragrid have already been extremely useful to my current research program.

XSEDE has been essential and a fantastic resource for my work! One minor issue I have run into is the somewhat old versions of various libraries available on some of the machines in XSEDE (such as graphical C++ library Boost and linear algebra library Eleental).

XSEDE is an excellent resource. I am currently using XSEDE for a complex genome assembly problem and we do not have the resources to perform the analysis at our home institution.

XSEDE is fine, we need more people working on our research area to produce more results.

XSEDE is very helpful to my research program. Short of eliminating queue waiting, not much can be improved.

XSEDE is very useful to the research program of my PI overall, but I am primarily affected in that it frees up other computational resources for my use since a large number of group members have access to the XSEDE machines

You are already doing your best

You're doing fine currently.

Overall XSEDE users reported being reasonably aware of most aspects of XSEDE. Users report the greatest levels of familiarity with XSEDE computational services, the XSEDE User Portal and the main XSEDE portal; lower levels of awareness were reported for visualization services and education and outreach activities.

Using a scale of 1 to 5, with 1 being "Completely unaware" and 5 being "Completely aware," reposndants rated their awareness of the following aspects of XSEDE:								
	Mean	Number of Responses	Distribution (1 = completely unaware, 5 = competely aware)					Histogram
			1	2	3	4	5	
Mission	3.46	422	10.9%	10.9%	23.5%	31.0%	23.7%	
Computational resources	3.95	427	5.2%	6.6%	13.8%	37.2%	37.2%	
Support/Consulting desk services	3.66	428	7.5%	11.0%	18.9%	33.6%	29.0%	
Training Opportunities	3.39	424	9.2%	16.5%	21.9%	30.4%	21.9%	
Education & outreach activities	3.09	427	14.3%	17.3%	30.2%	21.1%	17.1%	
Main portal (xsede.org) and other Web-based resources	3.90	428	5.8%	6.3%	18.0%	31.8%	38.1%	
XSEDE User Portal (portal.xsede.org)	4.01	427	6.8%	6.6%	12.9%	26.7%	47.1%	
XSEDE documentation	3.55	424	7.8%	11.1%	23.8%	32.5%	24.8%	
Storage services	3.48	423	7.3%	11.6%	27.2%	33.3%	20.6%	
Visualization services	2.99	426	15.3%	20.7%	28.2%	22.1%	13.8%	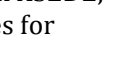

Users report a high degree of satisfaction with the resources and capabilities available through XSEDE, with notably high satisfaction for the capability (scalability) of XSEDE computational resources for simulation and data analysis, particularly for parallel processing.

Please rate how satisfied you are with the XSEDE resources listed below in the left column, where 1 is "Not at all satisfied" and 5 is "Very Satisfied" **								
	Mean	Number of Responses	Distribution (1 = not at all satisfied, 5 = very satisfied, **)					Histogram
			1	2	3	4	5	
Capability (scalability) of XSEDE computational resources for simulation, particularly parallel processing	4.29	331	1.5%	4.2%	8.5%	35.0%	50.8%	
Capability (scalability) of XSEDE computational resources for data analysis, particularly parallel processing applications	4.17	277	2.2%	3.6%	11.9%	39.7%	42.6%	
Capacity (in terms of high throughput computing) of computational resources for simulation	4.10	328	2.1%	4.9%	15.5%	36.0%	41.5%	
Capacity (in terms of high throughput computing) of computational resources for data analysis	4.06	270	1.5%	5.2%	18.1%	35.9%	39.3%	
Ability to utilize Science Gateways to access XSEDE resources	3.88	233	4.7%	3.4%	24.0%	35.2%	32.6%	
Data archiving capabilities of XSEDE resources	3.79	264	2.7%	8.7%	24.6%	35.2%	28.8%	
Visualization facilities and rendering capabilities of XSEDE resources	3.81	180	2.2%	9.4%	25.6%	30.6%	32.2%	
Ability to receive support/consulting services from XSEDE	4.02	322	3.1%	7.5%	15.5%	32.0%	41.9%	

*** scale has been transposed from original survey instrument to be consistent with higher numbers being positive and lower numbers being negative*

Satisfaction numbers were adjusted to reflect only those who evaluated a given service; those who indicated that their evaluation of a service was "not applicable" were extracted from the total number for the purpose of more accurately quantifying satisfaction levels.

XSEDE users report a near equal mix of local and XSEDE (or other NSF-funded grid resources) for their computational needs.

What percentage of the time do you use the following computational resources?							
	Number of Responses	Response Range					Histogram
		0% (never)	1%-25%	26%-50%	51%-75%	76%-100%	
Local Resources	371	*15.4%*	*26.1%*	*25.3%*	*9.2%*	*24.0%*	
State or Regional Resources	370	*84.4%*	*8.9%*	*5.1%*	*0.8%*	*0.5%*	
Commercial cloud resources	374	*93.8%*	*6.5%*	*0.3%*	*0.3%*	*0.0%*	
DoE-funded resources (e.g., INCITE program)	374	*88.7%*	*4.9%*	*5.7%*	*1.3%*	*0.3%*	
XSEDE or other NSF-funded grid resources	374	*19.9%*	*28.6%*	*24.5%*	*12.7%*	*15.1%*	
Resources outside the US	372	*84.9%*	*7.5%*	*2.7%*	*1.3%*	*3.8%*	
Other: (Please specify)	372	*94.6%*	*3.2%*	*0.5%*	*1.1%*	*0.8%*	

C.4. Satisfaction with XSEDE Resources

C.4.1. Satisfaction with Computation and Data Resources

Users report relatively high levels of satisfaction with all computational and data resources across all categories when evaluated individually (e.g. Ranger, Lonestar, Kraken, etc.) Similarly, users report relatively high levels of satisfaction when considering service dimensions (e.g., ease of use, reliability, etc.) across the aggregated XSEDE enterprise. Responses for each resource varied greatly and the number of respondents should be considered against actual usage when considering the overall score.

Computational Resources														
Resource	**Number of Respondents (703)**	Ease of use	Reliability (uptime)	Value of use of system to my research	Debugging tools	Performance	Scalability	Data transfer capability	Archival storage access and allottment	Available 3rd Party Software (applications and libraries)	Queue configuratio ns and operational policies	Documentation	User news	Average rating across all categories
Ranger	150	4.32	4.22	4.28	3.78	3.98	4.03	3.82	3.84	3.81	3.84	4.12	4.04	4.01
LoneStar	88	4.47	4.49	4.49	4.14	4.38	4.39	4.18	4.27	4.17	3.96	4.19	4.30	4.29
Kraken	146	4.22	4.23	4.31	3.83	4.15	4.20	4.00	3.94	3.95	3.77	4.03	4.12	4.06
Gordon	26	4.26	4.00	4.35	3.67	4.44	4.35	4.11	3.67	3.82	4.22	4.06	4.21	4.10
Future Grid	9	4.33	4.33	4.67	4.33	4.00	4.00	4.33	4.33	4.33	4.33	4.00	3.67	4.22
Blacklight	70	4.00	3.93	4.27	3.83	4.14	4.10	4.09	4.00	3.92	3.79	3.85	3.98	3.99
Open Science Grid	12	4.50	4.00	4.50	4.50	4.00	4.67	4.00	3.83	4.00	4.25	3.75	4.00	4.17
Nautilus	30	4.24	4.00	4.26	4.25	4.20	4.14	4.17	4.20	4.50	4.40	4.14	4.29	4.23
Forge	30	4.00	3.95	4.18	3.88	4.18	3.95	4.15	4.00	4.05	3.90	4.00	4.00	4.02
Trestles	54	4.20	4.00	4.27	3.94	4.11	4.15	4.00	3.75	3.85	3.94	3.87	4.07	4.01
Wispy	3	–	–	–	–	–	–	–	–	–	–	–	–	–
Steele	37	4.25	4.19	3.88	3.62	4.19	3.86	3.92	4.00	3.85	3.28	3.65	3.96	3.89
Purdue Condor	13	3.50	4.00	4.75	3.67	4.00	4.00	4.50	4.50	3.67	3.67	4.33	4.25	4.07
Longhorn	28	4.48	4.43	4.62	4.25	4.57	4.35	4.19	4.30	4.26	4.10	4.16	4.26	4.33
Spur	7	3.67	4.00	3.67	3.67	3.67	3.67	3.67	4.00	4.00	3.67	4.00	4.00	3.81
Average rating across all resources		4.17	4.13	4.32	3.95	4.14	4.13	4.08	4.05	4.01	3.94	4.01	4.08	

In addition to evaluating current resources, users were asked to comment on what other resources XSEDE could provide to enhance their work. The table below lists comments related to additional computing, storage, and/or visualization resources or support needed from XSEDE

What computing, storage, and/or visualization resources or support could XSEDE provide to help you with your research, teaching, education, and/or outreach activities?

Mount Lustre on both Kraken and Nautilus

Reliable tape storage (Several TB of data were "lost" twice in the last year on NICS HPSS); Increased resources for doing small-number-processor visualization/analysis without the long queue wait times of NICS Kraken; Better monitoring of login nods for abuse on NICS Kraken, to kick users who initiate 10 scp or file-system-intensive programs on a single login node; Better performance from the file system on NICS Kraken (lustre is very slow).

"Ranger+Ranch" fits my research.

Having at least a modest number of cores available for long periods of time to do some cutting edge research; Having a way to get high school and undergrad students involved at Univ of Washington in supercomputing applications; Better technical support (The staff are excellent but over stretched.)

A central archival system and a wide area network accessible data capacitor

A GPU cluster as efficient as Trestles, for example

A primer on "what do you want to do"? with help picking the right resources to request to meet your needs.

Additional computing time. My last grant was significantly cut due to oversubscription of the computing resources.

Again, I'll know better in about 3 weeks

All is well.

An increase in storage space would be nice.

An installation of QIIME would help considerably.

As I mentioned before, it has to provide longer wall time so we would not have keep restarting our job.

At NICS it would be very useful to have a shared file system between Kraken and Nautilus. When I want to check if a job is running OK, I need to run some analysis software on output. Currently I have to transfer data either to Nautilus or a local resource.

Better documentation for configuring thread/MPI software for optimum performance
Campus Champion program is very good. Would be good if XSEDE could provide long term storage services to help people comply with the NSF data management plans required for NSF grant applications.

Computer time

Currently available resources/support meet all my needs.

Data archiving is beginning to be an issue; I haven't used such resources on XSEDE yet.

Data-compression tools, data viz help

Ensight CFD

File transfer from or to other centers needs major improvement on Kraken. Currently it is too slow or not even working for some file transfer commands.
Frequent purging and small disk space hinder our progress. I/O remains a major issue on Kraken/Nautilus. Longhorn is much better and additional GPUs/nodes on Longhorn with additional disk space would be very helpful.

Gaussian.

GPU clusters and larger leadership class machines such as kraken

Greater integration of resources such that I am not "locked in" to NICS, TACC, NCSA etc.... I should be able to use whichever resource works best for the specific requirements of my research with little to no effort (e.g. compute at one place, viz at another, and store at yet another)

Having more nodes available, and ensuring that a large volume of smaller jobs (1 node, 5-10 hours) are not penalized in the queue scheduling relative to jobs which use several nodes for days or weeks at a time. Also, having a disk quota that's twice as large would speed up my research, by allowing me to submit more jobs at a single time before having to move the data off Steele.

I am happy with what I have currently.

I am interested in storing curated data products behind the XSEDE gateway - we need to have a discussion about such dedicated data resources, to the extent that third parties may purchase hardware for installation and operation by XSEDE personnel

I am satisfied with the available resources.

I believe that the resources and support are there, the problem is in gathering the data and then organizing it to optimize using it for visualization and simulation. We are just getting to that point in the project.

I do not require any visualization resources, and I find the computing resources sufficient.

It is not clear to me how one would utilize archival storage resources at sites other than those where computing allocations have been awarded, which is why I amusing other DOE-funded tape resources.

I do not require visualization resources and find the computing resources sufficient.

I am not as familiar with storage resources, but could benefit from more information about them.

I would like to see a maintained version of CHARMM on XSEDE resources (not CHARM++). Though I know it does not parallelize well, it is invaluable in my simulations.

I would like to use more visualization capabilities.

I'm satisfied with the resources that you offer currently.

Ideally, a single file system accessible by all XSEDE facilities would be great.

Improved scalability and operational time of Kraken. Machine is down too often.

Increase compute capacity to increase allocations. Allow graduate students to apply for individual allocation.
Increased number of sites with large memory
It has most of what is needed.

It would be great to have more storage capacity on Kraken and more information about visualization resources on Nautilus.

It would be nice to build and use our own specialized interactive visualization software at the XSEDE site rather than transferring all the output back to look at it. This would probably require some sort of higher speed interconnectivity to us.

LAMMPS, VASP

Large memory and high speed of remote visualization concerning CFD application

Large memory machines

Longer retention time for data

LS-DYNA, PARADYN, FEBio

Maybe more computational resources

More and faster processors

More computing power of the Ranger / Kraken / Lonestar class. Long-term archival storage with a decadal longevity.

More documentation on ease of use with examples for the students

More documentation on using the machines optimally for large jobs (I think the documentation provided by NERSC is somewhat easier to follow and more comprehensive than that provided for Ranger or, especially, Kraken.) Also, the ability to use debugging tools such as DDT with a larger number of processors would be very helpful.

More efficient job queue system

More widely available Gaussian 09

NCL, NCO, ncdump, ncview for data processing and visualization; more clear storage documentation; more clear modules documentation (and how to use the modules to compile another program/model)

Need to be able to keep large amounts of data and move them around with ease.

Need to manually "stage" files can be slow and occasionally difficult.Needs primarily met

None. Honestly, I am pretty happy with the systems. Except ranger, which is not really a great machine.

Overview of bioinformatics on Blacklight

Philip Blood!

Please give a guide to using the far archiver on PSC Blacklight for people who are UNIX novices. I know it exists but I don't understand how to use it.

Porting code to GPUs

Probably more mid-term range storage.

Provide a faster higher capacity Blacklight

Provide cost effective high memory capacity data analysis platforms. Speed is secondary.

Providing methods for showcasing active research, web-based mechanisms for sharing generated data

Provision of more storage space.

PSC Blacklight Archive

R

Ranch is frustrating, especially the process of staging and getting stuff back. Why should there be two different commands, one for staging and one for copying over? Why would I stage something and then not want to copy it somewhere, such as scratch? I should always automatically copy it somewhere once it's staged, instead of just saying nothing unless it is checked on, and then often unstaging by the time I get to it. And why can't I check on a queue or something to see roughly how many stage jobs are ahead of mine?

Ranch does have huge capacity which is nice, but it is very annoying to use.

Ranch on the Lonestar cluster

Ranger, Lonestar, Trestles

Resources that support third party software are important. storage is also very important. Support for gateway management and workflow execution are significantly needed.

Shorter queues

Show more demo

Smaller GPU machines for post processing.

Some computing nodes with memory more than typical will be helpful for memory-intensive simulations.

Stable long-term storage

Storage and high speed data transfer are a problem; for some computing problems high speed processors are more important than parallel

Strategic investments in systems that are robust and well run at massive scale.

SWIG

Systems with more temporary local disk storage on each node.

The current computing, storage and visualization resources are excellent. One thing that concerns our research is that we prefer a queuing policy that supports long time calculations, for example a wall time of 72 hours or 168 hours.

The Kraken and its HPCC storage is very helpful to my research

The main bottleneck is data transfer from machine (local or state) to or from NICS HPSS. This is, probably, not an XSEDE directive, though; with TB's of data needing to be moved around, days of uninterrupted network are required...

For example, at the en of a project or sub-project, we move all our data back home and store it to personal disks. Thus, if a program were made available where I could mail y'all my personal storage disks and y'all could USB3.0 my disk straight to HPSS, the problem of using networks to transfer data back and forth would be gone. As I'm probably not the only person doing such a thing, it could be hypothesized that a hell of a lot of bandwidth could be saved with such a program. Probably would be hella cheap to implement, compared to re-working infrastructures: buy 24 USB3.0 hard disk toasters and add >> 24 USB3.0 drives (somehow) onto HPSS, and hire 2 or 3 part-timers to pick up the drives from a mailroom, plug them into the toasters, execute an rsync script, eat burritos, unmount the disks, and mail the drives back to the owner....HK2012

There are no additional resources for computing, storage or visualization that XSEDE could give that would help with my work.	Would like to explore use of GPU's more. Would be great if XSEDE could more of these available
They already provide extensive training programs that have been very helpful.	XSEDE is already offering more than enough for my needs
Up to date libraries as well as open mpi stack.	You could provide a much greater variety of software for electronic structure calculations, both molecular and extended system. Data storage is limited. It would be great to have a computer system dedicated to "structure searches", with several methods
Training on how to utilize the debugging utilities.	
Visualization for analyzing large data sets with Visit	
Visualization of multi-dimensional plots.	
We are heavy users in Kenneland. A large GPU-based resources is sorely needed.	

C.4.2. Satisfaction with XSEDE Software

As part of evaluating current resources, users were asked to comment on what other resources XSEDE could provide to enhance their work. The following table lists comments related to additional software tools or libraries that would be useful from XSEDE.

Are there software tools or libraries needed for your work that are not available on XSEDE?	
image display program than imagemagick on Kraken	I sometimes must locally compile the GMP library, but that is not difficult.
A reliable job submission mechanism	I use Gaussian 09 mainly, which is installed on Carver
Amber Force Field	I would prefer to be able to instantiate my own images as needed. I think the core services of providing a robust infrastructure should be separated from custom user support.
Amsterdam Density Functional software	IDL on Kraken.
As far as I can recall, Steele has had all the libraries and tools necessary for the programs I've wished to run on Steele.	Libraries that if installed are typically many versions out-dated: HYPRE, SuperLU;
Debugging is a huge issue; we need ddt	Libraries that are usually never installed: SUNDIALS
Docking software, hadoop on demand	Ltrace
Gaussian 09	Matlab
GAUSSIAN, newest NWCHEM	Matlab and Tecplot
GIS. LiDAR processing tools	Matlab, Opencv
Having a parallel Matlab configuration would be very good.	MDAnalysis would be good. Then I could also do some analysis on the supercomputers
HEALPix	Moose Library for Perl
up to date gsl libraries	MUMPS
up to date boost libraries	My initial intent was to run simulations on Ranger with a new release of the Desmond molecular dynamics code, but the implementation on Ranger was very slow, according to my benchmarks. The situation was never resolved, so I retasked the resource for a project that can make effective use of NAMD. But --- an optimized
I built my own parallel IO capable netcdf library.	
I cannot be exactly sure what is on each system right now, but in the next year I will need python 2.7 or greater, pylons web service framework, and zoltan parallel library	

installation of Desmond would be a HUGE boost to my program.	SWIG
My research mostly involves NAMD.	CHARMM is software tool I would like to see on XSEDE resources.
NCL, ncview, ncdump don't seem to be available. NetCDF module is hard to work with for compiling WRF model.	Up to date versions of Matlab, Blender, visualization libraries, GPU libraries
Numpy, Scipy	VAMPIRtrace, VAMPIRserver, VAMPIR
ORCA	VASP module on Kraken
Parallel hdf5 with Fortran support on Gordon	Velvet; GATK; BWA; CLC
Perfection achieved, presently	Visit, Flash Code
Perl might be a good addition	VIsit is a visualization tools developed by LLNL. The tool is capable for analyzing large data sets in parallel.
Provide an additional site (e.g., Lonestar) to provide the program, Molpro, in additional to Blacklight.	VMD, MDAnalysis
Gaussview could also be useful.	We prefer Octave to be installed on XSEDE clusters.
QIIME and HyPhy	Yes, but that's because the software I use is fairly specific to my research institution.
R	Yes, but they are commercial software that may be hard to license for supercomputing centers. For example, Molpro, Schrodinger Jaguar, ... etc
Schroedinger Suite, Docking programs. However, I realize that these are exceedingly expensive	Yes, but they're easily installed.
Possibly Terachem	CAPS compilers for accelerators would be useful.
Some of them were not available but after request they were made available.	HDF5
Some proprietary software, such as Intel MKL and compilers, PGI compilers (not sure)	OpenFOAM
Some R parallel computing library and help information	I need constrained optimization.
Some, but most of the software is available	Yes. but I have been able to install my own

C.4.3. Computational Resources - Science Gateways

XSEDE Science Gateways are portals to computational and data services and resources across a wide range of science domains for researchers, engineers, educators, and students. Depending on the needs of the communities, a gateway may provide any of the following features: high-performance computation resources; workflow tools; general or domain-specific analytic and visualization software; collaborative interfaces; job submission tools; and/or education modules.

Only 13% of respondents reported using Science Gateways to access XSEDE resources. Gateways cited include: Biodrugscore, the CIPRES Portal, GridChem, SIDGrid, Linked Environments for Atmospheric Discovery, Globus Online, the OGCE Science Gateway Portal, Open Science Grid, QuakeSim, TeraGrid Geographic Information Science Gateway, and the Earth System Grid, among several others. None were cited with a frequency greater than 10 (Globus), and most were cited with a frequency of one or two. While usage was relatively low as a percentage of respondents, overall satisfaction among users is well above average.

Do you access XSEDE resources through Science Gateways?	Frequency	Percent
Yes	46	13.0%
No	308	87.0%
Total	354	

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your overall satisfaction with Science Gateways:								
	Mean	Number of Responses	Distribution					Histogram
			1	2	3	4	5	
Ease of use	4.26	35	0.0%	5.7%	8.6%	40.0%	45.7%	
Reliability	4.06	35	2.9%	8.6%	8.6%	40.0%	40.0%	
Input/Output Performance	4.13	32	3.1%	6.3%	12.5%	31.3%	46.9%	
Value of use of system to my research	4.27	33	0.0%	3.0%	18.2%	27.3%	51.5%	
Data transfer capability	4.16	32	3.1%	6.3%	12.5%	28.1%	50.0%	
Storage access and allotment	4.04	28	7.1%	3.6%	14.3%	28.6%	46.4%	
Operational policies	4.30	30	0.0%	3.3%	16.7%	26.7%	53.3%	
Job submission	4.39	31	0.0%	3.2%	9.7%	32.3%	54.8%	
Work flow management	4.17	30	0.0%	6.7%	16.7%	30.0%	46.7%	

Respondents who do not use Science Gateways in their works largely cite the lack of an appropriate gateway to fit their needs and/or lack of knowledge/awareness of their benefits and capabilities, rather than negative past experiences.

Please tell us why you are not currently using Science Gateways: (Check all that apply)*	Frequency
I have not found a gateway that meets my needs	108
Lack of knowledge/awareness of Science Gateways	50
Gateways are too complicated to use	40
No need current need for Science Gateways (e.g., current needs are met with other resources)	26
Haven't taken the time to learn about capabilities/benefits of Science Gateways	17
Gateways do not provide enough flexibility	16
Prefer direct command line access	16
Support level was not meeting my expectations	8
Gateways are not reliable	4
Other	27
*Total # of selections / # of potential respondents	312 / 308

C.4.4. Future Computational, Data, and Gateway Utilization Plans

In asking users to consider their planned resource usage in Year 2, responses aligned closely with, or slightly above, the current year's usage patterns as indicated by rate of response to survey questions related to resource use .

Which XSEDE services do you plan to use in the upcoming year? (Check all that apply)	Frequency
Kraken	141
Ranger	130
Lonestar	102
Blacklight	62
Tresles	55
TACC Ranch	46
NICS HPSS	40
Gordon	39
Science Gateways	29
Nautilus	26
Keeneland	24
Steele	24
Longhorn	23
NCSA MSS	19
Open Science Grid	16
PSC Data Archive	14
Future Grid	12
Purdue Condor	10
IU Data Capacitor	8
Albedo (and the Data Replication Service)	7
Spur	3
Wispy	2
Other (e.g., Forge, Stampede)	9
None	26
Undecided	5
*Total # of selections / # of potential respondents	873 / 443

C.4.5. Accelerators and Co-Processors

Users were asked to indicate which accelerators and co-processors, and associated programming models, they currently use, as well as which they plan to use in the next year.

Which of the following accelerators or co-processors do you currently use? (Please select all that apply)*	Frequency
GPGPUs	83
FPGAs	15
Other: *Cell BE and MIC each cited one time' others did not a cite a specific accelerator or co-processor*	11
None	5
Undecided/Unaware	3
*Total # of selections / # of potential respondents	117 / 443
Which of the following accelerators or co-processors do you plan to use in the next year? (Please select all that apply)*	**Frequency**
GPGPUs	115
Intel MIC co-processor	54
FPGAs	19
Other: none specifically cited	4
None	5
Undecided/Unaware	4
*Total # of selections / # of potential respondents	201 / 443
Which programming models do you currently use on accelerators or co-processors? (Please select all that apply)*	**Frequency**
CUDA	83
OpenCL	12
OpenACC	7
Other: Cuda Fortran, MPI each cited one time	11
None	6
Undecided/Unaware	11
*Total # of selections / # of potential respondents	130 / 443
Which programming models do you plan to use on accelerators or co-processors in the next year? (Please select all that apply)*	**Frequency**
CUDA	67
OpenCL	19
OpenACC	13
Other: *PGI Fortran cited one time; others did not cite a specific programming model*	11
OpenMP offload execution on Intel MIC	22
Native execution directly on Intel MIC	10
None	5
Undecided/Unaware	6
*Total # of selections / # of potential respondents	153 / 443

C.5. Satisfaction with XSEDE Support Services

C.5.1. User Training and Development

Users report a tendency to prefer independent, ad-hoc support resources (e.g., online tutorials and documentation, consultation with colleagues, etc.), rather than more formal methods (e.g., conference, workshops, etc.), in learning to use XSEDE and/or its predecessor, TeraGrid. This maps closely to how users indicate they prefer to receive training.

On a scale from 1 to 5, how useful are the following resources in helping you learn to use XSEDE and/or its predecessor, TeraGrid?								
	Mean	Number of Responses	Distribution (1 = Not use at all, 5 = Very useful)					Histogram
			1	2	3	4	5	
Independent, self-study (online documentation, man pages, trial-and-error, etc.)	4.42	326	2.1%	1.5%	6.7%	31.3%	58.3%	
Colleagues	4.04	294	3.7%	3.7%	18.7%	32.3%	41.5%	
Conferences	3.25	201	13.4%	14.4%	26.4%	24.9%	20.9%	
Attending onsite workshops	3.32	168	16.1%	11.9%	21.4%	25.0%	25.6%	
Synchronous online workshops	3.44	176	13.6%	10.8%	21.6%	26.1%	27.8%	
Online, self-paced tutorials	4.03	222	4.5%	5.0%	14.4%	35.1%	41.0%	
Other: (e.g., students, technical support, SP Forum)	3.90	29	10.3%	0.0%	24.1%	20.7%	44.8%	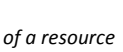

Usefulness numbers were adjusted to reflect only those who evaluated a given resource; those who indicated that their evaluation of a resource was "not applicable" were extracted from the total number for the purpose of more accurately quantifying usefulness levels.

How do you prefer to receive training? (Select all that apply)*	Frequency
Written documentation	247
Self-paced, online training (with hands-on components)	207
Live—online	122
Live—in person	120
Self-paced, online training (without hands-on components)	79
Other (e.g, on-the-fly, examples to edit, Google)	5
*Total # of selections / # of potential respondents	780 / 443

Having been asked to comment on the general training delivery methods, users were then asked to indicate what specific types of training (from a content perspective) would be of most interest in the coming year:

What specific types of training would be of most interest to you? (Select all that apply)*	Frequency
Tuning and Optimization	170
Debugging	124
Managing I/O	126
Visualization tools	132
Training on specific applications or libraries: (Please specify)*	39
Data analysis and management tools	115
Introductory programming topics (e.g., Fortran, C, C++)	56
Introduction to UNIX	29
Writing a successful XSEDE allocation request	97
Introduction to XSEDE	64
Programming Accelerators (GPGPU Programming, MIC Programmin	138
*Total # of selections / # of potential respondents	1090 / 443

***Unless noted otherwise, the specific applications or libraries were cited with a frequency of one: PETSc (3), R, VASP (2), Quantum Expresso, Trinity, Taverna, TAO, Airavata, MPI/Open MPI, LAPACK, MOLPRO, NWCHEM, BLAS, FFTW, LAMMPS, GROMACS, Intel Math Kernal Library, Flash Code, Gaussian 09*

Below is a complete list of user comments about training resources or support XSEDE could provide to aid research, teaching, education, and/or outreach activities.

What training resources or support could XSEDE provide to help with your research, teaching, education, and/or outreach activities?
A basic “Here’s where the important libraries are for each system,” rather than documentation scattered in many different locations, and in different patterns
A young, charismatic, real, and inspiring spokesman that makes their presence *very* public so that America knows what happens when taxes are spent well! (presumptuously, me in 5 years)
Adjusting environment variables as they relate to MPI. When is it appropriate to do so and what are the effects and side effects.
An "introduction to super computing" type tutorial would be helpful in training students in best practices.
An optimization online tutorial for myself and my students would be great! XSEDE "basics" tutorial for new student users would also be helpful.
Assistance migrating from machine to machine, since machines seem to only be supported for 1-2 years nowadays
Better documentation of obscure but necessary settings, e.g. for MPI and pthreads, on specific implementations.
Better responsiveness from support personnel (support was very good from SDSC; very poor on other systems).
Complete and up-to-date documentation is the most critical resource for me (and one where I have observed some lack).
Conversion from CPU to GPU architecture. "massive" parallelization schemes
Data visualization software for Gaussian03,09 outputs: Gaussview, Molden, Facio
General script formatting and usage
GPGPU programming.
Greater focus on High Performance Computing (supercomputers)
Hands on training for discipline specific topics would be useful.
High sufficient programming and software package example
How to use the PSC Data archiver
I am still not using all the possibilities I have
I believe that XSEDE should separate any training and support from a base level, focused organization that runs infrastructure in a cost effective manner. This is like the Federal Highway Department offering driving instruction. The highway department should maintain and develop

roads and bridges. Teaching people to drive is not their job. Same with XSEDE - it is trying to do too many things within one organization.

I cannot think about anything I really need at the moment, but am always curious about optimization/debugging and emerging technologies (like gpgpu).

I have not looked at the current resources and support closely enough to know what additional options would help me.

I think the resources currently provided are sufficient for my work.

I want to learn how to use TAU.

In addition, I would like to learn how to use cudaprof.

I will be starting my own lab soon, and I believe a very basic introduction (written or online tutorial) to high performance computing and using XSEDE resources would be invaluable to my future grad students.

I'm good with what is currently offered.

I'm not familiar enough with the offerings to be knowledgeable about this at this time.

I'm very happy with XSEDE's training resources and support.

introduction to different clusters

It is fine currently

It will be very helpful if I can know more about libraries(like mkl) and the routines available in these.

Keep online documentation updated and current

Knowledgeable Campus Champions.

More "intermediate" level classes. I know probably an intermediate amount in C++ and would like to improve, but the classes are usually too low level for me. I'd like to improve from the level I'm already at. The same goes for Unix.

More information on job scheduling/parallelization.

Most of it is taken care of, but newer developments in GPU computing are of interest to me

My students need simple, compact guides. A set of "How To" docs for common tools (gprof, tau, simple batch script examples, perf analysis tools and plots, etc.)

Online resources for self-paced study.

Online tutorial and/or workshop and phone services

Online, self-paced tutorials on the topics listed on the previous page of the survey would be great! Please provide tutorials in Fortran as well as in C or other languages.

Optimized parallel performance

Parallel scalability of a CFD code for large cases.

Queue optimization

Support in outreach and support establishing collaborations with researchers with complimentary expertise

Test programs for optimization and GPU/CPU porting

Transitioning between architectures (MPI to shared memory, for example)

Tutorials for classes of college students being exposed to a supercomputer

Understanding Unix.

UNIX for idiots

Use of new visualization tools

Utilization of Hybrid CPU/GPU systems

We had trouble getting our code to queue and run on Steele. A debugging queue would be invaluable. We also never figured out who to ask for help.

Xsede should support the statistical sciences and thus social sciences with support for R.

C.5.2. User Support Staff

When asked if there was a local resource person available to assist with their use of XSEDE, users' responses varied, with a significant portion indicating they did not have – or were unaware of – local support resources.

Is there a resource person at your institution available to assist with your use of XSEDE? (Check all that apply)*	Frequency
Colleague at my institution	118
Local IT support person	64
XSEDE Campus Champion	50
XSEDE staff member	35
No	108
I do not know	80
*Total # of selections / # of potential respondents	455 / 443

XSEDE user support staff receives high marks, with users reporting they are very or completely satisfied with most aspects of their experience well over 80% of the time:

On a 1-5 scale, with a 1 being "Very dissatisfied" and 5 being "Very satisfied," respondents rated thier experience with XSEDE user support staff:								
How satisfied are you with the…	Mean	Number of Responses	Distribution (1 = not at all satisfied, 5 = very satisfied)					Histogram
			1	2	3	4	5	
...effectiveness of problem resolution provided by XSEDE user support staff?	4.31	298	3.0%	4.0%	9.7%	25.8%	57.4%	
...helpfulness of XSEDE user support staff?	4.43	297	2.0%	3.4%	6.7%	24.9%	63.0%	
...knowledge of XSEDE user support staff?	4.40	291	1.7%	3.1%	8.6%	26.8%	59.8%	
...communication skills of XSEDE user support staff?	4.41	297	1.3%	2.7%	7.1%	31.0%	57.9%	
...timeliness of responses from XSEDE user support staff?	4.36	296	2.4%	2.4%	10.1%	26.7%	58.4%	
...notifications and announcements provided by XSEDE user support staff?	4.38	287	1.7%	1.4%	10.1%	30.7%	56.1%	
...assistance provided by your XSEDE Campus Champion, local XSEDE staff member, or other local support providers?	4.40	167	2.4%	1.8%	10.2%	24.6%	61.1%	
...courtesy and demeanor exhibited by XSEDE support staff	4.57	278	1.4%	1.8%	5.4%	20.9%	70.5%	

Satisfaction numbers were adjusted to reflect only those who evaluated a given service; those who indicated that their evaluation of a service was "not applicable" were extracted from the total number for the purpose of more accurately quantifying satisfaction levels.

C.5.3. XSEDE User Portal

Nearly 71% of XSEDE users report using the XSEDE User Portal, with high levels of satisfaction across all service dimensions (e.g., allocation management, file management, etc.).

Do you use the XSEDE User Portal (XUP)?	Frequency	Percent
Yes	244	70.9%
No	100	29.1%
Total	344	

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with the following XSEDE User Portal (XUP) features.

	Mean	Number of Responses	Distribution (1 = not at all satisfied, 5 = very satisfied)					Histogram
			1	2	3	4	5	
Allocation management	4.35	226	1.8%	1.8%	9.3%	33.6%	53.5%	
Logging into systems	4.04	194	2.6%	9.3%	16.0%	25.8%	46.4%	
File management	4.02	162	1.2%	10.5%	15.4%	30.9%	42.0%	
User forums	4.05	129	2.3%	4.7%	20.9%	29.5%	42.6%	
Resource Monitor	4.24	216	0.5%	3.2%	16.7%	31.0%	48.6%	
My Jobs	4.14	191	2.1%	4.7%	16.2%	30.9%	46.1%	
Queues prediction	3.81	167	4.2%	10.8%	22.8%	24.0%	38.3%	
User news	4.15	178	0.6%	4.5%	18.0%	33.7%	43.3%	

Satisfaction numbers were adjusted to reflect only those who evaluated a given service; those who indicated that their evaluation of a service was "not applicable" were extracted from the total number for the purpose of more accurately quantifying satisfaction levels.

Fewer than 4% of users reported having used the new XSEDE Mobile, with 37% indicating they were unaware of this service.

Do you use XUP Mobile?	Frequency	Percent
Yes	9	3.7%
No	145	59.2%
I am unaware of this service	91	37.1%
Total	245	

C.5.4. User Support Documentation

Over 71% of XSEDE users report using the XSEDE User Support Documentation, with 94% of those who responded indicating that the documentation resources were helpful in their work. User satisfaction in this area is well above average across all services and resources.

Do you use XSEDE User Support Documentation?	Frequency	Percent
Yes	242	71.8%
No	95	28.2%
Total	337	

Do you find XSEDE User Support Documentation helpful?	Frequency	Percent
Yes	233	94.0%
No	15	6.0%
Total	248	

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with the following activities and services.								
	Mean	Number of Responses	Distribution (1 = not at all satisfied, 5 = very satisfied)					Histogram
			1	2	3	4	5	
XSEDE.org web site	4.21	219	1.4%	2.7%	13.7%	37.4%	44.7%	
XSEDE User Portal (XUP)	4.27	204	2.0%	2.5%	9.8%	37.7%	48.0%	
XSEDE help desk services via help@xsede.org	4.34	201	2.5%	2.5%	8.5%	31.3%	55.2%	
XSEDE help desk services via telephone (866-907-2383)	4.41	236	1.3%	4.2%	11.0%	19.5%	64.0%	
XSEDE help desk services via XUP	4.34	106	2.8%	0.0%	11.3%	32.1%	53.8%	
XSEDE Knowledge Base	4.11	161	1.2%	5.0%	15.5%	37.9%	40.4%	
XSEDE user news	4.15	191	2.1%	1.0%	18.3%	36.6%	41.9%	
XSEDE online user services documentation (user guides, usage policies, etc.)	4.21	225	1.3%	4.4%	10.2%	40.0%	44.0%	

The following table is a full listing of user comments about which new documentation and portal feature they would find most useful.

Text Comments: What new features would you like to see in the XSEDE User Portal or in the documentation?	
Better listing of software available in the different facilities	More and better-written documentation
Simpler live queue of all resources	More descriptive examples
I would like to see when and how much CPU time is refunded because of system problem; I would like to see who used how much CPU time in our group; I would like to see job scheduling policy.	More information about compiling hybrid code would be welcome.
A more user friendly, intuitive portal.	More specific examples of job submission scripts (for example) with the different parameters actually explained in a way accessible to scientists who happen to use supercomputers.
Better predictions when my jobs will start.	More up to date and specific information on library linking would be of help.
Clearer.	Pointers to vendor on-line documents
Either more information on setting up/customizing your environment so it works for me or more help from the help desk/user forum.	Quick-reference lists of the system-specific parameters: queue names/parameters, names of various temporary storage directories, etc. Whenever starting on a new system, I spend way too long looking for these things. And too often I've ended up having to ask HelpDesk staff for assistance because some piece is missing or hard to find.
Everything is great for now.	Simplify access; the ssh portal is almost unusable.
First, I would like to use everything that exists.	Site is slow to load
How to set up remote ssh instead of using the java-based terminal	Support for Statistics.
I would like to be able to cc (or otherwise include) other users in help desk tickets submitted through the XSEDE User Portal. At present, when I report an issue that also affects my colleagues, I need to forward every message to them separately.	The documentation is not well organized. Too many clicks to find out about the specs of the HPC machines.
I would like to see more emphasis on delivering more cycles in a more responsive manner without layers of administration. I would like to see domain experts empowered to develop portals as they see fit via more aggressive use of dynamic imaging and user customization.	Third party software performance data on various resources that is current would be very useful.
I would like to see more extensive documentation on job scheduling (i.e. all details with more examples of job scheduling scripts).	Um, I'd love to see it not be incredibly slow any time I log in.
It needs to work faster with browsers.	That's a royal pain in the butt - it occasionally takes ten or fifteen seconds between when I click on a link and when I see the new page.
Keep information up to date. No links to broken or dead projects	This is intermittent, but crippling in terms of usability.
Live system status monitor	User guide demo
	Website organization can be confusing - at times I end up going around in circles looking for things

C.5.5. Allocations

Over 56% of users have used the POPS Allocation Process. Satisfaction is consistently high, with each service dimension scoring well above average.

Do you use the POPS allocation submission process?	Frequency	Percent
Yes	192	56.3%
No	149	43.7%
Total	341	

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with the allocation process:								
	Mean	Number of Responses	Distribution (1 = not at all satisfied, 5 = very satisfied)					Histogram
			1	2	3	4	5	
Process of applying for an XSEDE allocation	4.17	183	*1.6%*	*4.9%*	*12.0%*	*37.2%*	*44.3%*	
Speed of response for my XSEDE allocation	4.09	183	*1.1%*	*7.1%*	*16.9%*	*31.7%*	*43.2%*	
Response to my request for an XSEDE allocation (amount of resource allocated)	4.16	181	*1.7%*	*5.0%*	*14.4%*	*33.7%*	*45.3%*	
Feedback on my XSEDE allocation request	4.13	166	*1.8%*	*3.0%*	*17.5%*	*35.5%*	*42.2%*	

Satisfaction numbers were adjusted to reflect only those who evaluated a given service; those who indicated that their evaluation of a service was "not applicable" were extracted from the total number for the purpose of more accurately quantifying satisfaction levels.

C.5.6. Extended Collaborative Support Services

Only 28% of users responding reported being aware of XSEDE's Extended Collaborative Support Services (ECSS). Of these users, 72% are aware of how to request ECSS, but only 23% have done so. While only 15 users responded to the question regarding overall satisfaction with their ECSS experience, 67% report being very or completely satisfied. The mean score for overall satisfaction was 3.93 on a scale of 1-5, indicating an above average level of satisfaction.

Are you aware of the XSEDE Extended Collaborative Support Services (ECSS)?	Frequency	Percent
Yes	95	28.0%
No	244	72.0%
Total	339	

Do you know how to request Extended Collaborative Support Services (ECSS)?	Frequency	Percent
Yes	69	73.4%
No	25	26.6%
Total	94	

Have you received assistance from Extended Collaborative Support Services (ECSS)?	Frequency	Percent
Yes	22	23.7%
No	71	76.3%
Total	93	

In what area(s) have your projects received Extended Collaborative Support Services (ECSS)? (Check all that apply)*	Frequency
Advanced Research Team Support	5
Novel and Innovative Projects	2
Advanced Community Capabilities Support	0
Advanced Science Gateways Support	5
Advanced EOT support	2
I am not certain in which area(s) my project(s) received ECSS	7
*Total # of selections / # of potential respondents	21 / 95

What assistance did Extended Collaborative Support Services (ECSS) provide? (Check all that apply)*	Frequency
Performance Optimization (single process or parallel)	12
Application Scalability	6
Parallelism (MPI, OpenMP, or Accelerators)	10
I/O Optimization	7
Application Gateways (including grid computing)	6
Other	2
*Total # of selections / # of potential respondents	43 / 95

C.5.7. Extended Collaborative Support Services – Future Needs

Do you have projects that could benefit from collaboration with Extended Collaborative Support Services (ECSS)	Frequency	Percent
Yes	47	13.9%
No	41	12.1%
I don't know	250	74.0%
Total	338	

Do you intend to request Extended Collaborative Support Services (ECSS) for your project?	Frequency	Percent
Yes	35	11.9%
No	52	17.6%
I don't know	208	70.5%
Total	295	

What assistance do you expect to request from Extended Collaborative Support Services (ECSS)? (Check all that apply)*	Frequency
Performance optimization (single process or parallel)	105
Application scalability	67
Parallelism (MPI, OpenMP, or accelerators)	101
I/O optimization	62
Application gateways (including grid computing)	32
Other: (Please specify)**	11
Unsure/undecided	4
I don't know what ECSS is	4
*Total # of selections / # of potential respondents	386 / 443

***Assistance respondents expected to request cited with a frequency of one: Flash Code, Future Grid, Visualization, WorkFlow, MIC, CUDA.*

Have you ever received Extended Collaborative Support Services (ECSS) from someone at one XSEDE site for a system located at another XSEDE site?	Frequency	Percent
Yes	9	3.8%
No	150	63.8%
I don't know	76	32.3%
Total	235	

The following table is a complete listing of comments related to the process for obtaining Allocations and Extended Collaborative Support Services.

Do you have any comments about the processes for obtaining Allocations (POPS) and/or Extended Collaborative Support Services (ECSS), or other general comments about Allocations and ECSS?
ECSS has been extremely helpful. The staff is extremely knowledgeable and has made it possible to use more of XSEDE's resources.
I am not clear to ECSS.
I am very happy with the POPS system; it was very effective and in my case had a good outcome.
I do not know on what basis decisions re ECSS are made, and what kind of help is amenable to such a request.
I do not want to build a gateway. I would like XSEDE to stand up and operate gateways for popular disciplines and applications
I find the allocations process is something of a farce. The reviewers come across as generally second rate and it is a huge waste of everyone's time.
I found applying for Allocations somewhat daunting.
I think POPS is very good, and I do not know anything about ECSS.
I've been reasonably happy with Allocations.
It seems like things are a bit slower lately, post-XSEDE transition, but I imagine that's due to changes behind the scenes that mere users are not privy to.
It's certainly not affecting my ability to do my job!
It is my understanding that it is not too difficult and that the directions for application are fairly straightforward.
Like to see an improved process for allocations for classroom activities involving students
POPS turnaround should be significantly quicker. It is tough to design research without knowing if computing time will be available for several months.
POPS web site is very counterintuitive; very hard to use
The only thing I have to comment on was that I never received any notification or feedback after making a startup request. I happened to log in once and noticed that I had hours, which is quite indirect.
The POPS system is outdated and very clunky. A streamlined process would be much appreciated.
The process is very cumbersome and the long turnaround time greatly reduces XSEDE's usefulness for responding to research needs.
The selection of the reviewers for the allocation process needs to be matched better to the particular topic of the allocation request proposal. Some reviewers seem to be quite unfamiliar with the proposal topic, which is evident from some of the comments they make. Such reviewers often make unjustified claims, which then lead to significantly reduced allocations.
The system is okay. The POPS interface is clunky but I am used to it.
The user interface has improved a lot over the years (I like being able to save state). However, I did find it hard to really tell where I was in the process (what pages / sections were completed). I recently asked for an extension, and in the process of reloading the existing data, my co-PI data did not reload. Would be good if you presented a summary page of current data, and let me select what data needs to be updated.
This survey is the first time I've ever heard of ECSS. I don't know anything about it. As for POPS, I know my advisor pinged XSEDE several times over a period of several weeks to inquire about the status of our application, and she received no response whatsoever from XSEDE during that time, until our allocation request was granted. Not getting any response to inquiries for weeks was frustrating.

C.6. Final User Comments

In concluding the survey, users were asked to provide any final comments or suggestions on XSEDE (or other NSF-funded cyberinfrastructure), as well as thoughts on the value derived from the NSF's investment in XSEDE. Responses ranged broadly, but many users commented on the absolutely essentially nature of XSEDE to their research program and were in full support of the NSF's current and future investments.

Please provide any comments you have on the value derived from the NSF's investment in XSEDE	
A great resource, especially for starting researchers building their own research programs. Much better resource than is available at my current institution.	I think that NSF needs to look hard at providing more investment in support staff for XSEDE resources.
Being at a small liberal arts college with limited computational resources, my students and I find the XSEDE facilities and support to be indispensable to our research.	I couldn't do my work without XSEDE.
Computation resources are of vital importance to research of both basic and applied sciences.	I derive a tremendous value from XSEDE; simply put, if I didn't have access to it, I couldn't do the work that I do now. My entire field of research wouldn't exist if it weren't for federally funded open access to supercomputers. So, I'd say that is an awful lot of value derived.
I am very content with the current NSF investment in XSEDE and I hope that NSF could keep investing in large computation resources like XSEDE.	I got very good value because of NSF investment in XSEDE.
Essential or very important to our research, but because of continuity issues, excessive hardware downtime, poor documentation, out-of-date libraries, we are unable to derive a great benefit.	I greatly appreciate the support by NSF because I cannot do my research effectively without XSEDE.
Essential!	I think the NSF's investment in XSEDE is a fantastic national resource. It helps US researchers immensely. I would be happy to support this in any way from the University of Washington.
Excellent value	I think that it will be a very beneficial service to the academic research community. There is great hope also that it will excite those who are not currently using the national HPC resources to become engaged in order to more efficiently and productively further their research.
Extremely valuable!!	I think that there is enormous scope creep in the XSEDE community and that it's focus on infrastructure development and core operations is sub-optimal. Environments like the Amazon AWS services are vastly more satisfying to work with than XSEDE; however, they are not cost-effective. I do no understand why XSEDE isn't trying harder to offer a service that is as satisfying as Amazon, but with world-breaking price-performance. It should not be trying to do science and engineering itself, but too often that's what it appears to be doing. I think focus on increased specialization around delivering basic, but extremely high-performance, compute-storage-network services will lead to far more science and engineering discovery. Higher-level services can be and should be developed and cultivated outside of the XSEDE core. I think huge opportunities are being missed due to a lack of focus.
Great value!	I think the investment in XSEDE by the NSF is valuable to thousands of researchers and educators.
Great job.	It is a very complex infrastructure that needs to have better, more standardized interfaces.
Great value.	
Having access to large amounts of computing time is essential for the progress in my field of research. In particular, I benefit from large centralized computing resources, since they allow to test out several ideas in parallel and thus speed up discovery substantially.	
Highly valuable, would be difficult to do my work without it	
I am grateful that NSF has funded this resource and that it exists.	
I literally could not do my dissertation project without the massive amount of computational resources that are available with XSEDE.	
I believe this is a strong and strongly needed investment.	
However, it seems clear that NSF's investment in support staff and services has declined and I fear this is eroding the value of the XSEDE hardware.	

I think that users need to have a better idea of the big picture of what the purposes and directions of the administrators of the system are.

In addition to the direct benefit of my XSEDE allocation to my ongoing research, there is great value in the efficiency of operating so many high-performance computing sites within such a unified system with a single allocation and review process.

In my view, this investment on XSEDE by NSF is worthy to continue.

Is XSEDE just a pool of money given to the various sites to do with at they wish, or is XSEDE its own thing?

It is transformative. My research would not be possible without it.

It allows researchers from any institution to pursue projects that they may otherwise not do due to limited computational resources.

It has been phenomenal value for money for the NSF. XSEDE provides a fully operational, reasonably flexible ecosystem to facilitate a diversity of grid-enabled computing. It is a unique resource and I am very happy to acknowledge XSEDE's interest in reaching out to the life sciences (where I work).

It is a great resource.

It is a tremendous resource. More money should be spent funding XSEDE and U.S. supercomputing infrastructure.

It is tremendously valuable to a large group of people and a worthy investment by NSF.

It is very helpful that I get to use some of the fastest and best machines in the world, which are very crucial to my research. In addition, the workshops are also very useful to me.

It is very useful. We distribute 10 million CPUs hours to a community. They are able to produce 300 publications that would not have otherwise been possible. It is great.

It supports scientific explorations in a way that local resources cannot fulfill.

It would be hard to do without it

It's early yet, but if XSEDE works as well as TERAGRID, it is safe to say that we simply couldn't do our research without facilities like these.

It's invaluable and a great idea to centralize this instead of having to fund computer resources at every Ma and Pa university in the country.

It’s wonderful to have the power of a supercomputer available for research.

Making American science sustainably competitive in an exponential world

More focus on high performance computing

NSF is doing great help for researchers by providing computational resources.

NSF is getting a good return on its investment in XSEDE

NSF should definitely continue investing in XSEDE. It's a critical resource for computational scientists at small colleges and universities that cannot afford to provide much local computing.

NSF's investment in XSEDE allows those who would not otherwise have access to superior computing facilities the ability to successfully pursue their scientific inquiries in a simple and powerful manner. The value is extremely high.

Our national computing resources are absolutely essential to maintaining our international competiveness. Continued investment in infrastructure and next-generation computing technologies is imperative. The world moves fast; we must keep up

Please keep it going, and enhance it. This is very helpful. Also, some subjects by nature have a different requirement of hardware, like CFD (large scale and small scale problems), so please allocate the resources or invest on the resources more efficiently.

Re-branding NSF-funded cyberinfrastructure to Teragrid was understandable. Doing again for XSEDE seems like an administrator's dream, but offers nothing of substance.

Really helpful for my research. The technical support is offered on time and is really awesome.

Sorry, but this survey is too long. I did not expect that. Having a way to save in the event that I don't have time is good. I groaned when I realized that I had to a full survey page for _each_ resource I was using. You need to find a way to tighten that up and put multiple machines on one page to help speed up the survey process.

Terrific

The helpdesk people should be better trained.

The investment is extremely valuable, and indispensable for the scientific research programs that have benefited from the XSEDE resources. Without the resources, a lot of the state-of-the-art research would not be possible. I hope NSF continues its invaluable investment in XSEDE.

The NSF and XSEDE/Teragrid have ignored statistics.

The NSF's investment in XSEDE is invaluable to the scientific community in that it allows a large number of research groups access to computational resources and expert support for these resources in a centralized manner.

This model is highly efficient; instead of paying for a number of individual computational resources and the support staff needed to run them from grant funding, it centralizes the costs for both aspects. It is necessary for scientific computing and allows individual groups access

to blocks of computational time on demand, greatly increasing the efficiency at which these calculation can be done and minimizing the time during which these resources are idle.

The overall investment is great. My only complaint is that I often feel that too many of the systems are optimized for too similar types of computational problems, and greater hardware diversity could help.

The quantity and quality of XSEDE resources that have been available to our project has been invaluable. Many thanks!

The value of NSF's investment is evidenced in the ability for me to perform large-scale simulations integral to my research. XSEDE provides a very useful resource for setting up and performing simulations.

The value of well-maintained high performance supercomputing changes the landscape of science. It is becoming more important, as interesting system sizes grow, to have access to supercomputing facilities. Strides in research have been accomplished due to NSF funding of XSEDE.

The XSEDE is extremely valuable for my research.

There is a large benefit, especially to those on campuses without support for computational science.

This is a great program. Increase funding for greater benefit.

This is a great resource for smaller projects at smaller institutions that don't have computing resources. Without this, I would be really stuck and isolated.

This is a very valuable resource for computational projects funded by the NSF. One wish would be that once an award is made by NSF that the process of then getting an allocation on XSEDE would only involve a short (one page) statement of the project with required computing needs. The need to write an additional proposal (almost of the length of a regular NSF proposal) to obtain a computing time allocation on a resource already supported by NSF basically doubles the effort for a PI.

This is an incomparable resource that has kept US computational science at the forefront in the world.

This is an incredibly useful tool that has directly benefited my research. I look forward to using XSEDE for my research and outreach needs in the future.

This is an outstanding resource and I am more than happy with the capabilities and level of support. I think it is money well spent that has enabled research projects that otherwise would not be possible.

Valuable resources ... probably more valuable at this point to users at institutions where resources are located and institutions with little to no other HPC resources.

Value over Teragrid not noticeable

Very helpful and necessary.

Very helpful and worthwhile.

Very powerful, an important resource for my research group.

Very useful for a wide variety or researchers

Very useful in scientific research

Very valuable to my research projects

We have benefitted greatly from XSEDE Services in our research

We make heavy use of NAMD on XSEDE resources (Ranger/Kraken). NAMD simulations are essential for our research and we look forward to using it on newer and faster machines through XSEDE. We also backup all our simulation data in NCSA/MSS, a fantastic and essential resource for us.

When not overloaded, the tools offered by XSEDE have been extremely vital in accomplishing my research. The speed and size of the computers have been phenomenally helpful. However, recent overloading of certain assets have made working with XSEDE computers frustrating. In the future, loading limits should be reexamined.

Without XSEDE my modest needs would not be met. It is a good way to support smaller (including PUI institutions) researchers.

Without XSEDE resources, my research would be severely hampered it terms of both scope and pace. XSEDE especially stands out to me for providing resources for physical science and engineering research that is computationally costly but perhaps does not push the boundaries of feasibility (i.e. jobs on hundreds to thousands of cores rather than tens to hundreds of thousands). For these purposes and user groups, reliability, ease of use, good support staff, and quick turn around time are essential features that XSEDE does a good job of meeting,

XSEDE has been invaluable to the community.

XSEDE is a critically important resource. Without it, my research that spans basic materials physics to engineering applications would not be possible.

XSEDE is a fantastic resource for high-performance computing, allowing us to answer some of the most pressing questions of our time. Our simulations that provide critical support to coastal managers and emergency personnel would not be possible without facilities like XSEDE.

XSEDE is a great national resource. Quite a lot of that resource is currently wasted on users who should run small local jobs at their own sites. XSEDE should provide the resources that cannot be found locally. This might mean better regional infrastructure.

XSEDE is a valuable and important resource, but for a single investigator its value is significantly reduced do to the time-sink associated with overcoming technical barriers (e.g., login problems, being constantly shuffled

from one system to another and then having to debug new problems with setup and optimization settings on each new system, etc.)

XSEDE is a valuable asset to the scientific community. Without the XSEDE resources I would be unable to perform my research as my institution and my collaborators do not have the appropriate computational resources.

XSEDE is an extremely valuable resource that should continue to receive funding to expand its operations.

XSEDE is an essential resource for computationally intensive research, like climate modeling.

XSEDE is an excellent resource and has been incredibly valuable to my research in advanced thermoelectric materials.

XSEDE is an incredible service, and I plan to rely on it more as my current software projects leave the development stage and enter production. Without it, I might not have considered designing my programs for scalability and would not be as prepared to take advantage of modern computational power as I currently am; even using it only minimally, it has improved my experience as a graduate student qualitatively.

XSEDE is an initiative that is currently working amazingly well! In my research career, there has not been another initiative/resource, which has had such tremendous impact on my ability to do my research. I am truly grateful for its existence and availability! XSEDE is enabling numerous researches nationwide to do cutting-edge research!

XSEDE is great. I have to say, it is much better than the TeraGrid program. Far more reliable.'

XSEDE is very valuable, but would be more so with more ground level technical assistance.

XSEDE is vital to my research. The type of problems I study would be unapproachable without XSEDE. If possible, funding needs to be increased.

XSEDE provides a lot of resources/services that I have never heard of, don't use, and probably never will use. So it's hard for me to assess their value.

XSEDE provides an invaluable service to those of us in the numerical simulations field; large-scale supercomputing essentially for free. A vast part of the scientific endeavor could not be accomplished in the US without it.

XSEDE provides the "production"-like landing place for FutureGrid test bed users who are ready to tap into a larger pool of nodes/cycles.

XSEDE provides the highest value possible for what NSF is investing. A cyber-infrastructure of this magnitude (hardware/software and expertise) is indispensable for promoting research and education.

XSEDE represents an extremely valuable resource that would be unreasonable for many institutions to support themselves.

XSEDE resources have enabled my research. It would not have been possible without XSEDE

Do you have any other suggestions or comments regarding XSEDE and/or other NSF-funded cyberinfrastructure?

A mechanism for disseminating knowledge about new resources and their capabilities would help to propel my research program. As it is, I typically learn of such things from colleagues.

Adding new users still too slow

Budget needs to be allocated to local Campus Champions/Campus Bridging people and for the collaboration services participants beyond XSEDE personnel.

Continue the efforts...

Cyberinfrastructure demands are growing but NSF funding for research is not. We are going to be left behind in worldwide terms soon.

Decouple the soft aspects from the core infrastructure aspects and focus on building more robust and reliable facilities for compute, network and storage; leave the higher-level stuff to discipline areas.

Don't ignore statistics and the social sciences.

Greater integration and ease of usability across systems

I am looking forward to seeing more documentation and training resources becoming available. I would also really appreciate any additional queue management features that can be provided to help me decide where to submit jobs, and chose core counts and runtimes judiciously.

I have one specific piece of advice to offer: In our research, we often need to use cores for a long amount of time (e.g., we have a 32 core machine on which we run jobs for 30 days). We have never been able to run these jobs on XSEDE given the huge number of resources that XSEDE has. The primary reasons are that (i) the jobs abort after two days (the allotted time); (ii) the jobs do not parallelize well on large number (e.g., 1000) of cores.

I would like to be able to easily monitor job submissions on my award from approved users other than me.

I'd just like to compliment the XSEDE support staff. Every interaction I've had with them has been phenomenal; they're very helpful, prompt, and courteous.

It's fantastic!

Just quicken the turnaround time for POPS proposals. Having to wait for months to hear back is not acceptable, especially on smaller proposals like mine.

Maybe better organization of documentation, fewer clicks, not long, but short pages to browse.

More funding for long-term storage

NSF should really commit to production infrastructure on a long time scale, or give it up. It takes a long time to win user confidence, so 3- or 5-year funding cycles do not work for infrastructure. It just ends up making users hate you.

Provide opportunities to showcase science gateways to potential new communities for widening user base could be considered as part of training/out reach activities.

Smaller universities (like mine) do not have the resources to provide local HPC facilities. This use of NSF-funding has a high impact and high reward, and is utilized by researchers at more rural and underrepresented schools. This benefits those schools and_ the students at those schools.

So far, my experience with the staff is very satisfactory

Sometimes the wait time in the queue (e.g., in ranger) is more nearly 50 minutes. It would be better if that could be reduced.

Stop spending resources on rebranding.

Thank you very much for your dedication in maintaining this important national infrastructure!

The NIH should leave biomedically related cyberinfrastructure development and access to the experts, specifically via XSEDE.

The web/java-based user login is very inconvenient. Copy+paste and auto-complete are a pain compared to a proper Xterm. Often I have to log off and log in again to be able to see what I'm doing.

These infrastructures are absolutely essential for research progress in almost every field.

This is the way of the future. However, I think an accounting model like Amazon's EC2 might make more sense.

This survey is way too long.

We need to complete the rules of engagement between XSEDE and FutureGrid, so we can understand how best to utilize both of our resource pools effectively.

XSEDE is a very valuable resource for the scientific and research community in US. It is great to be able to provide access to this resource in a very open and easy way to researchers who are interested in using it.

XSEDE is much easer to start using than other clusters -- especially those at my home university

D. Appendices

Appendix Contents

- Appendix 1 – Recruitment and Reminder Letters
- Appendix 2 – Survey Interface (screenshots)
- Appendix 3 – Study Information Sheet (screenshots)
- Appendix 4 – Question Interfaces (representative screenshots)
- Appendix 5 – Complete Survey Text

D.1. Appendix 1 – Recruitment and Reminder Letters

D.1.1. Appendix 1a – Initial Recruitment Letter

From: Stewart, Craig Alan
Sent: Wednesday, May 02, 2012 3:56 PM
To: [email address]
Subject: 2012 XSEDE User Satisfaction Survey

Dear [first name] [last name]:

XSEDE (Extreme Science and Engineering Discovery Environment) is the most advanced, powerful, and robust collection of integrated advanced digital resources and services in the world -- a single virtual system used by researchers, technologists, and scientists, such as yourself, to interactively share computing resources, data, and expertise.

Your feedback is vital to the evolution of this important resource, and I am writing to ask for your participation in the 2012 XSEDE User Satisfaction Survey conducted by Indiana University.

The annual survey aims to assess users' current levels of satisfaction with the XSEDE computational environment and its associated services and activities (e.g., training, allocations, conferences, user support, etc.) Further, this survey aims to quantify the satisfaction with, and availability of, cyberinfrastructure resources and support among members of US scientific and engineering communities, regardless of their use of, or access to, XSEDE resources. Survey information will be used to improve and expand services to the XSEDE user community and to the US scientific and research communities, broadly, and to aid in the decision-making processes related to resource allocation.

The survey can be accessed here:

https://websurv.indiana.edu/xsede/login.cfm?id=5DE754AA3E

The Indiana University Center for Survey Research (CSR) administers the survey and assures that your responses will remain completely confidential. Neither your name nor your organization will be associated with any data or included in any reports.

If you have any questions about this survey or how the results will be used, please feel free to contact Julie Wernert, Information Manager, Indiana University, at jwernert@iu.edu, or call (812) 856-5517.

Thank you for your time and help with this important effort that will impact future decisions related to cyberinfrastructure funding for the US scientific and engineering communities.

Sincerely,

Craig A. Stewart, Ph.D.
Executive Director, Pervasive Technology Institute
Associate Dean, Research Technologies
Office of the Vice President for Information Technology
Indiana University

The IU Center for Survey Research is administering this questionnaire on behalf of the National Science Foundation-funded Extreme Science and Engineering Discovery Environment (XSEDE).

If you are unable to access the link listed above, please follow these instructions:

- In your web browser, type: https://websurv.indiana.edu/xsede
- In the Login Code box, enter: 5DE754AA3E

If you have any other difficulties logging in or have questions about the study, please e-mail csr@indiana.edu for assistance.

If you do not wish to participate or receive further notices about this study, please use the instructions above to access the survey site. After logging in, select the button marked "I do not wish to participate."

Reference ID: 4788323434

D.1.2. Appendix 1b – First Reminder Message

From: Stewart, Craig Alan
Sent: Friday, May 11, 2012 3:19 PM
To: [email address]
Subject: REMINDER: 2012 XSEDE User Satisfaction Survey

Dear [first name] [last name]:

XSEDE (Extreme Science and Engineering Discovery Environment) is the most advanced, powerful, and robust collection of integrated advanced digital resources and services in the world -- a single virtual system used by researchers, technologists, and scientists, such as yourself, to interactively share computing resources, data, and expertise.

Your feedback is vital to the evolution of this important resource, and I am writing to ask for your participation in the 2012 XSEDE User Satisfaction Survey conducted by Indiana University.

The annual survey aims to assess users' current levels of satisfaction with the XSEDE computational environment and its associated services and activities (e.g., training, allocations, conferences, user support, etc.). Further, this survey aims to quantify the satisfaction with, and availability of, cyberinfrastructure resources and support among members of US scientific and engineering communities, regardless of their use of, or access to, XSEDE resources. Survey information will be used to improve and expand services to the XSEDE user community and to the US scientific and research communities, broadly, and to aid in the decision-making processes related to resource allocation.

The survey can be accessed here:

https://websurv.indiana.edu/xsede/login.cfm?id=1AE7AE65EB

The Indiana University Center for Survey Research (CSR) administers the survey and assures that your responses will remain completely confidential. Neither your name nor your organization will be associated with any data or included in any reports.

If you have any questions about this survey or how the results will be used, please feel free to contact Julie Wernert, Information Manager, Indiana University, at jwernert@iu.edu, or call (812) 856-5517.

Thank you for your time and help with this important effort that will impact future decisions related to cyberinfrastructure funding for the US scientific and engineering communities.

Sincerely,

Craig A. Stewart, Ph.D.
Executive Director, Pervasive Technology Institute
Associate Dean, Research Technologies
Office of the Vice President for Information Technology
Indiana University

The IU Center for Survey Research is administering this questionnaire on behalf of the National Science Foundation-funded Extreme Science and Engineering Discovery Environment (XSEDE).

If you are unable to access the link listed above, please follow these instructions:

- In your web browser, type: websurv.indiana.edu/xsede
- In the Login Code box, enter: 1AE7AE65EB

If you have any other difficulties logging in or have questions about the study, please e-mail csr@indiana.edu for assistance.

If you do not wish to participate or receive further notices about this study, please use the instructions above to access the survey site. After logging in, select the button marked "I do not wish to participate."

Reference ID: 473434

D.1.3. Appendix 1c – Second Reminder Message

From: Stewart, Craig Alan
Sent: Wednesday, May 30, 2012 3:19 PM
To: [email address]
Subject: REMINDER: 2012 XSEDE User Satisfaction Survey

Dear [first name] [last name]:

Several weeks ago I wrote asking for your feedback on the XSEDE computational environment and its associated services and activities. Your feedback is vital to the evolution of this important resource, and I am writing again in the hope that you will take a few moments yet this week to complete the survey.

The survey can be accessed here:

https://websurv.indiana.edu/xsede/login.cfm?id=1AE7AE65EB

Thank you for your support and consideration.

Sincerely,

Craig A. Stewart, Ph.D.
Executive Director, Pervasive Technology Institute
Associate Dean, Research Technologies
Office of the Vice President for Information Technology
Indiana University

The Indiana University Center for Survey Research (CSR) administers the survey and assures that your responses will remain completely confidential. Neither your name nor your organization will be associated with any data or included in any reports.

If you have any questions about this survey or how the results will be used, please feel free to contact Julie Wernert, Information Manager, Indiana University, at jwernert@iu.edu, or call (812) 856-5517.

Thank you for your time and help with this important effort that will impact future decisions related to cyberinfrastructure funding for US scientific and engineering communities.

If you are unable to access the link listed above, please follow these instructions:

- In your web browser, type: websurv.indiana.edu/xsede
- In the Login Code box, enter: 1AE7AE65EB

If you have any other difficulties logging in or have questions about the study, please e-mail csr@indiana.edu for assistance.

If you do not wish to participate, or to receive further notices about this study, please follow the link to the survey and click on the button marked "I decline to participate."

Reference ID: 473434

D.1.4. Appendix 1d – Final Reminder Message

From: Stewart, Craig Alan
Sent: Tuesday, June 5, 2012 3:20 PM
To: [email address]
Subject: Final reminder: XSEDE Annual Survey to conclude Friday, June 8

Dear [first name] [last name]:

Several weeks ago I wrote asking for your feedback on the XSEDE computational environment and its associated services and activities. Your feedback is vital to the evolution of this important resource, and I am writing again in the hope that you will take a few moments yet this week to complete the survey.

The survey can be accessed here:

https://websurv.indiana.edu/xsede/login.cfm?id=1AE7AE65EB

Thank you for your support and consideration.

Sincerely,

Craig A. Stewart, Ph.D.
Executive Director, Pervasive Technology Institute
Associate Dean, Research Technologies
Office of the Vice President for Information Technology
Indiana University

The Indiana University Center for Survey Research (CSR) administers the survey and assures that your responses will remain completely confidential. Neither your name nor your organization will be associated with any data or included in any reports.

If you have any questions about this survey or how the results will be used, please feel free to contact Julie Wernert, Information Manager, Indiana University, at jwernert@iu.edu, or call (812) 856-5517.

Thank you for your time and help with this important effort that will impact future decisions related to cyberinfrastructure funding for US scientific and engineering communities.

If you are unable to access the link listed above, please follow these instructions:

- In your web browser, type: websurv.indiana.edu/xsede
- In the Login Code box, enter: 1AE7AE65EB

If you have any other difficulties logging in or have questions about the study, please e-mail csr@indiana.edu for assistance.

If you do not wish to participate, or to receive further notices about this study, please follow the link to the survey and click on the button marked "I decline to participate."

Reference ID: 473434

D.2. Appendix 2 – Survey Interface (Introductory and Conclusion Pages, Screenshots)

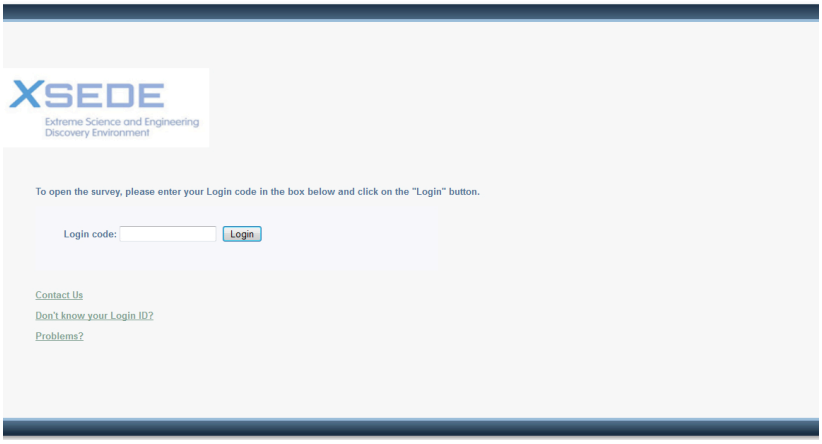

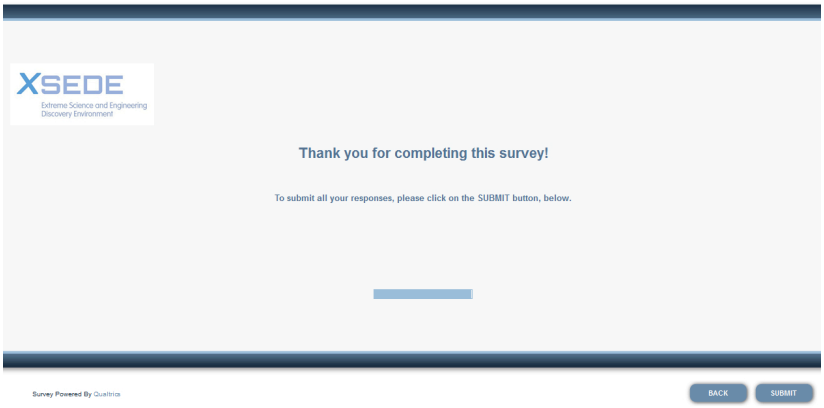

D.3. Appendix 3 – Study Information Sheet (screenshot)

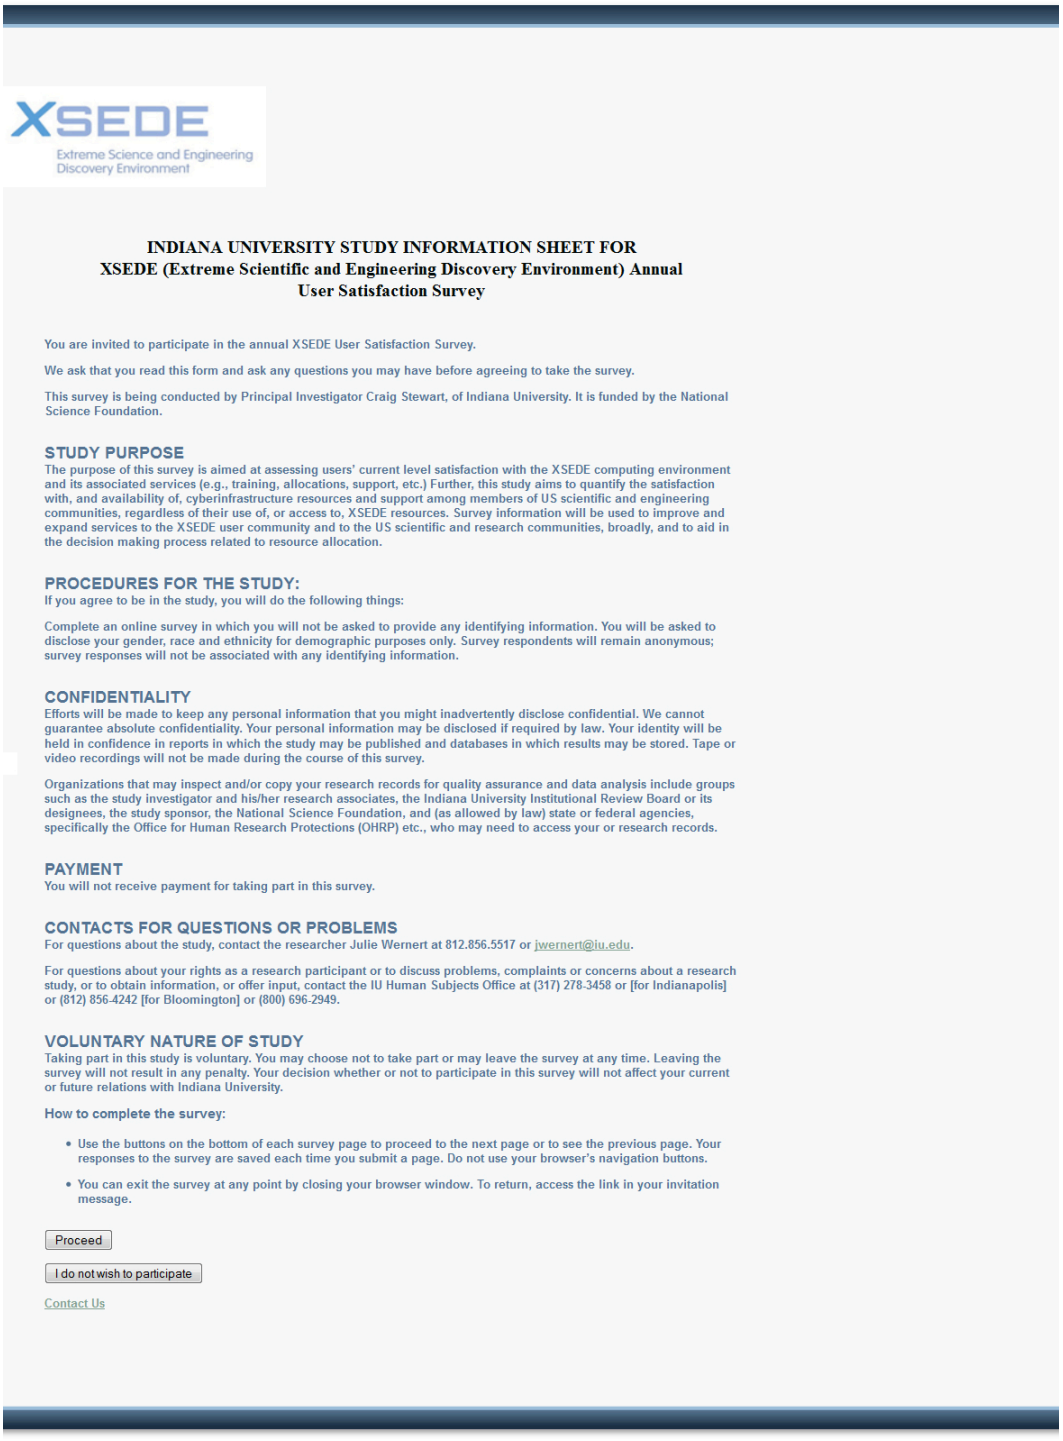

XSEDE
Extreme Science and Engineering Discovery Environment

INDIANA UNIVERSITY STUDY INFORMATION SHEET FOR
XSEDE (Extreme Scientific and Engineering Discovery Environment) Annual User Satisfaction Survey

You are invited to participate in the annual XSEDE User Satisfaction Survey.

We ask that you read this form and ask any questions you may have before agreeing to take the survey.

This survey is being conducted by Principal Investigator Craig Stewart, of Indiana University. It is funded by the National Science Foundation.

STUDY PURPOSE
The purpose of this survey is aimed at assessing users' current level satisfaction with the XSEDE computing environment and its associated services (e.g., training, allocations, support, etc.) Further, this study aims to quantify the satisfaction with, and availability of, cyberinfrastructure resources and support among members of US scientific and engineering communities, regardless of their use of, or access to, XSEDE resources. Survey information will be used to improve and expand services to the XSEDE user community and to the US scientific and research communities, broadly, and to aid in the decision making process related to resource allocation.

PROCEDURES FOR THE STUDY:
If you agree to be in the study, you will do the following things:

Complete an online survey in which you will not be asked to provide any identifying information. You will be asked to disclose your gender, race and ethnicity for demographic purposes only. Survey respondents will remain anonymous; survey responses will not be associated with any identifying information.

CONFIDENTIALITY
Efforts will be made to keep any personal information that you might inadvertently disclose confidential. We cannot guarantee absolute confidentiality. Your personal information may be disclosed if required by law. Your identity will be held in confidence in reports in which the study may be published and databases in which results may be stored. Tape or video recordings will not be made during the course of this survey.

Organizations that may inspect and/or copy your research records for quality assurance and data analysis include groups such as the study investigator and his/her research associates, the Indiana University Institutional Review Board or its designees, the study sponsor, the National Science Foundation, and (as allowed by law) state or federal agencies, specifically the Office for Human Research Protections (OHRP) etc., who may need to access your or research records.

PAYMENT
You will not receive payment for taking part in this survey.

CONTACTS FOR QUESTIONS OR PROBLEMS
For questions about the study, contact the researcher Julie Wernert at 812.856.5517 or jwernert@iu.edu.

For questions about your rights as a research participant or to discuss problems, complaints or concerns about a research study, or to obtain information, or offer input, contact the IU Human Subjects Office at (317) 278-3458 or [for Indianapolis] or (812) 856-4242 [for Bloomington] or (800) 696-2949.

VOLUNTARY NATURE OF STUDY
Taking part in this study is voluntary. You may choose not to take part or may leave the survey at any time. Leaving the survey will not result in any penalty. Your decision whether or not to participate in this survey will not affect your current or future relations with Indiana University.

How to complete the survey:

- Use the buttons on the bottom of each survey page to proceed to the next page or to see the previous page. Your responses to the survey are saved each time you submit a page. Do not use your browser's navigation buttons.
- You can exit the survey at any point by closing your browser window. To return, access the link in your invitation message.

Proceed

I do not wish to participate

Contact Us

D.4. Appendix 4 – Representative Question Interfaces (screenshots)

Exclusive choice (select only one)

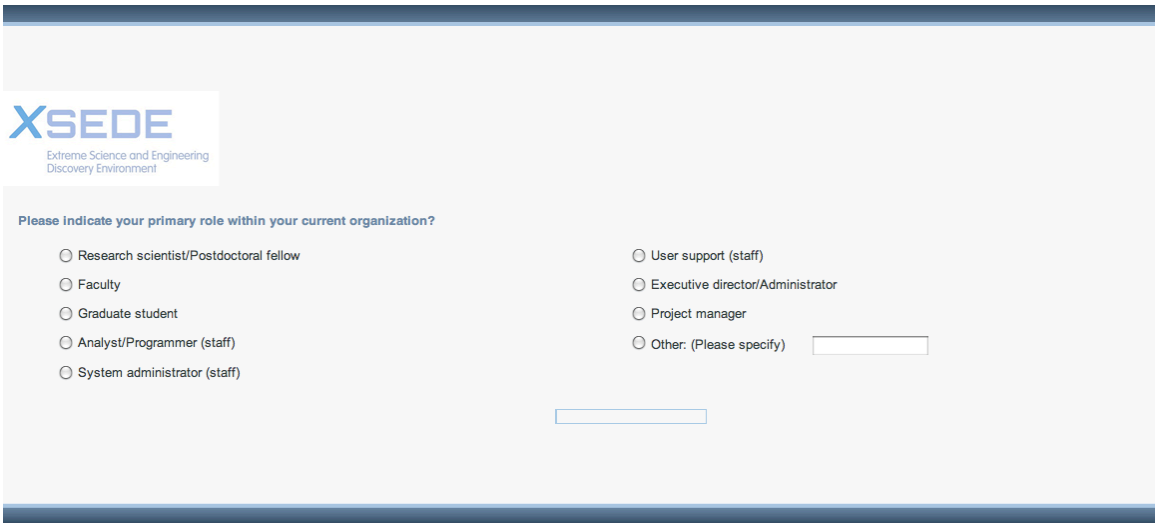

Multiple choice (check all that apply)

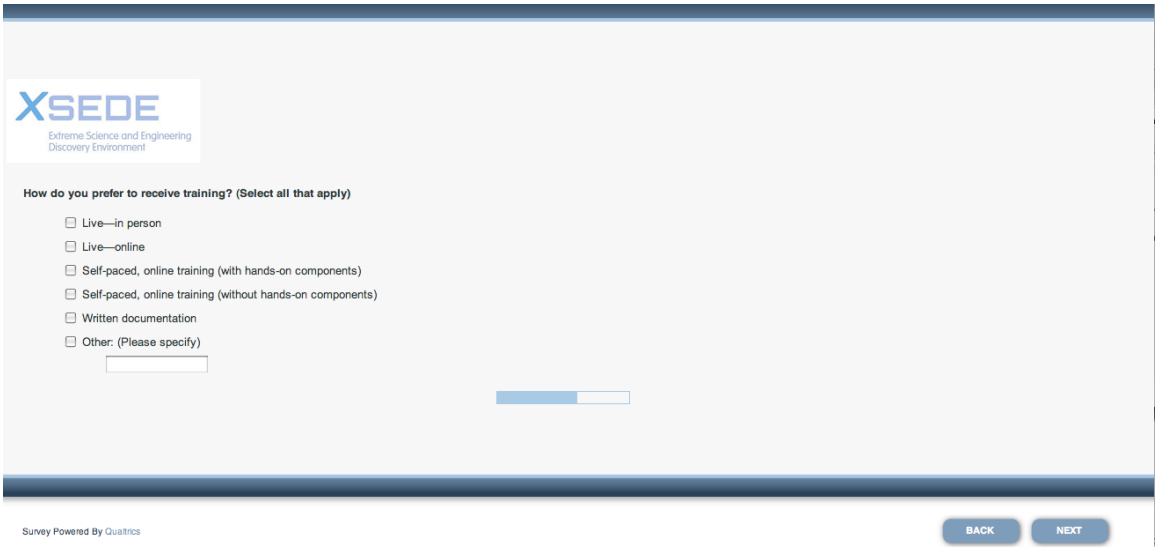

Satisfaction scale

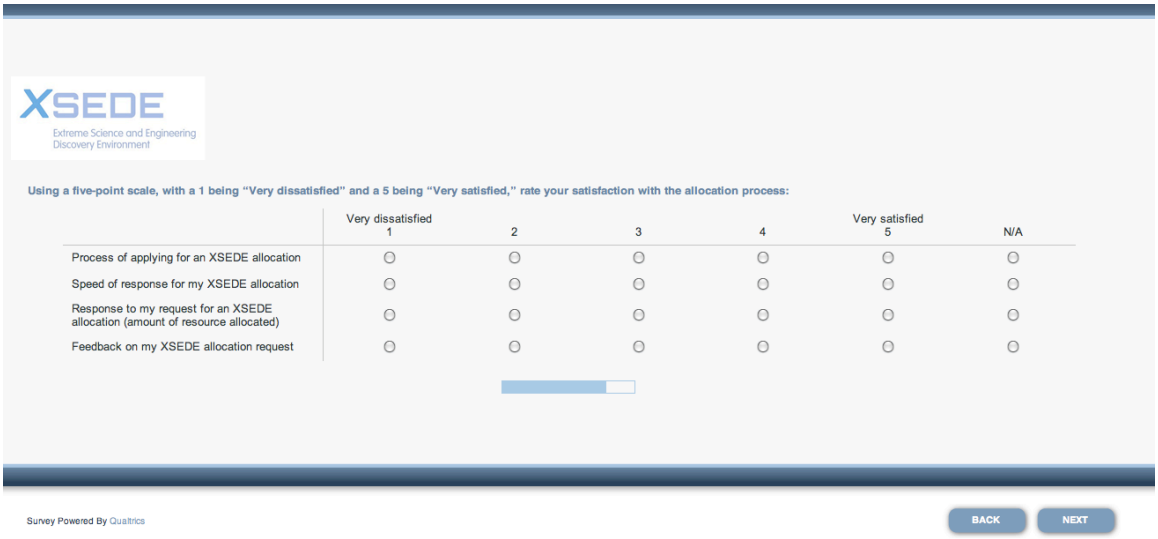

Open-ended, text response

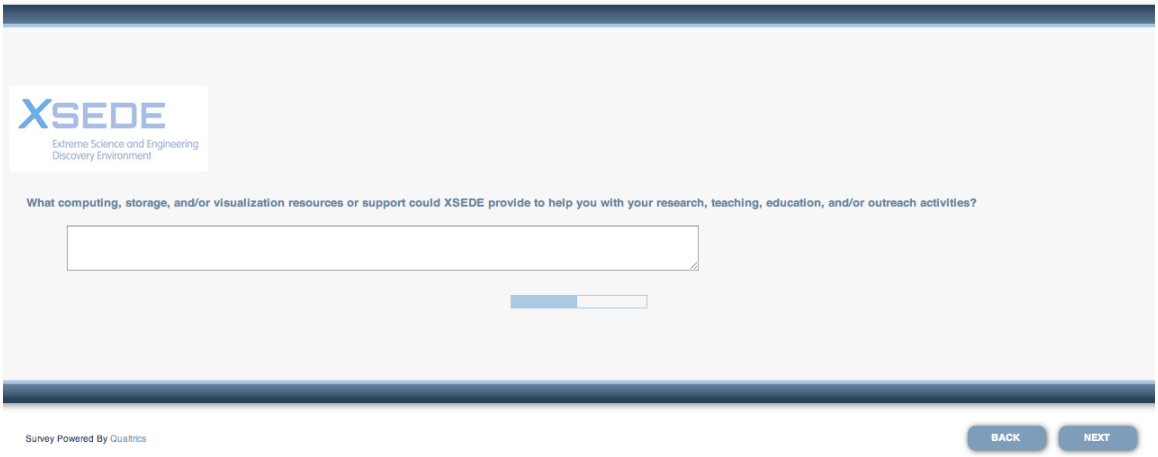

D.5. Appendix 5 - Complete Survey Text

What is your primary research field (as categorized by the National Science Foundation, National Institutes of Health, and/or the Department of Energy)?

- ❍ Astronomy
- ❍ Atmospheric Sciences
- ❍ Biology
- ❍ Chemistry
- ❍ Diseases
- ❍ Computer and Information Science
- ❍ Earth Science
- ❍ Engineering
- ❍ Health and Wellness
- ❍ Mathematics
- ❍ Medicine
- ❍ Physics
- ❍ Psychology
- ❍ Sociology
- ❍ Other: (Please specify) _____

Please indicate your primary role within your current organization?

- ❍ Research scientist/Postdoctoral fellow
- ❍ Faculty
- ❍ Graduate student
- ❍ Analyst/Programmer (staff)
- ❍ System administrator (staff)
- ❍ User support (staff)
- ❍ Executive director/Administrator
- ❍ Project manager
- ❍ Other: (Please specify) _____

What is your status relative to funding support from the National Science Foundation (NSF)? (Check all that apply)

- ❑ I am currently funded by an NSF award
- ❑ I am currently a Principal Investigator on an NSF award
- ❑ I am currently a Co-Investigator on an NSF award
- ❑ I am not currently funded by an NSF award, but I have been within the past five years
- ❑ Not applicable

What is your status relative to funding support from the National Institutes of Health (NIH)? (Check all that apply)

- ❑ I am currently funded by an NIH award
- ❑ I am currently a Principal Investigator on an NIH award
- ❑ I am currently a Co-Investigator on an NIH award
- ❑ I am not currently funded by an NIH award, but I have been within the past five years
- ❑ Not applicable

What is your status relative to funding support from the Department of Energy (DoE)? (Check all that apply)

- ❑ I am currently funded by a DoE award
- ❑ I am currently a Principal Investigator on a DoE award
- ❑ I am currently a Co-Investigator on a DoE award
- ❑ I am not currently funded by a DoE award, but I have been within the past five years
- ❑ Not applicable

Currently, how often do you have access to adequate cyberinfrastructure resources in support of your research activities?

- ❍ Never
- ❍ Some of the time
- ❍ Most of the time
- ❍ All of the time

Currently, how often do you have access to adequate technical support (e.g., information and consultants) in support of your use of cyberinfrastructure resources as part of your research activities?

- ❍ Never
- ❍ Some of the time
- ❍ Most of the time
- ❍ All of the time

How did you first learn about XSEDE?

- ❍ Colleague
- ❍ Advisor or principal investigator
- ❍ Internet
- ❍ Conference or workshop
- ❍ NSF announcement/website
- ❍ Legacy TeraGrid user
- ❍ Published research
- ❍ This survey
- ❍ Other: (Please specify) _____

If This survey instrument Is Selected, Then Skip To What is your gender?

Using the following table, please rate XSEDE for each of the aspects listed in the far left column. On a scale of 1 to 5, with 1 being "Completely unaware" and 5 being "Completely aware," please rate your awareness of the following aspects of XSEDE.

	Completely unaware 1		2	3	4
Mission	❍	❍	❍	❍	❍
Computational resources	❍	❍	❍	❍	❍
Support/Consulting desk services	❍	❍	❍	❍	❍
Training opportunities	❍	❍	❍	❍	❍
Education & outreach activities	❍	❍	❍	❍	❍
Main portal (xsede.org) and other Web-based resources	❍	❍	❍	❍	❍
XSEDE User Portal (portal.xsede.org)	❍	❍	❍	❍	❍
XSEDE documentation	❍	❍	❍	❍	❍
Storage services	❍	❍	❍	❍	❍
Visualization services	❍	❍	❍	❍	❍

How long have you used XSEDE (or TeraGrid) resources, and/or overseen the use of XSEDE (or TeraGrid) resources by others?

- ❍ Less than 6 months
- ❍ 6-11 months
- ❍ 1-2 years
- ❍ 3-5 years
- ❍ More than 5 years

How frequently did you use XSEDE and/or TeraGrid services in the past year?

- ❍ More than 50 times per year (once per week or more)
- ❍ 25-50 times per year
- ❍ 10-24 times per year
- ❍ 5-9 times per year
- ❍ 1-4 times per year
- ❍ I have never used XSEDE
- ❍ I have in the past, but no longer use XSEDE

If More than 50 times per year... Is Selected, Then Skip To Is your use of XSEDE primarily relate...If 25-50 times per year Is Selected, Then Skip To Is your use of XSEDE primarily relate...If 10-24 times per year Is Selected, Then Skip To Is your use of XSEDE primarily relate...If 5-10 times per year Is Selected, Then Skip To Is your use of XSEDE primarily relate...If 1-4 times per year Is Selected, Then Skip To Is your use of XSEDE primarily relate...

Why are you not currently using XSEDE services?

If Why are you not currently u... Is Displayed, Then Skip To What is your gender?

Is your use of XSEDE primarily related to research or education/outreach?

- ❍ Primarily research
- ❍ Primarily education
- ❍ Equal parts research and education

How important is XSEDE to your research?

- ❍ Essential; I would not be able to pursue my current research program without its use
- ❍ Very helpful; I would have difficulty pursuing my current research program without its use
- ❍ Helpful; It is useful, but I could pursue my research program without it
- ❍ Neutral; It is neither helpful nor unhelpful to my current research program
- ❍ Unhelpful; It is presently of negative net value to my research

How can we make XSEDE more useful to your research program?

What is your role relative to an XSEDE allocation? (Check all that apply)

- ❑ I am a Principal Investigator on an XSEDE allocation
- ❑ I am a Co-Investigator on an XSEDE allocation
- ❑ I am an XSEDE Allocation Manager
- ❑ I use XSEDE services personally
- ❑ I oversee the use of XSEDE by others
- ❑ Other: (Please specify) _____

What is your classification relative to an XSEDE? (Check all that apply)

- ❑ I am a Campus Champion or other volunteer
- ❑ Other: (Please specify) _____
- ❑ I am a faculty member
- ❑ I am a member of the research staff
- ❑ I am a postdoctoral fellow
- ❑ I am a graduate student
- ❑ I am an undergraduate student
- ❑ I am an XSEDE staff member

Please rate how satisfied you are with the XSEDE resources listed below in the left column, where 1 is "Very satisfied" and 5 is "Not at all satisfied".

	Very satisfied 1	2	3	4	Not at all satisfied 5	N/A
Capability (scalability) of XSEDE computational resources for simulation, particularly parallel processing applications	❍	❍	❍	❍	❍	❍
Capability (scalability) of XSEDE computational resources for data analysis, particularly parallel processing applications	❍	❍	❍	❍	❍	❍
Capacity (in terms of high throughput computing) of computational resources for simulation	❍	❍	❍	❍	❍	❍
Capacity (in terms of high throughput computing) of computational resources for data analysis	❍	❍	❍	❍	❍	❍
Ability to utilize Science Gateways to access XSEDE resources	❍	❍	❍	❍	❍	❍
Data archiving capabilities of XSEDE resources	❍	❍	❍	❍	❍	❍
Visualization facilities and rendering capabilities of XSEDE resources	❍	❍	❍	❍	❍	❍
Ability to receive support/consulting services from XSEDE	❍	❍	❍	❍	❍	❍
Other: (Please specify)	❍	❍	❍	❍	❍	❍

Resource Use - Computational All responses in this section should consider the period beginning June 2011 through the present.
What percentage of the time do you use the following computational resources. Please make sure your responses equal 100%.

_____ Local resources
_____ State or regional resources
_____ Commercial cloud resources
_____ DoE-funded resources (e.g., INCITE program)
_____ XSEDE or other NSF-funded grid resources
_____ Resources outside the US
_____ Other: (Please specify)

Which XSEDE computational services do you use? (Check all that apply)

- ❑ Ranger
- ❑ Lonestar
- ❑ Kraken
- ❑ Gordon
- ❑ Future Grid
- ❑ Blacklight
- ❑ Open Science Grid
- ❑ I do not use any of these resources
- ❑ Nautilus
- ❑ Forge
- ❑ Trestles
- ❑ Wispy
- ❑ Steele
- ❑ Purdue Condor
- ❑ Longhorn
- ❑ Spur

If I do not use any of these r... Is Selected, Then Skip To Which XSEDE data services do you use?...

Answer If Which XSEDE computational services do you use? (Check all... Ranger Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with Ranger:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability (uptime)	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Debugging tools	❍	❍	❍	❍	❍	❍
Performance	❍	❍	❍	❍	❍	❍
Scalability	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Archival storage access and allotment	❍	❍	❍	❍	❍	❍
Available 3rd Party Software (applications and libraries)	❍	❍	❍	❍	❍	❍
Queue configurations and operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE computational services do you use? (Check all... Lonestar Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with Lonestar:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability (uptime)	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Debugging tools	❍	❍	❍	❍	❍	❍
Performance	❍	❍	❍	❍	❍	❍
Scalability	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Archival storage access and allotment	❍	❍	❍	❍	❍	❍
Available 3rd Party Software (applications and libraries)	❍	❍	❍	❍	❍	❍
Queue configurations and operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE computational services do you use? (Check all... Kraken Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with Kraken:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability (uptime)	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Debugging tools	❍	❍	❍	❍	❍	❍
Performance	❍	❍	❍	❍	❍	❍
Scalability	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Archival storage access and allotment	❍	❍	❍	❍	❍	❍
Available 3rd Party Software (applications and libraries)	❍	❍	❍	❍	❍	❍
Queue configurations and operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE computational services do you use? (Check all... Gordon Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with Gordon:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability (uptime)	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Debugging tools	❍	❍	❍	❍	❍	❍
Performance	❍	❍	❍	❍	❍	❍
Scalability	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Archival storage access and allotment	❍	❍	❍	❍	❍	❍
Available 3rd Party Software (applications and libraries)	❍	❍	❍	❍	❍	❍
Queue configurations and operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE computational services do you use? (Check all... Blacklight Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with Blacklight:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability (uptime)	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Debugging tools	❍	❍	❍	❍	❍	❍
Performance	❍	❍	❍	❍	❍	❍
Scalability	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Archival storage access and allotment	❍	❍	❍	❍	❍	❍
Available 3rd Party Software (applications and libraries)	❍	❍	❍	❍	❍	❍
Queue configurations and operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE computational services do you use? (Check all... Trestles Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with Trestles:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability (uptime)	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Debugging tools	❍	❍	❍	❍	❍	❍
Performance	❍	❍	❍	❍	❍	❍
Scalability	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Archival storage access and allotment	❍	❍	❍	❍	❍	❍
Available 3rd Party Software (applications and libraries)	❍	❍	❍	❍	❍	❍
Queue configurations and operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE computational services do you use? (Check all... Steele Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with Steele:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability (uptime)	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Debugging tools	❍	❍	❍	❍	❍	❍
Performance	❍	❍	❍	❍	❍	❍
Scalability	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Archival storage access and allotment	❍	❍	❍	❍	❍	❍
Available 3rd Party Software (applications and libraries)	❍	❍	❍	❍	❍	❍
Queue configurations and operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE computational services do you use? (Check all... Forge Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with Forge:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability (uptime)	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Debugging tools	❍	❍	❍	❍	❍	❍
Performance	❍	❍	❍	❍	❍	❍
Scalability	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Archival storage access and allotment	❍	❍	❍	❍	❍	❍
Available 3rd Party Software (applications and libraries)	❍	❍	❍	❍	❍	❍
Queue configurations and operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE computational services do you use? (Check all... Purdue Condor Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with Purdue Condor:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability (uptime)	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Debugging tools	❍	❍	❍	❍	❍	❍
Performance	❍	❍	❍	❍	❍	❍
Scalability	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Archival storage access and allotment	❍	❍	❍	❍	❍	❍
Available 3rd Party Software (applications and libraries)	❍	❍	❍	❍	❍	❍
Queue configurations and operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE computational services do you use? (Check all... Open Science Grid Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with Open Science Grid:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability (uptime)	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Debugging tools	❍	❍	❍	❍	❍	❍
Performance	❍	❍	❍	❍	❍	❍
Scalability	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Archival storage access and allotment	❍	❍	❍	❍	❍	❍
Available 3rd Party Software (applications and libraries)	❍	❍	❍	❍	❍	❍
Queue configurations and operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE computational services do you use? (Check all... Wispy Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with Wispy:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability (uptime)	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Debugging tools	❍	❍	❍	❍	❍	❍
Performance	❍	❍	❍	❍	❍	❍
Scalability	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Archival storage access and allotment	❍	❍	❍	❍	❍	❍
Available 3rd Party Software (applications and libraries)	❍	❍	❍	❍	❍	❍
Queue configurations and operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE computational services do you use? (Check all... Future Grid Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with Future Grid:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability (uptime)	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Debugging tools	❍	❍	❍	❍	❍	❍
Performance	❍	❍	❍	❍	❍	❍
Scalability	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Archival storage access and allotment	❍	❍	❍	❍	❍	❍
Available 3rd Party Software (applications and libraries)	❍	❍	❍	❍	❍	❍
Queue configurations and operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE computational services do you use? (Check all... Nautilus Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with Nautilus:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability (uptime)	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Debugging tools	❍	❍	❍	❍	❍	❍
Performance	❍	❍	❍	❍	❍	❍
Scalability	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Archival storage access and allotment	❍	❍	❍	❍	❍	❍
Available 3rd Party Software (applications and libraries)	❍	❍	❍	❍	❍	❍
Queue configurations and operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE computational services do you use? (Check all... Longhorn Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with Longhorn:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability (uptime)	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Debugging tools	❍	❍	❍	❍	❍	❍
Performance	❍	❍	❍	❍	❍	❍
Scalability	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Archival storage access and allotment	❍	❍	❍	❍	❍	❍
Available 3rd Party Software (applications and libraries)	❍	❍	❍	❍	❍	❍
Queue configurations and operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE computational services do you use? (Check all... Spur Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with Spur:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability (uptime)	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Debugging tools	❍	❍	❍	❍	❍	❍
Performance	❍	❍	❍	❍	❍	❍
Scalability	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Archival storage access and allotment	❍	❍	❍	❍	❍	❍
Available 3rd Party Software (applications and libraries)	❍	❍	❍	❍	❍	❍
Queue configurations and operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Which XSEDE data services do you use? (Check all that apply)

- ❑ Albedo (and the Data Replication Service)
- ❑ IU Data Capacitor
- ❑ I do not use any of these resources
- ❑ TACC Ranch
- ❑ NCSA MSS
- ❑ PSC Data Archive
- ❑ NICS HPSS

If I do not use any of these r... Is Selected, Then Skip To What storage and computing resources ...

Answer If Which XSEDE data services do you use? (Check all that apply) Albedo (and the Data Replication Service) Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with ${q://QID24/ChoiceDescription/1}:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability	❍	❍	❍	❍	❍	❍
Input/output performance	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Storage access and allotment	❍	❍	❍	❍	❍	❍
Operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User News	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE data services do you use? (Check all that apply) IU Data Capacitor Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with ${q://QID24/ChoiceDescription/2}:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability	❍	❍	❍	❍	❍	❍
Input/output performance	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Storage access and allotment	❍	❍	❍	❍	❍	❍
Operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User News	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE data services do you use? (Check all that apply) TACC Ranch Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with TACC Ranch:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability	❍	❍	❍	❍	❍	❍
Input/output performance	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Storage access and allotment	❍	❍	❍	❍	❍	❍
Operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE data services do you use? (Check all that apply) NCSA MSS Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with NCSA MSS:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability	❍	❍	❍	❍	❍	❍
Input/output performance	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Storage access and allotment	❍	❍	❍	❍	❍	❍
Operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE data services do you use? (Check all that apply) PSC Data Archive Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with PSC Data Archive:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability	❍	❍	❍	❍	❍	❍
Input/output performance	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Storage access and allotment	❍	❍	❍	❍	❍	❍
Operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Answer If Which XSEDE data services do you use? (Check all that apply) NICS HPSS Is Selected

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with NICS HPSS:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability	❍	❍	❍	❍	❍	❍
Input/output performance	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Storage access and allotment	❍	❍	❍	❍	❍	❍
Operational policies	❍	❍	❍	❍	❍	❍
Documentation	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

What computing, storage, and/or visualization resources or support could XSEDE provide to help you with your research, teaching, education, and/or outreach activities?

Are there software tools or libraries needed for your work that are not available on XSEDE?

Science Gateways XSEDE Science Gateways are portals to computational and data services and resources across a wide range of science domains for researchers, engineers, educators, and students. Depending on the needs of the communities, a gateway may provide any of the following features: high-performance computation resources; workflow tools; general or domain-specific analytic and visualization software; collaborative interfaces; job submission tools; and/or education modules. See the full list of Science Gateways accessible through XSEDE at: xsede.org/web/guest/gateways-listing Do you access XSEDE resources through Science Gateways?

- ❍ Yes
- ❍ No

If No Is Selected, Then Skip To Please tell us why you are not curr...

Which science gateways do you use? (Check all that apply) Skin #SkinContent { width:1400 px;

- ❑ Asteroseismic Modeling Portal
- ❑ Biodrugscore: A portal for customized scoring and ranking of molecules docked to the human proteome
- ❑ Center for Multiscale Modeling of Atmospheric Processes
- ❑ Chemical Informatics and Cyberinfrastructure Collaboratory
- ❑ CIG Science Gateway for the Geodynamics Community
- ❑ CIPRES Portal for inference of large phylogenetic trees
- ❑ Community Climate System Model (CCSM) TeraGrid Gateway
- ❑ Computational Chemistry Grid (GridChem)
- ❑ Cyberinfrastructure for End-to-End Environmental Exploration Portal
- ❑ Dark Energy Survey Data Management
- ❑ Developing Social Informatics Data Grid (SIDGrid)
- ❑ EPSCoR Desktop to TeraGrid EcoSystem
- ❑ Globus Online
- ❑ High Resolution Daily Temperature and Precipitation Data for the Northeast United States
- ❑ High-Resolution Modeling of Hydrodynamic Experiments with UltraScan
- ❑ Indiana University Centralized Life Sciences Data
- ❑ Isoscapes modeling, analysis and prediction (IsoMAP)
- ❑ Linked Environments for Atmospheric Discovery
- ❑ Massive Pulsar Surveys using the Arecibo L-band Feed Array (ALFA)
- ❑ National Biomedical Computation Resource
- ❑ Network for Computational Nanotechnology and nanoHUB
- ❑ Network for Earthquake Engineering Simulation
- ❑ Neutron Science TeraGrid Gateway
- ❑ OGCE Science Gateway Portal
- ❑ Open Science Grid
- ❑ Purdue Environmental Data Portal
- ❑ QuakeSim
- ❑ ROBETTA: Automated Prediction of Protein Structure and Interactions
- ❑ SCEC Earthworks Project
- ❑ Science Gateway for Diffraction Facilities, Data and Methods
- ❑ Social Science Gateway
- ❑ TeraGrid Geographic Information Science Gateway
- ❑ The Earth System Grid
- ❑ User-Friendly Security Solutions for Grid Environments
- ❑ VLab - Virtual Laboratory for Earth and Planetary Materials
- ❑ Other: (Please specify) _____

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your overall satisfaction with Science Gateways:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Ease of use	❍	❍	❍	❍	❍	❍
Reliability	❍	❍	❍	❍	❍	❍
Input/Output Performance	❍	❍	❍	❍	❍	❍
Value of use of system to my research	❍	❍	❍	❍	❍	❍
Data transfer capability	❍	❍	❍	❍	❍	❍
Storage access and allotment	❍	❍	❍	❍	❍	❍
Operational policies	❍	❍	❍	❍	❍	❍
Job submission	❍	❍	❍	❍	❍	❍
Work flow management	❍	❍	❍	❍	❍	❍

How do you prefer to access XSEDE resources?

- ❍ Logging in directly and issuing commands via ssh shell
- ❍ Using Web or desktop graphical tools to assist in submitting/monitoring jobs and accessing files
- ❍ Initially setting up application codes, with subsequent routine access via graphical tools to run jobs, move data, etc.
- ❍ Other: (Please specify) _____

Answer If Science Gateways XSEDE Science Gateways are portals ... No Is Selected

Please tell us why you are not currently using Science Gateways: (Check all that apply)

- ❑ I have not found a gateway that meets my needs
- ❑ Gateways are too complicated to use
- ❑ Gateways do not provide enough flexibility
- ❑ Gateways are not reliable
- ❑ Support level was not meeting my expectations
- ❑ Other: (please specify) _____

Which XSEDE services do you plan to use in the upcoming year? (Check all that apply)

- ❑ Ranger
- ❑ Lonestar
- ❑ Kraken
- ❑ Gordon
- ❑ Blacklight
- ❑ Open Science Grid
- ❑ Albedo (and the Data Replication Service)
- ❑ IU Data Capacitor
- ❑ PSC Data Archive
- ❑ Other: (Please specify) _____
- ❑ Nautilus
- ❑ Keeneland
- ❑ Trestles
- ❑ Wispy
- ❑ Longhorn
- ❑ TACC Ranch
- ❑ NCSA MSS
- ❑ NICS HPSS
- ❑ I do not plan to use any of these resources
- ❑ Science Gateways
- ❑ Steele
- ❑ Purdue Condor
- ❑ Future Grid
- ❑ Spur

Resource Use - Support All responses in this section should consider the period beginning June 2011 through the present. On a scale from 1 to 5, how useful are the following resources in helping you learn to use XSEDE and/or its predecessor, TeraGrid?

	Not at all useful 1	2	3	4	Very useful 5	N/A
Independent, self-study (online documentation, man pages, trial-and-error, etc.)	❍	❍	❍	❍	❍	❍
Colleagues	❍	❍	❍	❍	❍	❍
Conferences	❍	❍	❍	❍	❍	❍
Attending onsite workshops	❍	❍	❍	❍	❍	❍
Synchronous online workshops	❍	❍	❍	❍	❍	❍
Online, self-paced tutorials	❍	❍	❍	❍	❍	❍
Other: (Please specify)	❍	❍	❍	❍	❍	❍

How do you prefer to receive training? (Select all that apply)

- ❑ Live—in person
- ❑ Live—online
- ❑ Self-paced, online training (with hands-on components)
- ❑ Written documentation
- ❑ Other: (Please specify) _____
- ❑ Self-paced, online training (without hands-on components)

What specific types of training would be of most interest to you? (Select all that apply) Skin #SkinContent { width:1400 px;

- ❑ Tuning and Optimization
- ❑ Debugging
- ❑ Managing I/O
- ❑ Visualization tools
- ❑ Training on specific applications or libraries: (Please specify) _____
- ❑ Data analysis and management tools
- ❑ Introductory programming topics (e.g., Fortran, C, C++)
- ❑ Introduction to UNIX
- ❑ Writing a successful XSEDE allocation request
- ❑ Introduction to XSEDE
- ❑ Programming Accelerators (GPGPU Programming, MIC Programming)

What training resources or support could XSEDE provide to help with your research, teaching, education, and/or outreach activities?

Is there a resource person at your institution available to assist with your use of XSEDE? (Check all that apply)

- ❑ XSEDE staff member
- ❑ XSEDE Campus Champion
- ❑ Local IT support person (i.e., an individual not designated as an XSEDE Campus Champion)
- ❑ No
- ❑ I do not know
- ❑ Colleague (faculty, post-doc, graduate student, etc.) at my institution

Answer If Is there a resource person at your institution available ... XSEDE Campus Champion Is Selected

You indicated that your institution has an XSEDE Campus Champion. Are you the XSEDE Campus Champion for your institution?

- ❍ Yes
- ❍ No

Answer If Is there a resource person at your institution available ... Local IT support person (i.e., an individual not designated as an XSEDE Campus Champion) Is Selected

You indicated that you have a local IT support person available to assist you with XSEDE. Are you this person?

- ❍ Yes
- ❍ No

On a 1-5 scale, with a 1 being “Very dissatisfied” and 5 being “Very satisfied,” please rate your experience with XSEDE user support staff:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
How satisfied are you with the effectiveness of problem resolution provided by XSEDE user support staff?	❍	❍	❍	❍	❍	❍
How satisfied are you with the helpfulness of XSEDE user support staff?	❍	❍	❍	❍	❍	❍
How satisfied are you with the knowledge of XSEDE user support staff?	❍	❍	❍	❍	❍	❍
How satisfied are you with the communication skills of XSEDE user support staff?	❍	❍	❍	❍	❍	❍
How satisfied are you with the timeliness of responses from XSEDE user support staff?	❍	❍	❍	❍	❍	❍
How satisfied are you with notifications and announcements provided by XSEDE user support staff?	❍	❍	❍	❍	❍	❍
How satisfied are you with the assistance provided by your XSEDE Campus Champion, local XSEDE staff member, or other local support providers?	❍	❍	❍	❍	❍	❍
How satisfied are you with the courtesy and demeanor exhibited by XSEDE support staff	❍	❍	❍	❍	❍	❍

Do you use the XSEDE User Portal (XUP)?

- ❍ Yes
- ❍ No

If No Is Selected, Then Skip To Do you use the XSEDE user support doc...

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with the following XSEDE User Portal (XUP) features.

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Allocation management	❍	❍	❍	❍	❍	❍
Logging into systems	❍	❍	❍	❍	❍	❍
File management	❍	❍	❍	❍	❍	❍
User forums	❍	❍	❍	❍	❍	❍
Resource monitor	❍	❍	❍	❍	❍	❍
My Jobs	❍	❍	❍	❍	❍	❍
Queues prediction	❍	❍	❍	❍	❍	❍
User news	❍	❍	❍	❍	❍	❍

Do you use XUP Mobile?

- ❍ Yes
- ❍ No
- ❍ I am unaware of this service

Do you use the XSEDE user support documentation?

- ❍ Yes
- ❍ No

If No Is Selected, Then Skip To Resource Use - Allocations ...

Did you find the XSEDE user support documentation helpful?

- ❍ Yes
- ❍ No

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with the following activities and services.

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
XSEDE.org web site	❍	❍	❍	❍	❍	❍
XSEDE User Portal (XUP)	❍	❍	❍	❍	❍	❍
XSEDE help desk services via help@xsede.org	❍	❍	❍	❍	❍	❍
XSEDE help desk services via telephone (866-907-2383)	❍	❍	❍	❍	❍	❍
XSEDE help desk services via XUP	❍	❍	❍	❍	❍	❍
XSEDE Knowledge Base	❍	❍	❍	❍	❍	❍
XSEDE user news	❍	❍	❍	❍	❍	❍
XSEDE online user services documentation (user guides, usage policies, etc.)	❍	❍	❍	❍	❍	❍

What new features would you like to see in the XSEDE User Portal or in the documentation?

Resource Use - Allocations All responses in this section should consider the period beginning June 2011 through the present. Do you use the POPS allocation submission process? That is, have you submitted a request for an XSEDE allocation?

- ❍ Yes
- ❍ No

If No Is Selected, Then Skip To Resource Use - Extended Collaborati...

Using a five-point scale, with a 1 being "Very dissatisfied" and a 5 being "Very satisfied," rate your satisfaction with the allocation process:

	Very dissatisfied 1	2	3	4	Very satisfied 5	N/A
Process of applying for an XSEDE allocation						
Speed of response for my XSEDE allocation						
Response to my request for an XSEDE allocation (amount of resource allocated)						
Feedback on my XSEDE allocation request						

Resource Use - Extended Collaborative Support Services, formerly Advanced User Support All responses in this section should consider the period beginning June 2011 through the present. Are you aware of the XSEDE Extended Collaborative Support Services (ECSS)?

- ❍ Yes
- ❍ No

If No Is Selected, Then Skip To Future Needs Do you have projects t...

Do you know how to request Extended Collaborative Support Services (ECSS)?

- ❍ Yes
- ❍ No

Have you received assistance from Extended Collaborative Support Services (ECSS)?

- ❍ Yes
- ❍ No

If No Is Selected, Then Skip To Future Needs Do you have projects t...

In what area(s) have your projects received Extended Collaborative Support Services (ECSS)? (Check all that apply)

- Advanced Research Team Support
- Novel and Innovative Projects
- Advanced Community Capabilities Support
- Advanced Science Gateways Support
- Advanced EOT support
- I am not certain in which area(s) my project(s) received ECSS

What assistance did Extended Collaborative Support Services (ECSS) provide? (Check all that apply)

- Performance Optimization (single process or parallel)
- Application Scalability
- Parallelism (MPI, OpenMP, or Accelerators)
- I/O Optimization
- Application Gateways (including grid computing)
- Other: (Please specify) _____

On a 1-5 scale, with a 1 being "Very dissatisfied" and 5 being "Very satisfied," how satisfied are you with your Extended Collaborative Support Services (ECSS) experience?

- 1 Very dissatisfied
- 2
- 3
- 4
- 5 Very satisfied
- N/A

Future Needs Do you have projects that could benefit from collaboration with Extended Collaborative Support Services (ECSS)?

- Yes
- No
- I don't know

If No Is Selected, Then Skip To Accelerator availability and support

Do you intend to request Extended Collaborative Support Services (ECSS) for your project?

- Yes
- No
- I don't know

If No Is Selected, Then Skip To Accelerator availability and support

What assistance do you expect to request from Extended Collaborative Support Services (ECSS)? (Check all that apply)

- Performance optimization (single process or parallel)
- Application scalability
- Parallelism (MPI, OpenMP, or accelerators)
- I/O optimization
- Application gateways (including grid computing)
- Other: (Please specify) _____

Have you ever received Extended Collaborative Support Services (ECSS) from someone at one XSEDE site for a system located at another XSEDE site?

- Yes
- No
- I don't know

Do you have any comments about the processes for obtaining Allocations (POPS) and/or Extended Collaborative Support Services (ECSS), or other general comments about Allocations and ECSS?

- Accelerator availability and support Which of the following accelerators or co-processors do you currently use? (Please select all that apply)
- GPGPUs
- FPGAs
- Other _____

Which of the following accelerators or co-processors do you plan to use in the next year? (Please select all that apply)

- GPGPUs
- Intel MIC co-processor
- FPGAs
- Other _____

Which programming models do you currently use on accelerators or co-processors?

- CUDA
- OpenCL
- Other: (Please specify) _____
- OpenACC

Which programming models do you plan to use on accelerators or co-processors in the next year?

- CUDA
- OpenCL
- Other: (Please specify) _____
- OpenACC
- OpenMP offload execution on Intel MIC
- Native execution directly on Intel MIC

Please provide any comments you have on the value derived from the NSF's investment in XSEDE:

Do you have any other suggestions or comments regarding XSEDE and/or other NSF-funded cyberinfrastructure?

Now, we have some questions about you.

What is your gender?

- Male
- Female

What is your ethnicity? For definitions of Hispanic or Latino, place your cursor over the Hispanic or Latino category.

- Hispanic or Latino
- Not Hispanic or Latino

What is your race? Please select all that apply. For definitions of each race, place your cursor over each category.

- American Indian or Alaska Native
- Asian
- Black or African-American
- Native Hawaiian or Other Pacific Islander
- White

Thank you for completing this survey! To submit all your responses, please click on the SUBMIT button, below.