2018 Oak Ridge Leadership Computing Facility User Survey

Findings and Recommendations

PREPARED BY:

Ann M. Martin, Ph.D.
Senior Evaluator
Scientific Assessment & Workforce Development

Kelly P. Townsend
Evaluation Specialist
Scientific Assessment & Workforce Development

Karen Hawley
Senior Program Specialist
Scientific Assessment & Workforce Development

Erin M. Burr, Ph.D.
Senior Evaluator and Section Manager, Assessment & Evaluation
Scientific Assessment & Workforce Development

PREPARED FOR:

Oak Ridge Leadership Computing Facility
Oak Ridge National Laboratory
February 2019

The opinions expressed herein do not necessarily reflect the opinions of the sponsoring institutions of Oak Ridge Associated Universities.

This report was prepared as an account of work sponsored by the United States Government. Neither the United States Government nor the U.S. Department of Energy, nor any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe on privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, mark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement or recommendation, or favor by the U.S. Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the U.S. Government or any agency thereof.

The Oak Ridge Institute for Science and Education (ORISE) is a U.S. Department of Energy (DOE) asset that is dedicated to enabling critical scientific, research, and health initiatives of the department and its laboratory system by providing world-class expertise in STEM workforce development, scientific and technical reviews, and the evaluation of radiation exposure and environmental contamination.

ORISE is managed by ORAU, a 501(c)(3) nonprofit corporation and federal contractor, for DOE's Office of Science. The single largest supporter of basic research in the physical sciences in the United States, the Office of Science is working to address some of the most pressing challenges of our time. For more information, please visit science.energy.gov.

ORAU provides innovative scientific and technical solutions to advance national priorities in science, education, security and health. Through specialized teams of experts, unique laboratory capabilities and access to a consortium of more than 100 major Ph.D.-granting institutions, ORAU works with federal, state, local and commercial customers to advance national priorities and serve the public interest. A 501(c)(3) nonprofit corporation and federal contractor, ORAU manages the Oak Ridge Institute for Science and Education (ORISE) for the U.S. Department of Energy (DOE). Learn more about ORAU at www.orau.org.

Table of Contents

List of Tables

List of Figures

Executive Summary

In an effort to promote continual improvement at the Oak Ridge Leadership Computing Facility (OLCF), users were sent a survey soliciting their feedback regarding their experience as a user of the facilities and support services.

Respondents

At the end of the nine-week survey period, 422 users completed the survey out of 1,230 possible respondents, giving an overall response rate of 34.3%. Respondents' projects were supported by Director's Discretion (64%), INCITE (38%), ALCC (21%), ECP (19%), and Other (4%) sources.

Findings Highlights

Overall Evaluation

The proportions of all respondents *satisfied*, *or very satisfied* with OLCF resources/services, ranged from 91% to 96% for "overall" evaluation items. Specifically, ratings for major categories of resources/services were a) OLCF (96%; continuing a slow, but steady increase from 86% in 2007), b) Compute Resources (94%), c) Data Resources (91%), and d) Support Services (94%). Thematic analysis of open-ended comments identified *computing power/performance* (56% of respondents) and *user technical support /staff* (36% of respondents) as the most valued OLCF qualities.

The table below indicates satisfaction (*satisfied* or *very satisfied*) ratings. The color scale indicates the relative magnitude of cell values: high-medium-low = green-yellow-red.

	All	PI	Non-PI	INCITE	DD	ALCC	ECP
Max N responding:	*419*	*87*	*332*	*159*	*271*	*88*	*80*
OLCF	96%	97%	95%	94%	97%	95%	91%
Compute Resources	94%	94%	94%	95%	94%	94%	92%
Titan	95%	96%	95%	94%	95%	95%	93%
Eos	99%	100%	98%	100%	98%	100%	100%
Rhea	100%	100%	100%	100%	100%	100%	100%
Data Resources	91%	96%	89%	88%	91%	91%	80%
Data Transfer Nodes	87%	85%	88%	86%	88%	91%	92%
HPSS	95%	91%	96%	93%	94%	91%	100%
Lustre/Spider	92%	91%	92%	89%	92%	88%	84%
Support Services	94%	97%	93%	94%	95%	94%	91%
User Assistance	97%	98%	97%	98%	97%	95%	94%
Account Services	96%	100%	93%	97%	95%	100%	96%
Data Analysis and Visualization Support Services	100%	100%	100%	100%	100%	100%	100%
INCITE Liaison	95%	93%	96%	94%	97%	84%	78%
Communication	91%	98%	89%	89%	91%	91%	93%
Training and Technical Reference Documentation	94%	90%	95%	92%	93%	94%	92%
OLCF Website	93%	89%	94%	92%	93%	89%	90%

OLCF Systems, Data Resources, and Compute Resources

Titan, Eos, and Rhea are all used at similar rates compared to results from 2014-2017. Most users (74%) noted no changes in overall OLCF computing performance over the last year, while 24% cited improved performance; only 2% noted a decrease in performance compared to 2017. Overall satisfaction across the compute resources and data resources ranged from 87% (data transfer nodes) to 100% (Rhea) of users either *satisfied* or *very satisfied*. 89% of users were *satisfied* or *very satisfied* with project disk space, and notice for scheduled maintenance and bandwidth offered by the OLCF were rated this highly by more than 90% of users (94% and 91%, respectively). Given the opportunity to rank the importance of potential future data services or features, the highest ranked options were *long-term data retention* (64%), *access for your specific OLCF project members to your data over the web* (46%), and *long-term data curation* (46%).

Support Services

The User Assistance Center (UAC) was the most highly utilized support service (58%), and nearly all users were satisfied with it (97%). This was followed by 23% using the Account Services, 21% with assigned INCITE Scientific Computing/Liaisons, and 4% using Data Analysis and Visualization support services. When considering all users of support services, satisfaction levels ranged from 91% for communication to 100% for data analysis and visualization support services.

Communication with Users

91% of respondents were overall *satisfied* or *very satisfied* with communication from OLCF. The communication activity that received the least positive rating were the *monthly conference calls* (86%). Nearly all respondents felt adequately informed about *OLCF changes* (97%), *events* (98%), and *current issues* (96%).

Training and Technical Reference Documentation

94% of respondents were satisfied with OLCF training and technical reference documentation overall, with the *training calendar* (85% satisfaction) receiving the lowest satisfaction rating among specific aspects of training. The most preferred ways of receiving training were via *online documentation* (79%), or in an *online training format* (59%). Most expressed no preference as to time of year (56%); among those with a preference, two-thirds chose the summer. Respondents suggested future training topics in 21 categories. The most frequently suggested topics were *common tasks and OLCF basics, software setup for project-specific needs, Summit, GPU resources, Python,* and *example scripts* (16%, 10%, 8%, 8%, 8%, 8%, respectively).

Web Site

30% of respondents indicated they visit the OLCF website (http://olcf.ornl.gov) once a week or more frequently. More than 9 in 10 respondents indicated they were satisfied with the OLCF Website (93%). *Search capabilities* were the lowest rated aspect of website usability (86% satisfaction). The most commonly requested feature for the new MyOLCF portal was *allocation, quotas, and usage* information (51%).

Data Analysis, Visualization, and Workflow

100% of respondents were *satisfied* or *very satisfied* with the analysis and visualization support services and with the quality of technical support. Satisfaction with data analysis, visualization, and workflow at OLCF ranged from 73% for *sufficiency of software tools* to 79% for *ability to perform data analysis* and *sufficiency of hardware* for analysis and visualization needs. With respect to workflow and analysis, the largest proportion of respondents (34%) analyzed all of their data "elsewhere" and the smallest proportion (7%) analyzed it all at OLCF. About 12% analyzed most of their data at OLCF. Only 20% of users indicated that all of their data was sourced from their OLCF projects, while 47% of users were working with data mostly from outside OLCF.

Introduction

A survey was conducted to gather information about the users of the Oak Ridge Leadership Computing Facility (OLCF) at Oak Ridge National Laboratory (ORNL). The survey collected feedback about user needs, preferences, and experience with OLCF and its support capabilities. Attitudes and opinions on the performance, availability, and possible improvements of OLCF resources/services were also solicited. The survey was created by the Assessment and Evaluation team within Oak Ridge Associated Universities (ORAU), in collaboration with OLCF staff. OLCF staff also provided email addresses and data on the characteristics of OLCF users.

This report first briefly describes the data collection and analysis procedures. It then presents findings with respect to user characteristics, patterns of OLCF resource use, and satisfaction ratings of OLCF resources/services. The report also provides longitudinal comparisons of user responses from 2006 through 2018. Finally, recommendations for possible improvements are offered.

Data Collection and Analysis

Data Collection

The survey sampling frame was constituted by first collecting the names of individuals who had logged into an OLCF system between 1/1/2018 through 9/30/2018. OLCF staff and vendors as well as individuals with invalid email addresses were then removed from the list. An additional 11 users were identified and added to the user group after they had responded, since visitors to the OLCF website and others on OLCF distribution lists could also access the survey. Overall, this process resulted in a sampling frame with 1,230 OLCF users.

ORAU invited all OLCF users from this list to participate in the survey, which was hosted online beginning on October 9, 2018 and remained open for completion through December 12, 2018 (Appendix B: Survey Administration Timeline and Appendix C: Survey). A total of 422 users completed or partially completed the survey, resulting in a response rate of 34.3%. Figure 26, within Appendix B: Survey Administration Timeline, highlights the value of each reminder email in increasing the response rate.

The survey first asked respondents about their experience and patterns of use with OLCF resources/services, and then asked for their satisfaction with resources/services in the following main categories (bold) and subcategories (Appendix C: Survey):

OLCF (Overall)

OLCF Computing Resources

- Titan
- Eos
- Rhea

OLCF Data Resources

- Data Transfer Nodes (DTNs)
- HPSS
- Lustre/Spider

OLCF Support Services

- User Assistance
- Account Management
- INCITE Scientific Computing Liaison
- Communication with users
- Training and Technical Reference Documentation
- OLCF Website and MyOLCF
- Data analysis and visualization

Data Analysis

The findings section typically presents results summarized numerically that report responded levels of satisfaction. This is followed by a verbal summary of the open-ended comments from individuals that indicated being dissatisfied (via the scaled reply) with a resource or service (note: not all dissatisfied individuals supplied open-ended comments).

As noted, the survey assessed satisfaction with OLCF resources/services using a 5-point scale, from *Very Dissatisfied* (1) to *Very Satisfied* (5). These **closed-ended responses** were summarized using frequency distributions, proportions, means, and standard deviations. The proportion of respondents indicating either a 4 (*Satisfied*) or 5 *(Very satisfied)* on an item was also typically reported as %Sat to provide a summary measure. This measure was also used to indicate the relative satisfaction with resources/services within categories. Respondents that were *Very dissatisfied* or *Dissatisfied* with OLCF resources/services were asked to provide comments explaining their dissatisfaction (see below).

In order to better understand the types of OLCF users and how needs and preferences varied, closed-ended responses were frequently broken out by **principal investigator** (PI) status and by **project allocation**. Respondents were categorized according to the following project allocations:

INCITE The Department of Energy's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program aims to accelerate scientific discoveries and technological innovations by awarding, on a competitive basis, time on supercomputers to researchers with large-scale, computationally intensive projects that address "grand challenges" in science and engineering;

DD The National Center for Computational Sciences' Director's Discretion (DD) program is designed to give new researchers an opportunity to carry out a program of scalability and productivity enhancements to their scientific codes;

ALCC The Advanced Scientific Computing Research (ASCR) Leadership Computing Challenge (ALCC) program is open to scientists from the research community in national laboratories, academia and industry, and allocates up to 30% of the computational resources at National Energy Research Scientific Computing Center (NERSC) and the Leadership Computing Facilities at Argonne and Oak Ridge for special situations of interest to the Department's energy mission, with an emphasis on high-risk, high-payoff simulations;

ECP The Exascale Computing Project (ECP) is focused on accelerating the delivery of a capable exascale computing ecosystem that delivers 50 times more computational science and data analytic application power than possible with DOE HPC systems such as Titan (ORNL) and Sequoia (LLNL). The ECP is a collaborative effort of two U.S. Department of Energy organizations – the Office of Science (DOE-SC) and the National Nuclear Security Administration (NNSA); and

Other Other programs include Vendor and General projects.

Finally, tables and figures will include one or more of the following data elements:

- *N* = Total number of respondents who answered the question
- *n* = Total number of respondents who answered the specific item in the question or who provided a specific response
- *M* = the arithmetic average of respondents' scores from 1 (*Very Dissatisfied)* to 5 (*Very Satisfied)*
- *SD* = Standard deviation (indicating average deviation from the mean)
- %Sat = percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very Satisfied*) on satisfaction scales
- %Imp = percentage of respondents indicating 4 (*Very Important*) or 5 (*Extremely Important*) on importance scales

Color coding has been used in the report tables as below:

- Cell values in **green** are the highest %Sat values in the column
- Cell values in **red** are the lowest %Sat values in the column

This **color coding has not been applied** in cases where ratings are too similar or are identical in the column, or in cases where only three items are presented in a table.

As noted above, **open-ended responses** were typically information provided by respondents who were dissatisfied with a service/resource (i.e., responded as *Dissatisfied* or *Very Dissatisfied* on the satisfaction scale); other questions were open-ended invitations for suggestions or future needs. All open-ended responses were examined using categorical content analysis with complete thoughts in responses as the unit of analysis (note that percentages of response categories may add up to more than 100% when respondents provided multiple complete thoughts in a response).[1] Complete thoughts were sorted into categories for the purposes of counting, comparisons, and other forms of analysis.

Some response content categories were derived *a priori* from survey questions or OLCF website categories (e.g. *Data Management*). Other categories were developed inductively through an iterative process of grouping and regrouping similar content units (e.g., *Containers* or *Training and Tutorials*). Subcategories were elaborated as new relevant concepts or useful distinctions were identified, and are organized within major categories of closely related concepts. Table 1 provides a summary of major categories and subcategories used to organize open-ended replies. These are used to the extent possible, with variations as needed to accommodate differences in the focus of specific questions and year-to-year differences in users' specific and technical responses.

Examples of the most prominent themes are provided in the *Findings*, and all open-ended responses are provided in one of Appendices D-F.

[1] Complete thoughts (CTs) were simply response text that could stand alone as a meaningful reply to survey questions. CTs were not limited to any specific grammatical unit and could vary from a single word, to a phrase, sentence fragment or complete sentence.

Table 1. *Major Categories and Subcategories Used to Organize Open-Ended Responses*

Access
Accounts
Allocations
Login-connect
Hardware Computing Resources
Performance Upgrade
Capacities
Architecture
GPU Resources
Stability/Reliability
Running Jobs
Containers
Workflow
Scheduling Policy
Queue Time
Wall/Run Time
Data Management
Data Retention/Purge Policy and Procedures
Data Storage
Data Transfer
File Systems
Software
Libraries & Updates
Compilers
Debugging Tools
Development Tools
Visualization
Testing Capabilities
User Support
Documentation
User Guides
Tutorials
Training
Tech Support
Website
Communication
Example Additional Categories
Satisfaction
Miscellaneous
Survey Suggestions
Project Management/Planning

Findings

Respondents

Over 80% of respondents were **affiliated with** either a university or a DOE/Laboratory/ Government facility (Figure 1).

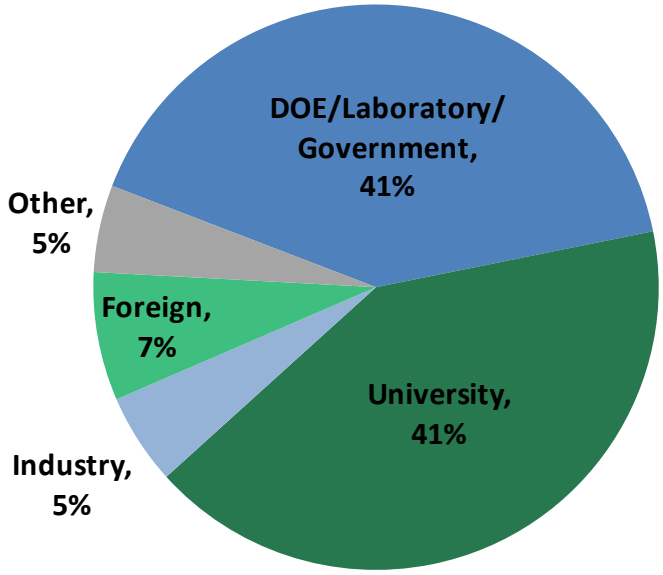

Figure 1. Respondent occupational affiliation (*N* = 422)

The distribution of OLCF users across **project allocations** is shown in Figure 2 and in greater detail in Table 2. The pool of survey respondents is generally representative of OLCF's distribution of users across various project lines. Note that the table categories are not exclusive (e.g., the INCITE category includes individuals assigned to INCITE, but who may also have been assigned to other projects). Note that 64% of respondents reported a single project allocation (i.e., assignment to only INCITE, only DD, only ALCC, or only ECP).

Table 2. *Project Allocations by OLCF Users and Survey Respondents*

	Survey Respondents (*N* = 422)		**OLCF Users (*N* =1230)**	
	Percentage	***n***	**Percentage**	***n***
INCITE	38%	161	34%	431
DD	64%	272	60%	754
ALCC	21%	88	18%	226
ECP	19%	80	15%	194
Other	4%	18	4%	55

Note: Percentages add to more than 100% as users are often affiliated with multiple projects.

The proportions of OLCF users and of 2018 survey respondents with PI status on at least one project are displayed in Figure 3. The survey respondent pool somewhat over-represents PIs. Throughout this report, tables separately report findings from respondents with PI status from those without PI status.

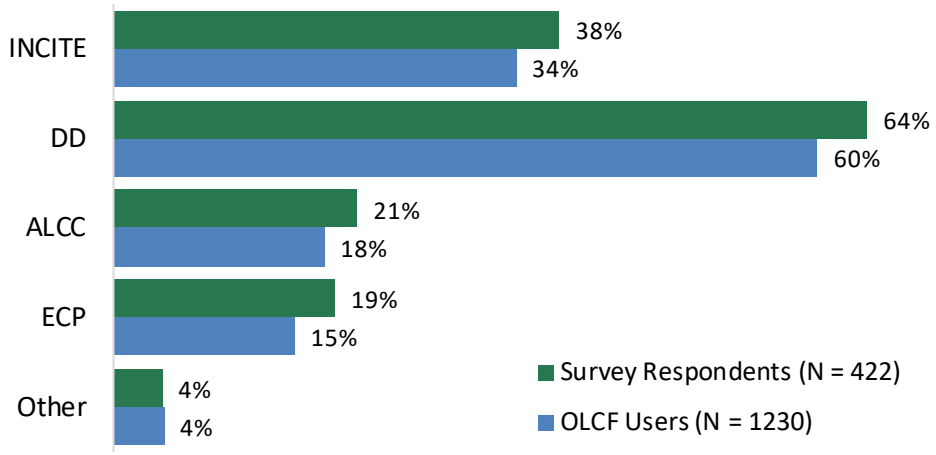

Figure 2. Project allocations for OLCF Users (*N* = 1230) and for Respondents (*N* = 422)

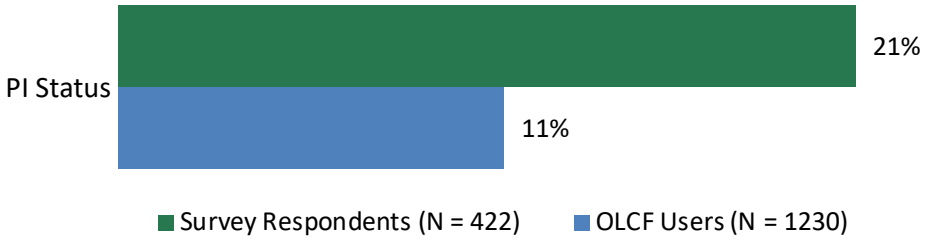

Figure 3. PI Status for OLCF Users (*N* = 1,230) and for Respondents (*N* = 422)

Resource Utilization

Overall experience using the OLCF was relatively evenly distributed across years of use. The largest proportion of respondents (more than one-half) had used the OLCF for more than 2 years (Figure 4).

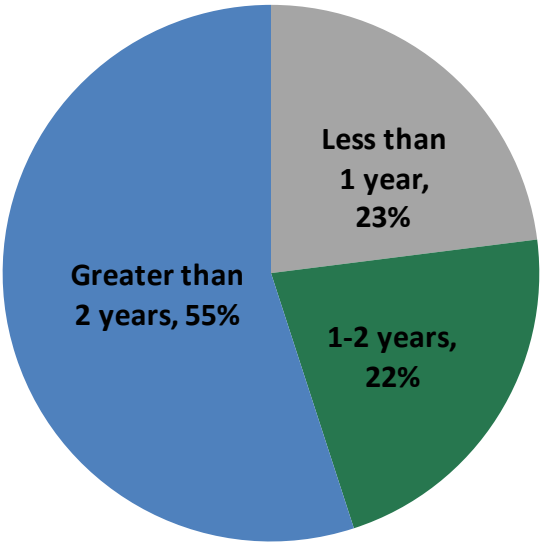

Figure 4. Experience using the OLCF (*N* = 422)

More specifically, respondents were asked to indicate **which OLCF HPC resources they utilized** during the 2018 calendar year. For all categories, the largest proportions of respondents indicated using Titan and Eos was utilized by the smallest proportion (Table 3).

Proportions of respondents **utilizing OLCF support services** during 2018 are presented in Table 4. The largest number of respondents indicated using the User Assistance Center while the smallest proportion utilized Data Analysis and Visualization Support Services.

The sections below report respondent **satisfaction ratings for OLCF resources/services** in four main categories (Overall Satisfaction, Computing Resources, Data Resources, and Support Services) and their subcategories.

Table 3. *HPC Resources Used by PI status, Project Allocation and Overall Totals*

	PI Status		INCITE		DD		ALCC		ECP		Other		Total	
	n	% Users	*n*	% Users	*n*	% Users	*n*	% Users	*n*	% Users	*n*	% Users	*n*	% Users
Titan	67	77%	126	78%	192	71%	73	83%	61	78%	12	67%	300	72%
Eos	21	24%	39	24%	59	22%	13	15%	6	8%	1	6%	87	21%
Rhea	24	28%	45	28%	71	26%	14	16%	10	13%	3	17%	96	23%
Data Transfer Nodes	40	46%	64	40%	120	44%	32	37%	26	34%	7	39%	159	38%
HPSS	23	26%	44	28%	66	24%	23	26%	13	17%	4	22%	97	23%
Lustre/Spider	58	67%	118	74%	176	65%	60	68%	50	64%	10	56%	270	65%

Note: Users add up to more than 100% because some used more than one system.

Table 4. *Support Services Used by PI Status, Project Allocation and Overall Totals*

	PI Status		INCITE		DD		ALCC		ECP		Other		Total	
	n	% Users	*n*	% Users	*n*	% Users	*n*	% Users	*n*	% Users	*n*	% Users	*n*	% Users
User Assistance Center	59	69%	83	54%	169	64%	44	52%	53	72%	14	78%	233	58%
Account Services	30	35%	36	24%	59	22%	23	27%	26	36%	6	35%	91	23%
INCITE Scientific Computing Liaison	18	21%	68	44%	48	18%	19	22%	11	15%	1	6%	84	21%
Data Analysis and Visualization Support Services	6	7%	9	6%	9	4%	4	5%	2	3%	2	12%	15	4%

Note: Users add up to more than 100% because some used more than one system.

Overall Satisfaction

Users were asked to rate their **"overall" satisfaction with the OLCF**, and then with OLCF Compute Resources, Data Resources, and Support Services. In these responses, individuals were not asked to consider the specific resources/services in a category, but rather report their general sense of satisfaction with the category. First, most respondents reported being *very satisfied* in this overall sense for all categories of resources/services (Figure 5).

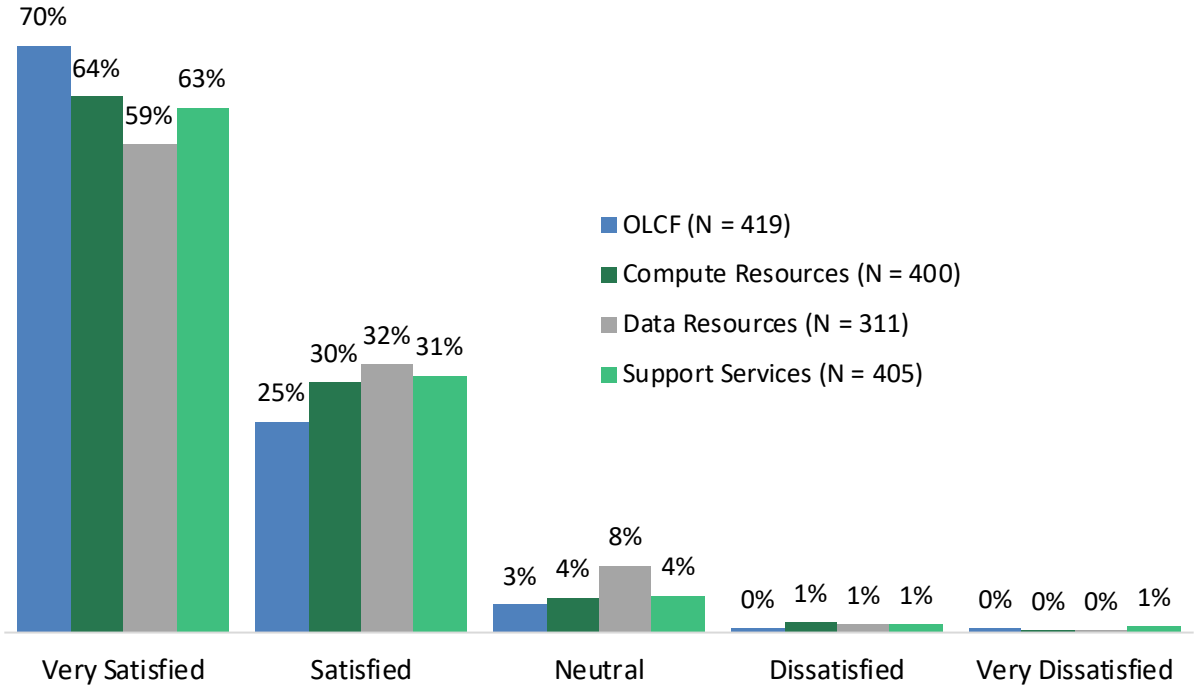

Figure 5. "Overall" Satisfaction with OLCF and its major resources/services (maximum *N* = 419)

Table 5 summarizes descriptive statistics for these overall satisfaction ratings for all respondents and broken down by PI status, while Table 6 reports satisfaction statistics across project allocations. The tables also include ratings of **specific compute resources** (i.e., Titan, Eos, and Rhea), data resources (i.e., Data Transfer Nodes, HPSS, and Lustre/Spider), and **support services** (i.e., User Assistance, Account Services, Data Analysis and Visualization Support Services, INCITE Liaison, Communication, Training, and aspects of the Website). Across 17 items and all categories of respondents, the tables show that:

- *%Sat* ranged from 87% to 100%,
- Means ranged from 4.2 to 4.7, and
- SDs ranged from 0.49 to 0.98.

Table 5. *Overall Satisfaction with OLCF and Its Major Resources/Services by PI Status and Totals*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
OLCF	**87**	**4.8**	**0.50**	**97%**	**332**	**4.6**	**0.64**	**95%**	**419**	**4.6**	**0.62**	**96%**
Compute Resources	**83**	**4.6**	**0.68**	**94%**	**317**	**4.6**	**0.66**	**94%**	**400**	**4.6**	**0.66**	**94%**
Titan	67	4.6	0.58	96%	231	4.4	0.64	95%	298	4.5	0.63	95%
Eos	21	4.5	0.51	100%	65	4.6	0.53	98%	86	4.5	0.52	99%
Rhea	24	4.7	0.46	100%	72	4.6	0.49	100%	96	4.6	0.49	100%
Data Resources	**67**	**4.6**	**0.62**	**96%**	**244**	**4.4**	**0.73**	**89%**	**311**	**4.5**	**0.71**	**91%**
Data Transfer Nodes	40	4.2	0.83	85%	119	4.2	1.02	88%	159	4.2	0.98	87%
HPSS	23	4.5	0.67	91%	74	4.5	0.69	96%	97	4.5	0.68	95%
Lustre/Spider	57	4.4	0.71	91%	210	4.4	0.75	92%	267	4.4	0.74	92%
Support Services	**87**	**4.7**	**0.58**	**97%**	**318**	**4.5**	**0.72**	**93%**	**405**	**4.5**	**0.70**	**94%**
User Assistance	57	4.7	0.51	98%	172	4.6	0.56	97%	229	4.6	0.54	97%
Account Services	30	4.6	0.49	100%	60	4.7	0.66	93%	90	4.7	0.60	96%
Data Analysis and Visualization Support Services	6	4.7	0.52	100%	7	4.3	0.49	100%	13	4.5	0.52	100%
INCITE Liaison	15	4.7	0.59	93%	48	4.8	0.53	96%	63	4.7	0.54	95%
Communication	81	4.5	0.55	98%	302	4.4	0.72	89%	383	4.4	0.69	91%
Training and Technical Reference Documentation	81	4.4	0.66	90%	288	4.5	0.63	95%	369	4.4	0.64	94%
OLCF Website	85	4.3	0.66	89%	297	4.4	0.61	94%	382	4.4	0.62	93%
Min	**6**	**4.2**	**0.46**	**85%**	**7**	**4.2**	**0.49**	**88%**	**13**	**4.2**	**0.49**	**87%**
Max	**87**	**4.8**	**0.83**	**100%**	**332**	**4.8**	**1.02**	**100%**	**419**	**4.7**	**0.98**	**100%**

Table 6. *Overall Satisfaction with OLCF and Its Major Resources/Services by Project Allocation*

	INCITE				DD				ALCC				ECP				Other*			
	N	*M*	*SD*	*%Sat*	*N*	*M*	*SD*	*%Sat*	*N*	*M*	*SD*	*%Sat*	*N*	*M*	*SD*	*%Sat*	*N*	*M*	*SD*	*%Sat*
OLCF	**159**	**4.6**	**0.59**	**94%**	**271**	**4.7**	**0.57**	**97%**	**88**	**4.6**	**0.67**	**95%**	**80**	**4.5**	**0.81**	**91%**	**18**	**4.6**	**0.50**	**100%**
Compute Resources	**155**	**4.5**	**0.64**	**95%**	**256**	**4.6**	**0.63**	**94%**	**87**	**4.4**	**0.69**	**94%**	**76**	**4.4**	**0.71**	**92%**	**17**	**4.6**	**0.51**	**100%**
Titan	126	4.4	0.59	94%	190	4.5	0.59	95%	73	4.3	0.68	95%	61	4.3	0.66	93%	12	4.5	0.52	100%
Eos	39	4.7	0.47	100%	58	4.5	0.54	98%	13	4.5	0.52	100%	6	4.3	0.52	100%	1	5.0	--	100%
Rhea	45	4.6	0.50	100%	71	4.6	0.49	100%	14	4.5	0.52	100%	10	4.6	0.52	100%	3	5.0	0.00	100%
Data Resources	**127**	**4.4**	**0.74**	**88%**	**209**	**4.5**	**0.68**	**91%**	**67**	**4.3**	**0.75**	**91%**	**55**	**4.3**	**0.84**	**80%**	**14**	**4.6**	**0.51**	**100%**
Data Transfer Nodes	64	4.1	1.21	86%	120	4.2	0.86	88%	32	4.1	0.95	91%	26	4.2	0.83	92%	7	3.7	1.38	71%
HPSS	44	4.5	0.63	93%	66	4.5	0.61	94%	23	4.3	0.92	91%	13	4.5	0.52	100%	4	4.3	0.50	100%
Lustre/Spider	116	4.3	0.72	89%	175	4.4	0.70	92%	58	4.1	0.82	88%	49	4.1	0.91	84%	10	4.4	0.52	100%
Support Services	**155**	**4.5**	**0.69**	**94%**	**263**	**4.6**	**0.69**	**95%**	**88**	**4.5**	**0.69**	**94%**	**76**	**4.4**	**0.79**	**91%**	**18**	**4.6**	**0.50**	**100%**
User Assistance	82	4.6	0.54	98%	166	4.6	0.54	97%	43	4.6	0.58	95%	52	4.7	0.58	94%	13	4.7	0.48	100%
Account Services	36	4.6	0.55	97%	59	4.6	0.58	95%	22	4.7	0.48	100%	26	4.7	0.68	96%	6	4.8	0.41	100%
Data Analysis & Visualization Support Services	8	4.4	0.52	100%	8	4.5	0.53	100%	4	4.5	0.58	100%	2	5.0	0.00	100%	2	5.0	0.00	100%
INCITE Liaison	50	4.8	0.55	94%	38	4.7	0.50	97%	19	4.4	0.76	84%	9	4.4	0.88	78%	1	5.0	--	100%
Communication	145	4.4	0.72	89%	254	4.4	0.67	91%	82	4.3	0.72	91%	71	4.5	0.63	93%	16	4.6	0.51	100%
Training and Technical Reference Documentation	141	4.4	0.65	92%	244	4.4	0.65	93%	81	4.3	0.69	94%	71	4.4	0.74	92%	16	4.5	0.52	100%
OLCF Website	144	4.4	0.62	92%	253	4.4	0.61	93%	83	4.2	0.67	89%	71	4.3	0.64	90%	16	4.4	0.51	100%
Min	**8**	**4.1**	**0.47**	**86%**	**8**	**4.2**	**0.49**	**88%**	**4**	**4.1**	**0.48**	**84%**	**2**	**4.1**	**0.00**	**78%**	**1**	**3.7**	**0.00**	**71%**
Max	**159**	**4.8**	**1.21**	**100%**	**271**	**4.7**	**0.86**	**100%**	**88**	**4.7**	**0.95**	**100%**	**80**	**5.0**	**0.91**	**100%**	**18**	**5.0**	**1.38**	**100%**

*18 respondents are allocated to "Other" projects, and some questions received too few responses for SD or %Sat to be meaningful.

Only 13 respondents reported reasons for **dissatisfaction with the OLCF overall or with its major resources/services.** Individuals cited problems with *support, training, and documentation* (*n* = 4), *software* (*n* = 4), *performance (n* = 4), and *usefulness/user-friendliness* (*n* = 3):

Support, Training, and Documentation	*"My group's allocation at OLCF disappeared August 15. My understanding was that new ECP allocations would start on OLCF systems as of Oct 1. When that didn't happen, it took weeks to resolve. . . Although my student has been told that he would get access to the P9+V100 hardware on Summit, it now appears that such access won't be granted until Jan 1."* *"The OLCF user website information has seemed to be outdated or simply not accurate for multiple uses recently - mainly in the area of how to set up and use the suggested debugging systems."* *"Training events are very infrequent and tend to very high level and not very useful at times. Documentation on website are helpful and expanding them would be of benefit to large user community."* *"We have found getting DL tools tensorflow, e.g., support to be less than strong, and peripheral tools like nccl update requests to be done grudgingly or not at all. We finally found solutions via an unpublicized website, which now appears to be behind a firewall."*
Software	*"Compared with ALCF, the OLCF resources usually have older versions of the visualization and postprocessing software, the software module system is also less user-friendly."* *"We have been using Titan as an ECP testbed for performance tools. The old kernel on Titan means that we can't test aspects of our tools such as collecting call stack samples in the kernel. What works on newer kernels doesn't work on Titan."* *"Several users have been asking whether OLCF could implement a JupyterHub service similar to NERSC's service (https://jupyter-dev.nersc.gov). In my experience, having interactive access to the simulations' data, via Jupyter, can oftentimes be key to the workflow of scientific discovery. Would it be possible to use a similar solution to the one implemented by NERSC?"* *"We have found getting DL tools tensorflow, e.g., support to be less than strong, and peripheral tools like nccl update requests to be done grudgingly or not at all."*

Performance	*"File system performance on Titan for compiling is poor (even in nfs mounted directories). Builds take several times longer than they would on a PC or at other DOE HPC resources. This makes the development/test/debug cycle much less efficient."* *"HPSS interface is at times "clunky" when trying to archive large molecular dynamics simulations."* *"Our experience has been that Titan performs rather poorly for jobs of this size, with frequent (~20% or more) job failures due to hardware issues such as node failures (often from the GPU) or the filesystem I/O hanging. Smaller jobs using our code perform much better, and significantly better performance is found on Summit. But for a platform nominal designed for, and pushed as, a resource for large-scale "capability" computing it has been underwhelming."* *"Summitdev isn't very usable. We regularly experience hangs for minutes at a time while trying to work on the system."*
Usefulness and User-Friendliness	*"Compared with ALCF, the OLCF resources usually have older versions of the visualization and postprocessing software, the software module system is also less user-friendly."* *"HPSS not userfriendly."* *"The OLCF resources are not suitable for my research."*

One user expressed satisfaction in the comments they provided, while other users were unhappy with token access and Early Science proposal processes. All open-ended responses are provided in Appendix E: User Dissatisfaction Explanations.

Finally, respondents described what they perceived to be **"the best qualities of OLCF."** Thematic analysis of user responses identified *computing power/performance* (56%) and *user tech support/staff* (36%) as the most valued qualities of the OLCF (Table 7; see Appendix D: Best Qualities of the OLCF for all responses by category; *N* = 349). Many illustrative examples praised multiple elements of OLCF:

> *"Powerful and reliable system. Consistently capable of providing a cutting edge service over the years. I have had the fortune of being able to use this system for the past 7-8 years, and published nearly 40 papers as a result, almost all of them including research that would have been impossible to carry out without access to OLCF facilities."*
>
> *"The OLCF and its staff are a world-class resource, critical to the advancement of science in the USA. The power of the computing systems and the skills of the staff are the facilities' best qualities."*

"OLCF is a one-of-a-kind facility. There are not comparable centers that provide the same scale of computational resources in conjunction with supporting infrastructure needed to support productive science."

"The total power of the resources is unmatched. Software availability is good. Overall it is easy for an experienced HPC user to get started and be productive."

"Machine capability and a flexible and responsive staff that are willing to consider changes needed to meet the needs of our domain science. The spirit of collaboration that exists between the facility and users thereof."

"Resources provided through OLCF are vital to perform work done within our research projects. It allows to run large scale simulations that could not be completed almost at any other systems in the world. The data analysis and visualization systems provide a convenient way to perform processing of the results."

While appreciation for the power and performance of the facilities may not come as a surprise, the relatively high frequency of positive references to OLCF User Support is perhaps more unexpected. These responses were re-examined, excluding individuals that mentioned **only** computing performance as the best quality (removing 72 responses). The relative frequency of comments reported by this group (*N* = 277), excluding references to computing power/performance is shown in the last column of Table 7. *Tech support/staff* is prominent as the perceived best OLCF quality when the responses are examined in this way, but there is significant spread across categories and variety in responses.

For example:

Tech Support/Staff	*"The staff are extremely helpful and responsive."* *"The best support team which is always ready to help OLCF users."* *"The people! User assistance and those who manage the projects are the bomb."*
Computing Performance	*"The compute capabilities are the most important quality. We use that for both data generation and analysis."* *"The OLCF represents something that does not exist elsewhere."* *"Providing continuing access to leadership class computing."*

Table 7. *Best Qualities of OLCF (ordered by % of all respondents, high to low)*

	All Responses (*N* = 349)	**Responses Excluding Computing Performance (*N* = 277)**
Computing power/HPC resources	56%	n/a
User support/staff	36%	45%
System documentation & website information	8%	10%
Stability/reliability	8%	10%
Tools (software, libraries, viz, & analysis)	7%	9%
Queue time	6%	8%
Supports scientific research	6%	8%
Availability/uptime	5%	6%
Resource management/infrastructure/ maintenance	5%	6%
GPU resources	4%	5%
Training/tutorials	4%	5%
Data storage/disk space	3%	4%
Accessibility	3%	4%
Overall satisfaction	3%	3%
Ease of use	3%	3%
Communication	2%	3%
Customer-focused	2%	3%
Shared filesystem/cross-system use cases	1%	1%
Data transfer	1%	2%
Miscellaneous/Other	5%	7%

Note: Users add up to more than 100% because many provided more than one theme in their response.

Compute and Data Resources

Respondents provided satisfaction ratings for several specific computing and data resources features:

- notice given prior to scheduled maintenance
- project disk space
- ease of transferring data to/from the OLCF
- bandwidth

Table 8 reports satisfaction for these features by PI status and overall, and Table 9 reports ratings by project allocation. The highest satisfaction rating (all respondents) was for *notice given prior to scheduled maintenance* (94% satisfied), and the lowest overall mean rating was for *ease of transferring data to/from the OLCF* (85% satisfied).

Of the 14 respondents that reported reasons for **dissatisfaction with one or more aspects of the OLCF compute and data resources,** the majority of complaints had to do with *data transfer* (*n* = 5), *disk space* (*n* = 4), and *access* (*n* = 3). All open-ended responses are provided in Appendix E: User Dissatisfaction Explanations.

Data Transfer	*"All of the data transfer options seem to be for internal ORNL use only. As an external collaborator I can only use scp to get simulation output onto my local machine. That wouldn't be so bad if OLCF allowed me to piggyback on an open ssh connection, but they have disabled that feature, so I have to actually reenter my PIN on every single transfer."* *"Data transfer pushes from the outside world to ORNL systems is hindered by the overly restrictive MFA and ssh/scp policies."* *"I completely understand why you don't, but it would be nice to support scp without using RSA."* *"I encountered problems transferring data, due to bandwidth problems at my institution. I then lost data as I did not transfer on time."* *"sFTP would be very useful for file transfer (maybe it is possible to use FTP and I am not aware of this?). I personally find rsync etc. quite awkward."*
Disk Space	*"I find the space limits in both home and project to be very limited (the main filesystem, HPSS is good)."* *"Our project's /ccs/proj/ space was already full from other projects/users."* *"Persistent project space that is large is desired."* *"The size is still too small if considering that we are in the big-data era."*

Access	*"As of my agreement as a foreign national and for security reasons, I only connect to OLCF from the same workstation on Campus. This limits the time I can use the service (our end). . . . Transferring rate is ok (scp) for my current coarse-grid tests, but is probably terrible for something serious. There was these couple of days in which everything was hanging."* *"Requirement to enter password for every ssh/rsync is cumbersome."* *"sFTP would be very useful for file transfer (maybe it is possible to use FTP and I am not aware of this?). I personally find rsync etc. quite awkward."*

Table 8. *Satisfaction Ratings for Features of the OLCF HPC Compute and Data Resources by PI Status and Overall Totals*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Notice given prior to scheduled maintenance	86	4.7	0.52	97%	320	4.5	0.70	93%	406	4.6	0.67	94%
Project disk space	87	4.5	0.76	92%	313	4.4	0.80	88%	400	4.4	0.79	89%
Ease of transferring data to/from the OLCF	84	4.4	0.81	88%	308	4.3	0.85	84%	392	4.4	0.85	85%
Bandwidth offered by the OLCF	82	4.6	0.74	91%	302	4.4	0.76	90%	384	4.5	0.75	91%

Table 9. *Satisfaction Ratings for Features of the OLCF HPC Compute and Data Resources by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Notice given prior to scheduled maintenance	155	4.6	0.64	95%	266	4.6	0.64	94%	87	4.4	0.74	93%	75	4.5	0.74	92%
Project disk space	155	4.3	0.84	86%	259	4.4	0.79	89%	86	4.3	0.80	86%	77	4.2	0.89	81%
Ease of transferring data to/from the OLCF	156	4.3	0.88	85%	254	4.3	0.83	86%	85	4.2	0.89	81%	72	4.3	0.87	85%
Bandwidth offered by the OLCF	152	4.4	0.79	88%	249	4.5	0.71	92%	84	4.4	0.70	94%	72	4.3	0.87	88%

In addition, respondents were asked to indicate their opinion regarding the **performance of computing and data resources compared to the previous year.** Overall, 24.3% reported *improvements*, just 1.8% perceived *decreases in performance*, and 73.8% reported *no change* (Figure 6). Some differences in these perceptions were observed across years of using the OLCF. Less experienced users (those with only 1-2 years' experience) were slightly less likely to report seeing a change over the last year, while more experienced users (more than 2 years' experience) were more likely to report an improvement in performance.

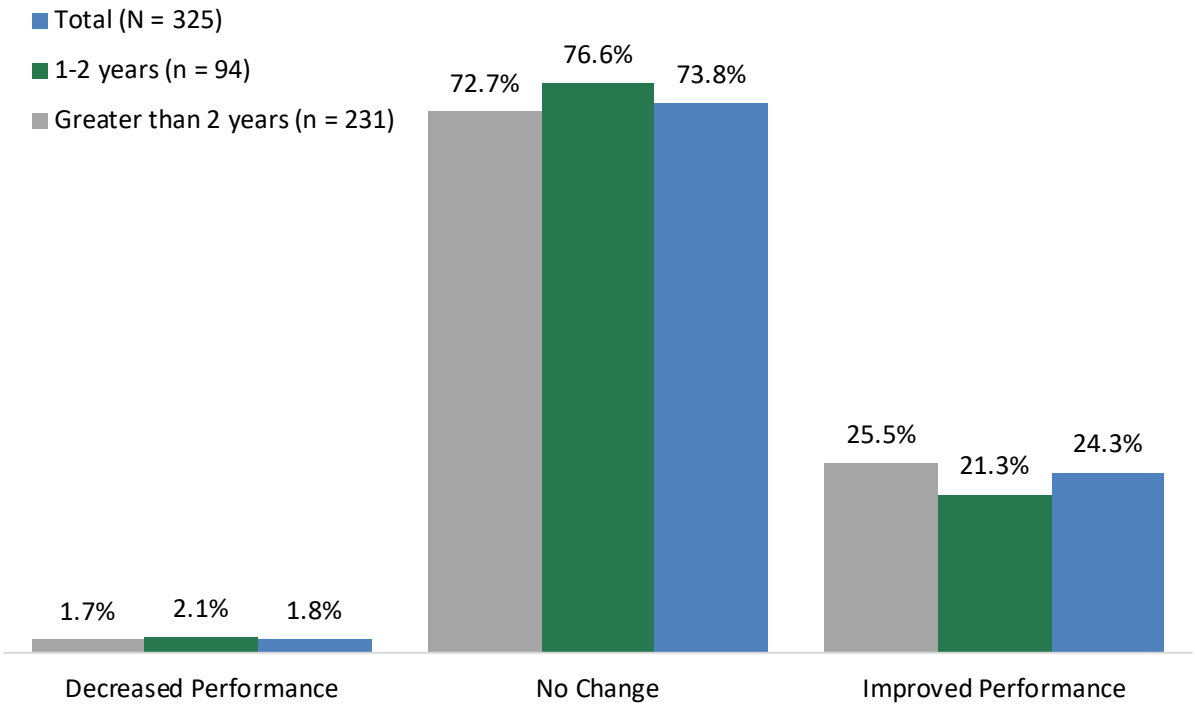

Figure 6. Perceived changes from 2017 computing/data resources performance by years using OLCF (*N* = 325)

Among the 6 respondents that provided comments **describing decreased performance,** *reliability, failures, and outages* was the most prominent theme with half of the comments reporting increased node failures or other forms of instability:

> *"I have observed more outages in Lustre this year than the last."*
>
> *"Libraries, modules, and container loading/support issues."*
>
> *"More nodes failing and intermittent failures. Also much slower filesystem access."*

All open-ended responses are provided in Appendix E: User Dissatisfaction Explanations.

Compute Resources

Titan

Titan was used by 72% of respondents during the 2018 calendar year (*N* = 418). Titan users were asked to provide satisfaction ratings for multiple aspects of the system, and descriptive statistics of these ratings are shown in Table 10, which also reports satisfaction by PI status. 95% of all respondents were either *satisfied* or *very satisfied* with the system. Table 11 summarizes these satisfaction statistics by project allocation.

The *job success rate, frequency of scheduled outages,* and the *programming environment* were the highest rated specific aspects, and lowest rated aspects were *debugging and performance tools* and *software/libraries.*

There were 27 Titan users who reported at least one reason for **dissatisfaction with Titan**. Over one-third of these users (*n* = 10) were unhappy with *outdated systems, tools, and libraries* that had an impact on the work they could conduct on the system. For example:

> *"A major problem with Titan is the very old OS which is making it difficult (or impossible) to run software that relies on newer glibc functionality. An example of this is the latest IntelMPI release which requires a newer glibc."*

> *"The kernel was quite old that makes it impossible to use new technologies developed in the lab."*

The next largest group of dissatisfied users (*n* = 6) were unhappy with the *performance* of the system or indicated a need for *performance upgrades*, with Lustre performance a common theme within these concerns. For example:

> *"I frequently have problems with Lustre performance, in particular the project shared and world shared areas."*

All open-ended responses are provided in Appendix E: User Dissatisfaction Explanations.

Table 10. *Satisfaction Ratings of Titan by PI Status and Overall*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Batch wait time	65	4.3	0.71	89%	222	4.2	0.78	86%	287	4.2	0.77	87%
Batch queue structure	64	4.3	0.69	91%	224	4.3	0.73	90%	288	4.3	0.72	90%
Job success rate	65	4.5	0.71	95%	222	4.4	0.74	91%	287	4.4	0.73	92%
Frequency of scheduled outages	67	4.4	0.61	94%	215	4.3	0.69	90%	282	4.4	0.67	91%
Frequency of unanticipated outages	66	4.4	0.74	88%	215	4.3	0.70	89%	281	4.3	0.71	89%
Debugging and performance tools	51	4.4	0.67	94%	185	4.2	0.80	82%	236	4.2	0.78	84%
Software/libraries	62	4.3	0.85	84%	220	4.2	0.79	86%	282	4.2	0.80	85%
Programming environment	58	4.5	0.57	97%	220	4.3	0.75	90%	278	4.3	0.72	91%
Scratch configuration	61	4.5	0.59	95%	216	4.3	0.79	87%	277	4.3	0.76	88%
I/O performance	65	4.5	0.61	94%	208	4.2	0.84	85%	273	4.3	0.79	87%
Overall satisfaction with Titan	67	4.6	0.58	96%	231	4.4	0.64	95%	298	4.5	0.63	95%

Table 11. *Satisfaction Ratings of Titan by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Batch wait time	121	4.2	0.72	86%	183	4.3	0.76	88%	70	4.2	0.80	87%	61	4.3	0.72	89%
Batch queue structure	122	4.3	0.66	92%	184	4.3	0.72	91%	70	4.2	0.83	87%	61	4.3	0.57	95%
Job success rate	124	4.3	0.82	88%	183	4.4	0.69	93%	69	4.2	0.79	87%	61	4.4	0.68	89%
Frequency of scheduled outages	124	4.3	0.68	87%	180	4.4	0.60	94%	70	4.2	0.77	87%	58	4.3	0.60	93%
Frequency of (unanticipated) unscheduled outages	122	4.2	0.68	86%	179	4.4	0.67	91%	69	4.2	0.76	88%	59	4.3	0.71	86%
Debugging and performance tools	104	4.1	0.80	79%	145	4.2	0.72	88%	58	4.1	0.79	81%	52	4.1	0.86	81%
Software/libraries	121	4.2	0.76	84%	176	4.2	0.82	86%	71	4.0	0.90	75%	58	4.1	0.84	81%
Programming environment	121	4.2	0.71	90%	172	4.4	0.71	93%	71	4.1	0.85	85%	59	4.2	0.85	83%
Scratch configuration	118	4.2	0.76	86%	174	4.3	0.69	90%	70	4.2	0.84	86%	59	4.2	0.78	85%
I/O performance	117	4.1	0.84	81%	175	4.3	0.76	89%	68	4.1	0.89	84%	55	3.9	1.01	75%
Overall satisfaction with Titan	126	4.4	0.59	94%	190	4.5	0.59	95%	73	4.3	0.68	95%	61	4.3	0.66	93%

Eos

Eos was used by 21% of respondents during the 2018 calendar year (*N* = 417). Eos users were asked to provide satisfaction ratings for multiple aspects of the system, and descriptive statistics of these ratings are shown in Table 12, which also reports satisfaction statistics by PI status. Almost all (99%) respondents were either *satisfied* or *very satisfied* with the system. Table 13 summarizes these satisfaction statistics by project allocation.

The *frequency of scheduled outages* was the highest rated specific aspect, and the lowest rated features were the *batch wait time* and *software/libraries.*

Six Eos users reported **reasons for dissatisfaction**. Two of these users were dissatisfied with *performance,* two with the *purge policy,* and one each with the *queue wait time* and *difficulty using the system:*

> *"Filesystem performance remains an issue."*
>
> *"I was unable to run an important part of my analysis on Eos due to node memory limitations. Basically a significant portion of each node's memory was occupied (by necessity) by a background mesh during grid adaptation, leaving a very small amount of memory remaining for other computations. This memory limit caused many of my grid adaptation sequence jobs to fail... I could not run the grid adaptation jobs on Rhea without significant re-writing of the code due to differences in the environments between Rhea and Eos."*
>
> *"I have found the purge policy on scratch to be quite restrictive. The short window before deletion can be problematic sometimes."*
>
> *"With current scratch purge policy, I'm often surprised to log in to find my files (but not directories) gone. But I expect that I am mis-estimating how long since files were created. HyperThreading controls (j1 and j2 flags to aprun, plus core numbering) seem broken. Definitely do not behave as expected. Documentation of core numbering in both cases (j1 and j2) is confusing, assuming those flags work."*
>
> *"The lack of machine learning libraries (e.g. tensorflow) is a major shortcoming. Batch wait time is longer than I like and apparently not scaled to the time remaining in my allocation, so I cannot get through my allocation over the course of the year due to the wait time in the queue."*
>
> *"Sometimes, missing a detail in your simulations makes it blow, for instance, and then a lot of hours passed while in the queue. But that is just user incompetence for not trying a smaller run with fewer nodes. Unfortunately, some cases required at least 2 nodes due to ram issues and had to wait anyway (getting 1 node was decently fast). Then you have to get smarter and play with the queue (showq). In my last few batches I was getting jobs done a bit better and I was given higher priority than other users, so it is appreciated."*

Table 12. *Satisfaction Ratings of Eos by PI Status and Overall*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Batch wait time	21	4.1	1.01	81%	64	4.3	0.73	84%	85	4.3	0.81	84%
Batch queue structure	21	4.4	0.50	100%	64	4.4	0.66	91%	85	4.4	0.62	93%
Job success rate	21	4.4	0.75	95%	64	4.5	0.69	92%	85	4.5	0.70	93%
Frequency of scheduled outages	21	4.5	0.51	100%	62	4.5	0.53	98%	83	4.5	0.53	99%
Frequency of (unanticipated) unscheduled outages	21	4.4	0.80	90%	60	4.4	0.62	93%	81	4.4	0.67	93%
Debugging and performance tools	15	4.6	0.51	100%	47	4.3	0.67	87%	62	4.3	0.65	90%
Software/libraries	19	4.3	0.81	89%	60	4.2	0.70	85%	79	4.2	0.72	86%
Programming environment	17	4.5	0.51	100%	62	4.3	0.70	90%	79	4.4	0.67	92%
Scratch configuration	20	4.4	0.75	85%	59	4.3	0.75	90%	79	4.3	0.74	89%
I/O performance	21	4.4	0.68	90%	62	4.4	0.68	92%	83	4.4	0.68	92%
Overall satisfaction with Eos	21	4.5	0.51	100%	65	4.6	0.53	98%	86	4.5	0.52	99%

Table 13. *Satisfaction Ratings of Eos by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Batch wait time	39	4.3	0.74	85%	58	4.2	0.86	79%	12	4.7	0.49	100%	6	4.5	0.84	83%
Batch queue structure	39	4.5	0.56	97%	58	4.3	0.63	91%	12	4.4	0.67	92%	6	4.7	0.52	100%
Job success rate	39	4.7	0.51	97%	58	4.3	0.76	90%	12	4.7	0.49	100%	6	4.7	0.52	100%
Frequency of scheduled outages	37	4.6	0.49	100%	57	4.4	0.54	98%	13	4.6	0.51	100%	6	4.5	0.55	100%
Frequency of (unanticipated) unscheduled outages	37	4.6	0.50	100%	55	4.3	0.72	89%	12	4.5	0.52	100%	6	4.3	0.52	100%
Debugging and performance tools	28	4.5	0.64	93%	40	4.2	0.64	88%	9	4.2	0.67	89%	4	4.3	0.50	100%
Software/libraries	36	4.3	0.74	83%	53	4.2	0.72	85%	12	4.3	0.75	83%	5	4.0	0.00	100%
Programming environment	35	4.4	0.77	89%	54	4.3	0.61	93%	11	4.2	0.98	82%	5	4.2	0.45	100%
Scratch configuration	34	4.3	0.86	85%	53	4.3	0.66	89%	12	4.0	1.13	75%	5	4.4	0.55	100%
I/O performance	37	4.4	0.60	95%	57	4.2	0.71	88%	12	4.6	0.51	100%	5	3.8	1.10	80%
Overall satisfaction with Eos	39	4.7	0.47	100%	58	4.5	0.54	98%	13	4.5	0.52	100%	6	4.3	0.52	100%

Rhea

Rhea was used by 23% of respondents during the 2018 calendar year (*N* = 417). Rhea users were asked to provide satisfaction ratings for multiple aspects of the system, and descriptive statistics of these ratings are shown in Table 14, which also reports satisfaction statistics by PI status. All respondents (100%) were either *satisfied* or *very satisfied* with the system. Table 15 summarizes these satisfaction statistics by project allocation.

The *batch queue structure* was the highest rated specific aspect, and the lowest rated features were *software/libraries* and the *frequency of unscheduled outages.*

Only 4 Rhea users expressed **reasons for dissatisfaction**, all related to performance and maintenance issues:

> *"It's difficult to maintain anaconda/python distributions (which I need only because analysis modules I use are unavailable in the machine modules), and I've also had difficulties using jupyter notebooks to perform analysis on Rhea. This is a hangup because our model generates a lot of data, and it's an analysis bottleneck for me to transfer that data elsewhere to perform analysis that (I think) should be doable on Rhea."*
>
> *"Libraries support/loading issues."*
>
> *"Sometimes the shared file system is slow."*
>
> *"The GLIBC on rhea is out of date, several of the modules and libraries are also out of date. This necessitates a significant amount of effort on re-compiling software and initializing local environments to compensate."*

Table 14. *Satisfaction Ratings of Rhea by PI Status and Overall Totals*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Batch wait time	23	4.7	0.57	96%	68	4.5	0.61	94%	91	4.6	0.60	95%
Batch queue structure	22	4.7	0.48	100%	68	4.6	0.61	94%	90	4.6	0.58	96%
Job success rate	23	4.7	0.47	100%	70	4.6	0.63	93%	93	4.6	0.59	95%
Frequency of scheduled outages	24	4.5	0.66	92%	69	4.4	0.69	88%	93	4.4	0.68	89%
Frequency of (unanticipated) unscheduled outages	24	4.5	0.59	96%	67	4.3	0.75	84%	91	4.4	0.71	87%
Debugging and performance tools	17	4.6	0.51	100%	51	4.3	0.71	86%	68	4.4	0.67	90%
Software/libraries	22	4.2	0.85	82%	65	4.3	0.77	88%	87	4.3	0.79	86%
Programming environment	18	4.5	0.62	94%	67	4.4	0.70	91%	85	4.4	0.68	92%
Scratch configuration	22	4.6	0.59	95%	60	4.5	0.62	93%	82	4.5	0.61	94%
I/O performance	24	4.6	0.58	96%	69	4.4	0.75	91%	93	4.4	0.71	92%
Overall satisfaction with Rhea	24	4.7	0.46	100%	72	4.6	0.49	100%	96	4.6	0.49	100%

Table 15. *Satisfaction Ratings of Rhea by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Batch wait time	42	4.6	0.54	98%	68	4.6	0.58	96%	12	4.5	0.67	92%	9	4.7	0.71	89%
Batch queue structure	42	4.6	0.53	98%	67	4.6	0.55	97%	12	4.4	0.67	92%	9	4.7	0.71	89%
Job success rate	44	4.6	0.53	98%	69	4.6	0.58	96%	13	4.5	0.66	92%	9	4.7	0.71	89%
Frequency of scheduled outages	44	4.4	0.69	89%	69	4.4	0.69	88%	14	4.6	0.65	93%	10	4.7	0.48	100%
Frequency of (unanticipated) unscheduled outages	44	4.3	0.71	86%	67	4.3	0.73	85%	14	4.7	0.47	100%	10	4.7	0.48	100%
Debugging and performance tools	29	4.3	0.72	86%	51	4.3	0.68	88%	11	4.5	0.69	91%	7	4.3	0.76	86%
Software/libraries	41	4.3	0.84	85%	65	4.2	0.77	86%	13	4.2	0.83	77%	9	4.1	1.05	78%
Programming environment	41	4.3	0.79	85%	63	4.4	0.71	90%	11	4.5	0.52	100%	9	4.6	0.73	89%
Scratch configuration	39	4.5	0.64	92%	60	4.5	0.62	93%	12	4.7	0.49	100%	8	4.5	0.76	88%
I/O performance	43	4.3	0.84	88%	70	4.4	0.63	93%	13	4.3	1.11	92%	9	3.9	1.27	78%
Overall satisfaction with Rhea	45	4.6	0.50	100%	71	4.6	0.49	100%	14	4.5	0.52	100%	10	4.6	0.52	100%

Data Resources

Data Transfer Nodes

DTNs were used by 38% of respondents during the 2018 calendar year (*N* = 416), and 87% were either *satisfied* or *very satisfied* with the DTNs. Satisfaction did not vary substantially by PI status or project allocation. Nine users **indicated they were dissatisfied** with DTNs, but did not provide comments explaining their ratings. However, three users who rated their satisfaction as *Neutral* or *Satisfied* provided the following comments:

> *"I think it works pretty well. I did not notice stalls or really bad transfer rates, just about ok."*
>
> *"It was somewhat difficult to determine how to maximize the bandwidth used by the DTNs. It doesn't seem like this needs to be so difficult."*
>
> *"The transfer was for Summit testing, so it was not production quality."*

HPSS

HPSS was used by 23% of respondents during the 2018 calendar year (*N* = 417). HPSS users were asked to provide satisfaction ratings for multiple aspects of the system, and descriptive statistics of these ratings are shown in Table 16, which also reports satisfaction statistics by PI status. 95% of respondents were overall either *satisfied* or *very satisfied* with the system. The highest rated items were *reliability (data integrity)* and the *frequency of scheduled outages.* The lowest rated aspects were the *htar interface* and the *hsi interface.* Table 17 summarizes these satisfaction statistics by project allocation.

Of the six respondents that reported **reasons for dissatisfaction** with HPSS, all mentioned htar or hsi limitations or inconveniences:

> *"hsi and htar interface are difficult to use when archiving multiple large (tens to hundreds of GB) simulations."*
>
> *"hsi interface is tedious to use and hasn't change in forever."*
>
> *"I have stopped using htar as I cannot incrementally save with it."*
>
> *"I love a good commandline tool. But hsi/htar need some love. They don't have any of the conveniences of modern terminals. It's like using a terminal from the 80s or something. So I say modernize hsi/htar or just go with Globus like a lot of folks seem to be doing now (I'm not necessarily in love with it either, but better than hsi/htar)."*
>
> *"There are times when the htar process takes very long, for some unknown reason. There are also frequent outages of the HPSS. The interface is also very un-intuitive."*
>
> *"Tough to find the files you are looking for, since tab to autocomplete does not work."*

Table 16. *Satisfaction Ratings of HPSS by PI Status and Overall*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
hsi interface	23	4.3	0.93	82%	74	4.4	0.97	86%	97	4.4	0.96	85%
htar interface	23	4.2	1.11	71%	74	4.5	1.05	86%	97	4.4	1.07	82%
Ability to store files	23	4.4	0.66	91%	74	4.5	0.71	95%	97	4.5	0.69	94%
Ability to retrieve files	23	4.4	0.84	91%	74	4.4	0.74	95%	97	4.4	0.76	94%
Reliability (data integrity)	23	4.7	0.63	95%	74	4.6	0.70	97%	97	4.6	0.68	97%
Time to store files	23	4.4	0.73	87%	74	4.4	0.79	91%	97	4.4	0.78	90%
Time to retrieve files	23	4.4	0.79	86%	74	4.3	0.85	88%	97	4.4	0.83	88%
Frequency of scheduled outages	23	4.7	0.70	95%	74	4.5	0.67	95%	97	4.6	0.68	95%
Frequency of (unanticipated) unscheduled outages	23	4.7	0.70	95%	74	4.5	0.81	92%	97	4.5	0.79	92%
Overall satisfaction with HPSS	23	4.5	0.67	91%	74	4.5	0.69	96%	97	4.5	0.68	95%

Table 17. *Satisfaction Ratings of HPSS by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
hsi interface	44	4.1	0.95	79%	66	4.4	0.91	87%	23	4.1	1.10	82%	13	4.5	0.66	100%
htar interface	44	4.3	1.09	76%	66	4.4	1.02	83%	23	4.3	1.19	85%	13	4.8	0.73	100%
Ability to store files	44	4.5	0.59	95%	66	4.5	0.64	92%	23	4.3	0.86	96%	13	4.5	0.52	100%
Ability to retrieve files	44	4.4	0.69	95%	66	4.4	0.74	92%	23	4.3	0.86	96%	13	4.5	0.52	100%
Reliability (data integrity)	44	4.6	0.54	100%	66	4.6	0.62	97%	23	4.4	0.89	96%	13	4.6	0.51	100%
Time to store files	44	4.4	0.75	89%	66	4.4	0.74	88%	23	4.2	0.85	96%	13	4.5	0.52	100%
Time to retrieve files	44	4.3	0.76	88%	66	4.3	0.83	85%	23	4.1	0.81	96%	13	4.3	0.63	92%
Frequency of scheduled outages	44	4.5	0.63	95%	66	4.6	0.66	94%	23	4.3	0.82	91%	13	4.8	0.60	100%
Frequency of (unanticipated) unscheduled outages	44	4.5	0.73	93%	66	4.6	0.75	92%	23	4.3	0.96	91%	13	4.8	0.60	100%
Overall satisfaction with HPSS	44	4.5	0.63	93%	66	4.5	0.61	94%	23	4.3	0.92	91%	13	4.5	0.52	100%

Lustre/Spider Scratch Filesystem

Lustre/Spider was used by 65% of respondents during the 2018 calendar year (*N* = 417). Lustre/Spider users were asked to provide satisfaction ratings for multiple aspects of the system, and descriptive statistics of these ratings are shown in Table 18, which also reports satisfaction statistics by PI status. 92% of respondents were overall either *satisfied* or *very satisfied* with the system. The *size* was the highest rated Lustre/Spider feature, and the lowest rated features were the *File and directory operations* and *the frequency of unscheduled outages*. Table 19 summarizes these satisfaction statistics by project allocation.

There were 20 users who **indicated dissatisfaction** with at least one aspect of the Lustre/Spider Scratch Filesystem and most comments were concerned either with *filesystem performance* or with the *stability/reliability and frequency of outages*. For example:

> *"Latency on Lustre is bewilderingly poor."*
>
> *"Of course, it is always hard to do file operations and search/find data on Lustre. It's slow."*
>
> *"Lustre hiccups seem to occur really frequently and have a huge impact on job performance. This may be a more recent thing but since mid-summer or so it has been frustrating to commonly have Lustre hang."*
>
> *"As mentioned in my previous comment, I have had tens of jobs crash and waste compute time because the filesystem had an unscheduled outage, and I would estimate the number of days where my work was negatively impacted by unscheduled Lustre outages as greater than 10 days."*
>
> *"Purging of persistent large input data can interrupt jobs and cause them to fail."*

See Appendix E: User Dissatisfaction Explanations for all replies by category.

Table 18. *Satisfaction Ratings of Lustre/Spider Scratch Filesystem by PI Status and Overall*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Size	57	4.6	0.60	95%	209	4.5	0.65	95%	266	4.5	0.64	95%
I/O bandwidth	58	4.5	0.71	91%	207	4.4	0.70	93%	265	4.4	0.71	92%
File and directory operations	58	4.3	0.84	86%	209	4.3	0.87	89%	267	4.3	0.86	88%
Reliability (data integrity)	58	4.5	0.66	95%	209	4.4	0.78	92%	267	4.5	0.75	93%
Frequency of scheduled outages	55	4.4	0.65	91%	204	4.5	0.63	94%	259	4.4	0.63	93%
Frequency of (unanticipated) unscheduled outages	55	4.3	0.77	89%	202	4.4	0.75	90%	257	4.4	0.75	90%
Overall satisfaction with Lustre/Spider filesystem	57	4.4	0.71	91%	210	4.4	0.75	92%	267	4.4	0.74	92%

Table 19. *Satisfaction Ratings of Lustre/Spider Scratch Filesystem by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Size	116	4.5	0.65	93%	174	4.5	0.62	95%	58	4.3	0.82	90%	49	4.4	0.64	96%
I/O bandwidth	115	4.4	0.68	90%	174	4.4	0.67	92%	58	4.2	0.80	90%	49	4.3	0.86	88%
File and directory operations	116	4.3	0.88	88%	175	4.3	0.86	86%	58	4.1	0.96	84%	49	4.0	1.12	80%
Reliability (data integrity)	117	4.5	0.74	93%	174	4.4	0.73	91%	59	4.2	0.94	88%	49	4.2	0.90	92%
Frequency of scheduled outages	113	4.5	0.61	94%	171	4.5	0.62	94%	57	4.3	0.63	95%	47	4.4	0.70	87%
Frequency of (unanticipated) unscheduled outages	112	4.4	0.71	90%	169	4.4	0.75	91%	57	4.2	0.75	93%	48	4.2	0.83	83%
Overall satisfaction with Lustre/Spider filesystem	116	4.3	0.72	89%	175	4.4	0.70	92%	58	4.1	0.82	88%	49	4.1	0.91	84%

Support Services

Figure 7 shows how frequently respondents **submitted queries** to OLCF in 2018. Half submitted between 1 and 5 requests, while **one-third had not submitted any queries at all.**

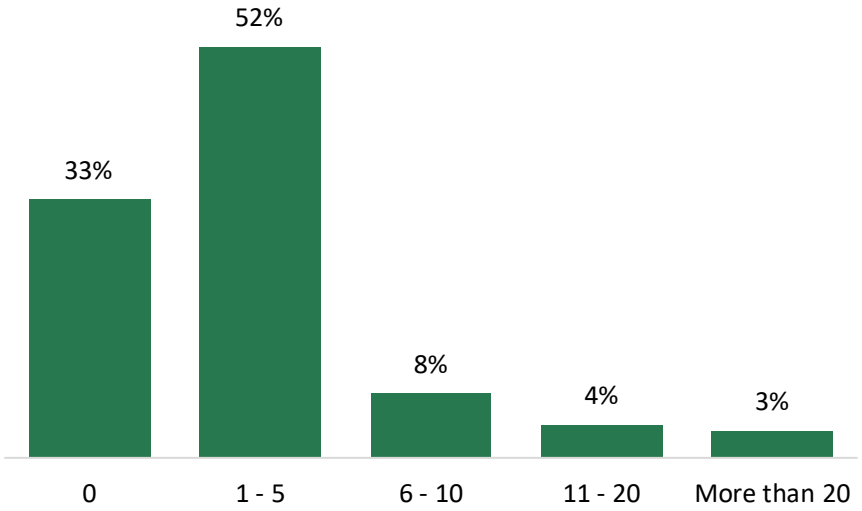

Figure 7. Distribution of number of queries submitted to OLCF in 2018 (*N* = 406)

User Assistance Center

The User Assistance Center was used by 58% of respondents during the 2018 calendar year (*N* = 405). Nearly all respondents that used the OLCF User Assistance Center (97%) were either *satisfied* or *very satisfied* (Table 20 and Table 21). Three respondents reported **reasons for dissatisfaction**:

> *"My request was about power profiling of a Summit-Dev node (for both CPU and GPU). I got a response saying that no such feature is available. I'm not sure if this has changed, though."*
>
> *"There was a significant wait - almost a week - for a response after the initial acknowledgement for support, for one of my requests. I was unable to determine if it was because the request was not sent on to the correct party or if it simply took a while to get a response. I received a response with a solution within a few hours of my second email asking when I would receive it."*
>
> *"Unfortunately, one of my tickets got no response at all."*

Table 20. *Satisfaction Ratings of the User Assistance Center by PI Status and Overall*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Speed of initial response to queries	58	4.6	0.68	98%	172	4.6	0.53	98%	230	4.6	0.57	98%
Speed of final resolution to queries	58	4.6	0.53	98%	172	4.5	0.70	92%	230	4.5	0.66	93%
Quality of technical information	57	4.7	0.57	98%	166	4.6	0.64	92%	223	4.6	0.62	94%
Response to special requests (e.g., scheduling exceptions, quota increases, software installations, etc.)	47	4.6	0.57	96%	127	4.5	0.70	88%	174	4.5	0.67	90%
Overall support from User Assistance	57	4.7	0.51	98%	172	4.6	0.56	97%	229	4.6	0.54	97%

Table 21. *Satisfaction Ratings of the User Assistance Center by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Speed of initial response to queries	83	4.6	0.49	100%	166	4.6	0.61	97%	44	4.5	0.55	98%	52	4.7	0.47	100%
Speed of final resolution to queries	83	4.5	0.61	96%	166	4.5	0.69	93%	44	4.5	0.59	95%	52	4.6	0.70	92%
Quality of technical information	80	4.5	0.64	93%	162	4.6	0.60	95%	41	4.6	0.59	95%	51	4.6	0.67	90%
Response to special requests (e.g., scheduling exceptions, quota increases, software installations, etc.)	61	4.4	0.69	89%	130	4.6	0.64	92%	32	4.4	0.71	88%	33	4.6	0.66	91%
Overall support from User Assistance	82	4.6	0.54	98%	166	4.6	0.54	97%	43	4.6	0.58	95%	52	4.7	0.58	94%

Account Management

23% of users utilized Account Management services in 2018 (*N* = 402). Users were asked to rate their satisfaction with two aspects of Account Management as well as provide an overall rating. Descriptive statistics for ratings by PI status and overall and by project allocation are shown in Table 22 and Table 23. Across various categories of users, mean ratings for *speed of responses to account management queries*, *effectiveness of response to account management queries*, and *overall account services* were similar (all means between 4.6 and 4.7, with satisfaction percentages from 93% to 100%).

Only one respondent provided a reason for **dissatisfaction:**

> *"I have spent the semester trying to get my Chinese student access to resources. I began pursuing this in earnest on Oct 19 when my student reported that he was unable to log in to summitdev (he had access last spring, but it seems to have vanished Aug 15.) It took until after the first week in November until he had access to Summitdev restored. After submitting a request for access to summit through [name redacted's] project . . . on Nov 15, I was just informed that we likely won't have access until Jan 1. We've spent much of the fall just waiting."*

Table 22. *Satisfaction Ratings of Account Management by PI Status and Overall Totals*

	PI Status (*N* = 30)			Non-PI Status (*N* = 61)			Total (*N* = 91)		
	M	***SD***	**%Sat**	***M***	***SD***	**%Sat**	***M***	***SD***	**%Sat**
Speed of responses to queries	4.6	0.57	97%	4.7	0.60	97%	4.6	0.59	97%
Effectiveness of response to queries	4.6	0.56	97%	4.7	0.66	93%	4.6	0.62	95%
Overall support from accounts team	4.6	0.49	100%	4.7	0.66	93%[1]	4.7	0.60	96%[2]

[1]n = 60; [2]n= 90

Table 23. *Satisfaction Ratings of Account Management by Project Allocation*

	INCITE (*N* = 36)			DD (*N* = 59)			ALCC (*N* = 23)			ECP (*N* = 26)		
	M	***SD***	**%Sat**	***M***	***SD***	**%Sat**	***M***	***SD***	**%Sat**	***M***	***SD***	**%Sat**
Speed of responses to queries	4.6	0.49	100 %	4.6	0.55	97%	4.7	0.49	100%	4.7	0.68	96%
Effectiveness of response to queries	4.7	0.48	100 %	4.6	0.61	93%	4.7	0.47	100%	4.7	0.67	96%
Overall support from accounts team	4.6	0.55	97%	4.6	0.58	95%	4.7	0.48	100%[1]	4.7	0.68	96%

[1]n = 22

INCITE Scientific Computing Liaisons

Just over a fifth (21%) of users (*N* = 403) indicated that their project has an assigned INCITE scientific computing liaison and Table 24 shows that 95% of respondents were overall either s*atisfied* or v*ery satisfied* with their liaison (see also Table 25).

There were no users who indicated they were dissatisfied with INCITE liaisons, but one user who rated their satisfaction as *Neutral* provided the following comment:

> *"We don't hear much from our liaison. Essentially all OLCF news and INCITE requirements come via OLCF general and/or INCITE reporting emails."*

Three other users whose projects do have INCITE liaisons provided comments indicating that they have not had ongoing contact with their liaisons:

> *"I do not interact directly with this person."*
>
> *"I don't recall any significant interactions with our liaison this year, in part because the grant funding this work includes support for staff at ORNL and General Atomics to help on code issues."*
>
> *"I'm not sure who my Liaison is. This is probably more my fault than theirs. I can't focus on using OLCF as much as I used to."*

Table 24. *Satisfaction Ratings of INCITE Liaisons by PI Status and Overall Totals*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Speed of initial response to queries	14	4.8	0.43	100%	46	4.8	0.52	96%	60	4.8	0.50	97%
Speed of final resolution to queries	14	4.8	0.43	100%	46	4.7	0.58	93%	60	4.7	0.55	95%
Quality of technical support	14	4.9	0.36	100%	46	4.8	0.52	96%	60	4.8	0.49	97%
Response to special requests (e.g., scheduling exceptions, quota increases, software installations, etc.)	13	4.8	0.38	100%	43	4.7	0.59	93%	56	4.8	0.55	95%
Overall support from your INCITE Scientific Computing Liaison	15	4.7	0.59	93%	48	4.8	0.53	96%	63	4.7	0.54	95%

Table 25. *Satisfaction Ratings of INCITE Liaisons by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Speed of initial response to queries	47	4.8	0.50	96%	36	4.8	0.50	97%	18	4.4	0.70	89%	9	4.3	0.87	78%
Speed of final resolution to queries	47	4.8	0.56	94%	36	4.7	0.58	94%	18	4.4	0.70	89%	9	4.3	0.87	78%
Quality of technical support	47	4.8	0.48	96%	36	4.8	0.50	97%	18	4.4	0.70	89%	9	4.4	0.88	78%
Response to special requests (e.g., scheduling exceptions, quota increases, software installations, etc.)	44	4.8	0.57	93%	34	4.7	0.58	94%	18	4.5	0.71	89%	9	4.4	0.88	78%
Overall support from your INCITE Scientific Computing Liaison	50	4.8	0.55	94%	38	4.7	0.50	97%	19	4.4	0.76	84%	9	4.4	0.88	78%

Communication with Users

As seen in Table 26, most respondents (91%) were either *satisfied* or *very satisfied* with overall OLCF communication. Ratings for the *monthly conference calls* were among the lowest across all categories of users (Table 26 and Table 27). Five respondents provided explanations for their **dissatisfaction with one or more aspects of OLCF communication,** primarily related to *too frequent communication* (*n* = 2) or *ineffective communication* (*n* = 2):

> *"Discarded all those emails. Never paid attention to these comms. I just want to use the computer as developer in multiphase flows, the rest if for other people."*
>
> *"I receive weekly emails when I am not at the lab for months on end, which is annoying."*
>
> *"Many of the weekly announcements could be changed to a monthly newsletter and weekly urgent updates."*
>
> *"Slides should be distributed at the time of presentation or right after. Currently they are released several days/week later and we are asked to check event page later for updates. NOT effective."*
>
> *"There have been at least a handful of times when the OLCF center status indicators indicated no outages when in fact the Lustre filesystem was not responding."*

All open-ended responses are provided in Appendix E: User Dissatisfaction Explanations.

Table 26. *Satisfaction Ratings of Communication by PI Status and Overall Totals*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Weekly e-mail	82	4.5	0.67	93%	299	4.4	0.72	90%	381	4.4	0.71	91%
Monthly conference calls	70	4.4	0.67	90%	227	4.3	0.76	85%	297	4.3	0.74	86%
OLCF Center announcements	75	4.5	0.64	92%	254	4.3	0.74	85%	329	4.3	0.72	87%
OLCF Center status	80	4.5	0.62	96%	277	4.3	0.72	87%	357	4.4	0.70	89%
Overall communications	81	4.5	0.55	98%	302	4.4	0.72	89%	383	4.4	0.69	91%

Table 27. *Satisfaction Ratings of Communications by Project Allocation*

	INCITE				DD			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Weekly e-mail	144	4.4	0.70	89%	253	4.4	0.71	91%
Monthly conference calls	110	4.4	0.74	86%	204	4.3	0.71	86%
OLCF Center announcements	125	4.4	0.70	87%	220	4.3	0.69	88%
OLCF Center status	139	4.4	0.65	91%	235	4.4	0.68	90%
Overall communications	145	4.4	0.72	89%	254	4.4	0.67	91%
	ALCC				**ECP**			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Weekly e-mail	81	4.4	0.73	91%	70	4.5	0.61	94%
Monthly conference calls	63	4.2	0.87	83%	53	4.4	0.74	89%
OLCF Center announcements	72	4.2	0.77	86%	57	4.4	0.68	89%
OLCF Center status	80	4.3	0.76	88%	65	4.5	0.64	92%
Overall communications	82	4.3	0.72	91%	71	4.5	0.63	93%

In addition, nearly all of the 400 users that responded felt **well informed** about OLCF *changes* (97%), *events* (98%), and *current issues* (96%). Four users who indicated they were **not well informed** about OLCF changes provided comments:

> *"I didn't notice the change of Titan libraries when I re-compile and re-run my code."*
>
> *"I would notice them if I start a job and it does not work, then I worry. Otherwise, I do not care."*
>
> *"I'm too new to ORNL to know all the intricacies of OLCF."*
>
> *"The process of moving to Summit has been somewhat opaque, as were calls for Summit Early Science Proposals."*

One user who indicated they were not well informed about OLCF events indicated, *"I could, but I do not pay attention to those."* Finally, two respondents provided feedback about communication regarding current issues:

> *"Don't keep track unless something goes wrong."*
>
> *"Current issues with Summitdev are not communicated to users."*

Training and Technical Reference Documentation

Users were asked to provide satisfaction ratings for their overall satisfaction with OLCF training and technical reference documentation, and five specific aspects. 94% of respondents were overall either *satisfied* or *very satisfied* with the system. The *Getting Started Guide* was the highest rated specific aspect, while the *training calendar* was the lowest rated feature. This pattern was largely observed regardless of PI status or project allocation, although ECP users were least satisfied with *software pages* (Table 28 and Table 29.)

Table 28. *Satisfaction Ratings of Training and Technical Documentation Aspects by PI Status and Overall Totals*

	PI Status				Non-PI Status				Total			
	N	*M*	*SD*	%Sat	*N*	*M*	*SD*	%Sat	*N*	*M*	*SD*	%Sat
Getting Started Guide	79	4.4	0.57	97%	291	4.5	0.61	96%	370	4.5	0.60	96%
System User Guides	81	4.4	0.59	96%	290	4.5	0.66	94%	371	4.5	0.64	95%
Training calendar	59	4.2	0.76	81%	217	4.3	0.72	87%	276	4.3	0.73	85%
Archived training event slides	60	4.3	0.77	81%	220	4.4	0.72	88%	280	4.3	0.73	86%
Software pages	71	4.2	0.82	81%	250	4.3	0.70	88%	321	4.3	0.73	86%
Overall satisfaction	81	4.4	0.66	90%	288	4.5	0.63	95%	369	4.4	0.64	94%

Table 29. *Satisfaction Ratings of Training and Technical Documentation Aspects by Project Allocation*

	INCITE				DD			
	N	*M*	*SD*	%Sat	*N*	*M*	*SD*	%Sat
Getting Started Guide	142	4.5	0.59	97%	243	4.4	0.59	96%
System User Guides	143	4.4	0.66	94%	244	4.4	0.63	95%
Training calendar	106	4.3	0.75	85%	184	4.3	0.73	84%
Archived training event slides	104	4.4	0.72	88%	187	4.3	0.73	85%
Software pages	124	4.3	0.74	85%	208	4.2	0.72	86%
Overall satisfaction	141	4.4	0.65	92%	244	4.4	0.65	93%
	ALCC				**ECP**			
	N	*M*	*SD*	%Sat	*N*	*M*	*SD*	%Sat
Getting Started Guide	80	4.4	0.60	96%	72	4.5	0.67	94%
System User Guides	82	4.3	0.68	93%	71	4.5	0.69	93%
Training calendar	63	4.2	0.72	86%	51	4.3	0.81	84%
Archived training event slides	63	4.3	0.71	89%	51	4.3	0.79	86%
Software pages	72	4.2	0.72	89%	60	4.2	0.78	83%
Overall satisfaction	81	4.3	0.69	94%	71	4.4	0.74	92%

Six respondents who were **dissatisfied with training and/or technical reference documentation** provided the following comments:

> *"Some of the documentation, especially for python, is really very bare bones. It would be nice to have more detailed documentation."*

"Online documentation for cumulus is scarce."

"Mentioned previously - user guides are not up to date on some aspects, debugging specifically."

"I wish issues and other info were more up to date. Other sources than OLCF's page is often where I find what I am looking for."

"I was not aware about many of these services until taking this questionnaire. Starting work here can be pretty overwhelming with the masses of training procedures and literature to read."

"Earlier this year the OLCF website was changed, and many support/documentation hyperlinks were broken and many appear to still be."

Respondents also reported their preferences with respect to **mode and timing of training**:

- The most popular modes of training were *online documentation* and *online training* (Figure 8)
- Respondents had the option to suggest another mode of training that was not listed; while 6 respondents selected *other*, only 4 provided a suggestion. These included *videos; seeking training on my own;* using *email to research specific capabilities*; and an indication of no interest in training.
- Most expressed *no preference* as to time of year (56%, *N* = 400), and among those with a preference, two-thirds chose the summer (Figure 9).

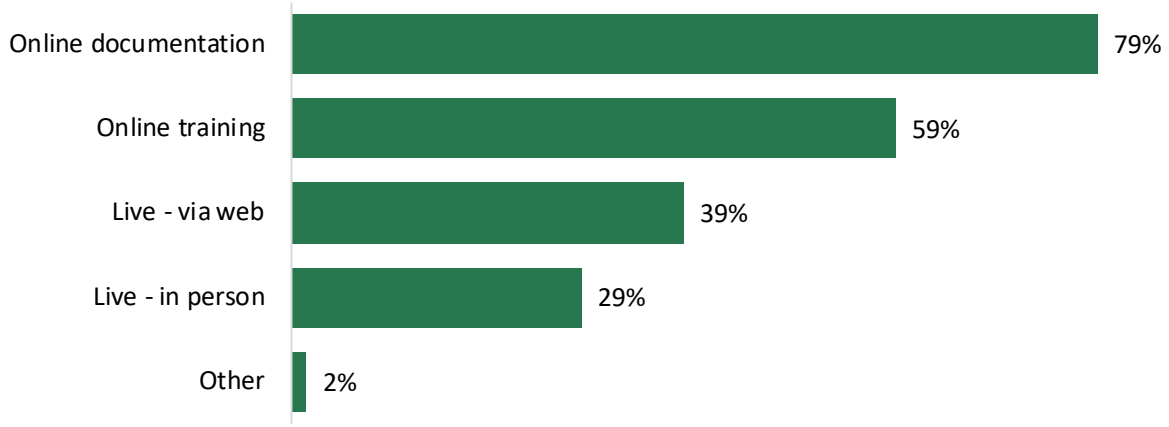

Figure 8. Training preferences of OLCF users (*N* =400)

Note: Percentages add to more than 100% because users could indicate multiple preferences.

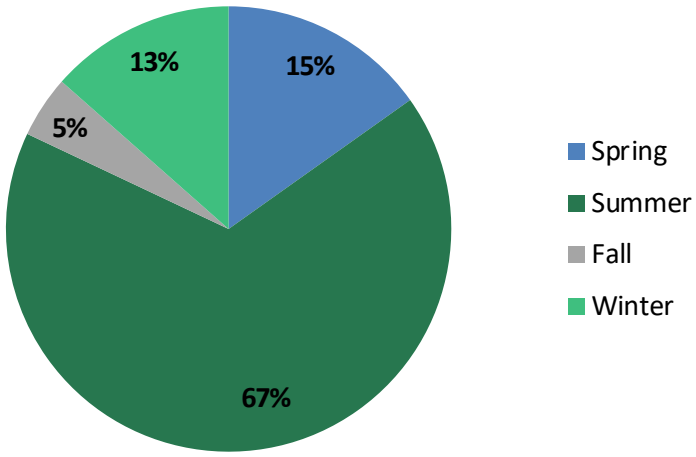

Figure 9. Most convenient time to attend a training event (*N* = 400), disregarding the 56% of respondents who indicated no preference

Finally, 50 respondents suggested future **training topics** in 21 categories (Table 30). The most frequently suggested topics were *common tasks and OLCF basics, software setup for project-specific needs, Summit, GPU resources, Python,* and *example scripts* (16%, 10%, 8%, 8%, 8%, 8%, respectively). See Appendix F: User Suggestions for Improvement for all topic suggestions organized by category.

Table 30. *Users' Suggestions for Training Topics (N = 50)*

Topic	*n*	Percentage
Common tasks/OLCF basics	8	16%
Software for project-specific use	5	10%
Summit	4	8%
GPU resources and programming	4	8%
Python	4	8%
Example scripts	4	8%
Satisfied with available training	3	6%
Remote visualization	3	6%
Updates to existing documentation	3	6%
jsrun	2	4%
Debugging tools	2	4%
HPC resources	2	4%
Machine learning	2	4%
Hackathon	1	2%
OpenMP	1	2%
Startup training	1	2%
Coding/Code Optimization	1	2%

Using containers	1	2%
Data Transfer	1	2%
Compiling	1	2%
Other/Miscellaneous	10	20%

Note: Users add up to more than 100% because some provided more than one theme in their response.

OLCF Website

Before indicating their satisfaction with various aspects of the website, users were asked how frequently they visit the OLCF website (http://olcf.ornl.gov), as displayed in Figure 10. The majority of users visit the website *monthly or less*, and 2% of respondents indicated that they had never visited the site.

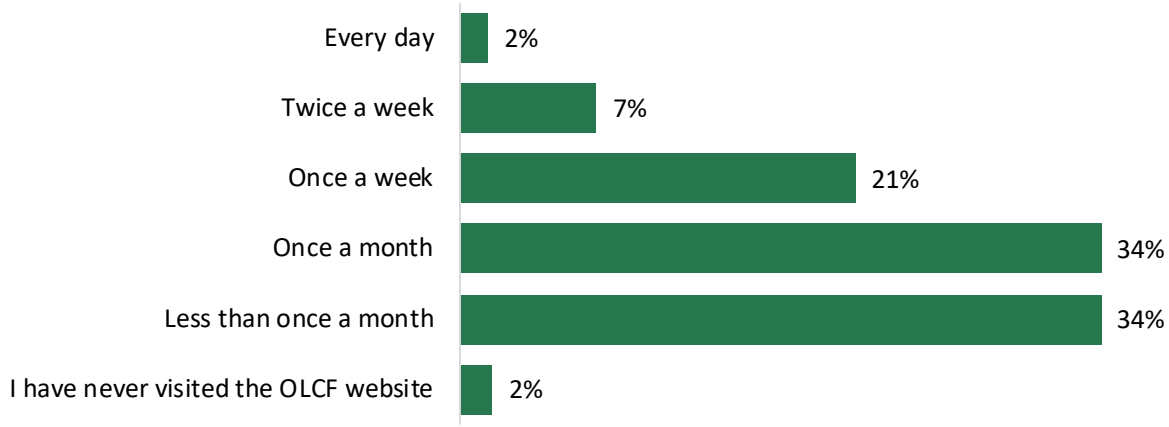

Figure 10. Frequency with which OLCF users visit the OLCF website (*N* = 398)

Users rated several aspects of the website. First, 93% of respondents were either s*atisfied* or v*ery satisfied* overall with the website (Table 31 and Table 32). The highest rated specific aspect of the website was *accuracy of information,* while the lowest rated aspect was *search capabilities.*

Table 31. *Satisfaction Ratings of the OLCF Website by PI Status and Overall Totals*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Ease of navigation	85	4.4	0.65	91%	294	4.3	0.66	91%	379	4.3	0.65	91%
Search capabilities	81	4.2	0.74	83%	277	4.2	0.68	87%	358	4.2	0.70	86%
Accuracy of information	85	4.4	0.72	88%	294	4.4	0.63	93%	379	4.4	0.65	92%
Timeliness of information	83	4.3	0.80	87%	283	4.3	0.64	91%	366	4.3	0.68	90%
Overall satisfaction with the OLCF website	85	4.3	0.66	89%	297	4.4	0.61	94%	382	4.4	0.62	93%

Table 32. *Satisfaction Ratings of the OLCF Website by Project Allocation*

	INCITE				DD			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Ease of navigation	141	4.3	0.67	91%	253	4.3	0.64	91%
Search capabilities	134	4.2	0.70	84%	237	4.2	0.68	87%
Accuracy of information	142	4.3	0.66	91%	252	4.4	0.66	92%
Timeliness of information	138	4.3	0.75	89%	242	4.3	0.70	90%
Overall satisfaction with the OLCF website	144	4.4	0.62	92%	253	4.4	0.61	93%
	ALCC				**ECP**			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Ease of navigation	82	4.2	0.67	90%	71	4.3	0.60	93%
Search capabilities	80	4.1	0.66	86%	66	4.2	0.63	89%
Accuracy of information	83	4.3	0.66	90%	70	4.3	0.73	87%
Timeliness of information	82	4.2	0.68	88%	69	4.2	0.69	88%
Overall satisfaction with the OLCF website	83	4.2	0.67	89%	71	4.3	0.64	90%

There were eight users who reported explanations for their **dissatisfaction with one or more aspects of the website**. The most common complaint was *outdated or incorrect information* (*n* = 4; see Appendix E: User Dissatisfaction Explanations for all comments by category):

> *"There have been some examples in the documentation that don't work. I don't know if they're still there, though."*
>
> *"There are many broken links unfortunately."*
>
> *"The information on the Early Science Program for Summit is stale and likely inaccurate."*
>
> *"Again, it could be more up to date."*

Users were also asked to share their perception of the website updates that were completed in 2018. The largest proportion of respondents were unsure or did not have an opinion (44%), but the next largest proportion found the new website to be an improvement compared to the previous site (35%; see Figure 11).

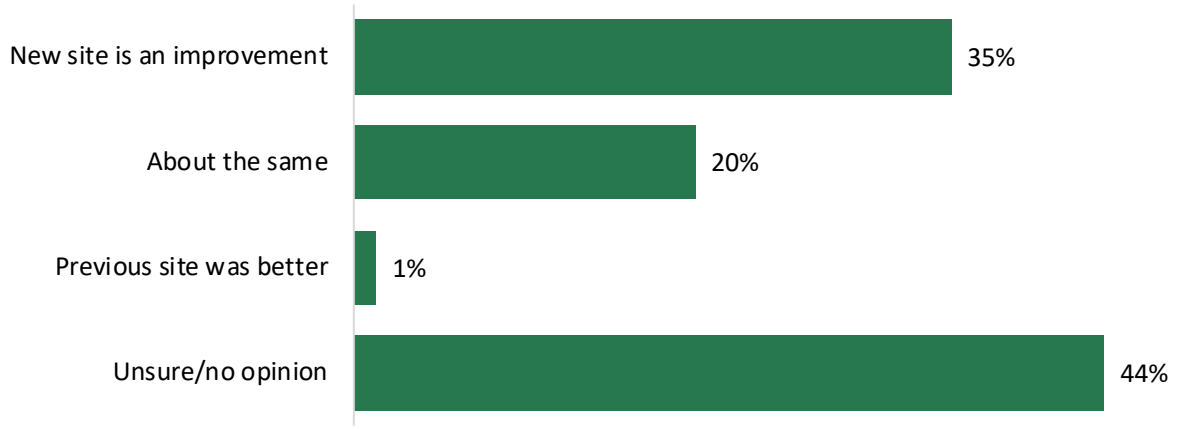

Figure 11. Users' opinions of the 2018 OLCF website updates (*N* = 381)

Four users who indicated that they found the previous version of the website to be better provided the following explanations:

> *"Unfortunately a lot of support/documentation hyperlinks were broken when the update happened."*
>
> *"The floating navigation bar takes up screen space and is distracting. The hover menus also occasionally malfunction and block the page's content."*
>
> *"My feeling is that some information is more difficult to find in the new site compared to that in the previous one."*
>
> *"Globus and DTN docs have gone missing."*

Data Analysis and Visualization

Data analysis and visualization services were used by 4% (15 of 387) of respondents during the 2018 calendar year. Service users were asked for overall satisfaction ratings and ratings for multiple specific aspects of the data analysis and visualization support services (Table 33 and Table 34). Table 33 shows that 100% of respondents were overall either *satisfied* or *very satisfied* with the support they received. Responses from PIs differed somewhat from responses from non-PIs. Across project allocations, DD and ALCC respondents provided lower ratings on some items (Table 34).

Satisfaction with several specific aspects of data analysis, visualization and workflow are summarized in Table 35 and Table 36, which show that satisfaction ranged:

- from 73% to 79% across all respondents,
- from 71% to 84% for PIs, and
- from 66% to 86% across project allocations.

Users gave the lowest ratings to *ability to perform project workflows* and *sufficiency of software tools.*

Table 33. *Satisfaction Ratings for Data Analysis and Visualization Support Services by PI Status and Overall Totals*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Speed of initial response to queries	6	4.3	1.03	67%	7	4.4	0.53	100%	13	4.4	0.77	85%
Speed of final resolution to queries	6	4.5	0.84	83%	7	4.4	0.53	100%	13	4.5	0.66	92%
Quality of technical support	6	4.7	0.52	100%	7	4.4	0.53	100%	13	4.5	0.52	100%
Overall support from the data analysis and visualization support personnel	6	4.7	0.52	100%	7	4.3	0.49	100%	13	4.5	0.52	100%

Table 34. *Satisfaction Ratings for Data Analysis and Visualization Support Services by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Speed of initial response to queries	8	4.5	0.53	100%	8	4.4	0.74	88%	4	4.0	1.15	50%	2	5.0	0.00	100%
Speed of final resolution to queries	8	4.5	0.53	100%	8	4.5	0.53	100%	4	4.3	0.96	75%	2	5.0	0.00	100%
Quality of technical support	8	4.5	0.53	100%	8	4.5	0.53	100%	4	4.5	0.58	100%	2	5.0	0.00	100%
Overall support from the data analysis and visualization support personnel	8	4.4	0.52	100%	8	4.5	0.53	100%	4	4.5	0.58	100%	2	5.0	0.00	100%

Table 35. *Satisfaction Ratings for Data Analysis, Visualization and Workflow by PI Status and Overall Totals*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Ability to perform data analysis	38	4.2	0.87	78%	123	4.1	0.77	79%	161	4.1	0.79	79%
Ability to perform project workflows	35	4.1	0.84	71%	117	4.0	0.71	76%	152	4.0	0.74	75%
Sufficiency of the OLCF hardware for your data analysis, visualization, and workflow needs	39	4.2	0.74	84%	125	4.1	0.73	78%	164	4.1	0.73	79%
Sufficiency of software tools for your data analysis, visualization, and workflow needs	38	4.0	0.82	76%	125	4.0	0.76	72%	163	4.0	0.77	73%

Table 36. *Satisfaction Ratings for Data Analysis and Visualization, and Workflow by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Ability to perform data analysis	71	4.2	0.70	86%	107	4.0	0.77	78%	39	4.2	0.76	85%	33	4.1	0.83	72%
Ability to perform project workflows	69	4.1	0.71	82%	100	4.0	0.72	72%	38	4.0	0.68	76%	30	3.9	0.83	66%
Sufficiency of the OLCF hardware for your data analysis, visualization, and workflow needs	73	4.2	0.71	85%	108	4.1	0.71	78%	41	4.1	0.70	80%	33	4.2	0.76	81%
Sufficiency of software tools for your data analysis, visualization, and workflow needs	74	4.0	0.77	75%	107	3.9	0.72	73%	40	4.0	0.70	78%	33	4.0	0.85	72%

Users were then asked to indicate where they analyze data produced by OLCF jobs. Of the 387 respondents who answered this question, 17% (*n* = 64) indicated that they do not need or use data analysis from OLCF. These responses were then excluded from further analysis, and the distribution of other responses in Figure 12 shows that the largest proportion of users analyzed all of their data "elsewhere" and the smallest proportion analyzed it all at OLCF.

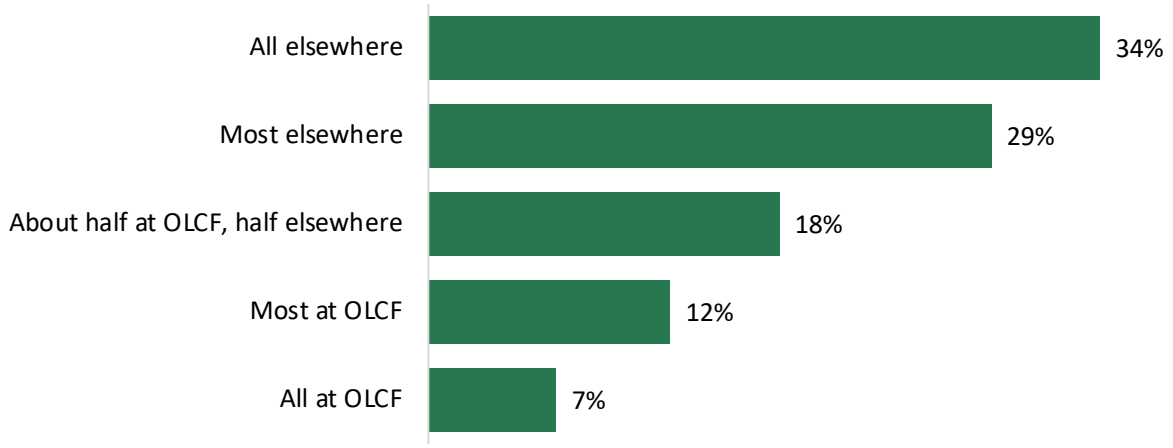

Figure 12. Locations for analysis of data produced by OLCF jobs (*N* = 323)

To put these results in context, users were also asked about the source of their data, displayed in Figure 13. The largest proportion of users are working with data that is primarily sourced from outside OLCF jobs.

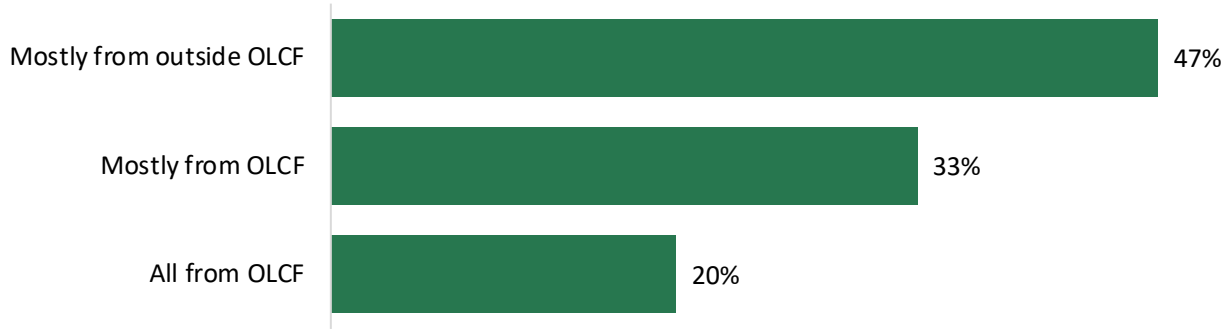

Figure 13. Source of user data (*N* = 387)

Data Services/Feature Priorities

In this section of the survey, users were asked to rate **the importance of 18 different data services/features**. All but two of the options were assessed previously, including on the 2017 survey. Two new options were added and were ranked by users for the first time on the 2018 survey.

Table 37 shows the *%Imp* (the percentage of respondents indicating *very important* or *extremely important*) for all respondents and broken down by both PI status and project allocation (color scale

indicates relative magnitude: high-med-low = green-yellow-red). The table is sorted from top to bottom from the most important to the least important services/features based on the total across all users. Examination of the table shows that regardless of respondent category, the **consistently most valued item** was *long-term data retention* (*%Imp* ranging from 51% to 73%). Generally, the **least valued item** was t*he availability of Parallel Big Data R (pbdR) Deep Learning and Machine Learning Consultation for your science use-case* (ranging from 7% to 16%), except that PIs value that item slightly more than their least valued item, *interactive analysis and/or integrated simulation including data from other experimental facilities* (12%).

Table 38 and Table 39 show the same findings, but with greater detail, displaying not only *%Imp* but *M* and *SD* as well. *%Imp* ranged across all respondent categories from 7% to 73%. The rank-ordered importance of services/features across all respondents was mirrored very closely across PI status and project allocations, as above. Many of the items were consistently ranked as unimportant by 60%-80% of all respondents. In future years, some of the options in this question could be selected from the responses to the previous year's survey to focus this investigation on highly relevant features.

Finally, users were asked whether they were interested in scheduling one-on-one conversations with OLCF analysis and visualization specialists in order to consult on needs and approaches. Twenty-four (6%, *N* = 351) respondents indicated they were interested, and their contact information was referred to the OLCF to arrange consultations.

Table 37. *Data Service/Feature Importance (%Imp) Rankings*

	Total (*N* = 385)	PI (*N* = 84)	Non-PI (*N* = 301)	INCITE (*N* = 141)	DD (*N* = 257)	ALCC (*N* = 83)	ECP (*N* = 73)
Long-term data retention	64%	62%	64%	73%	64%	57%	51%
Access for your specific OLCF project members to your data over the web	46%	36%	49%	48%	45%	49%	43%
Long-term data curation	46%	51%	44%	46%	46%	37%	34%
Access for collaborators to your data over the web	41%	31%	43%	45%	37%	49%	44%
Access to a large shared-memory system for data analysis & visualization	33%	23%	36%	41%	34%	27%	25%
Out-of-the-box workflow tools/libraries	32%	23%	34%	29%	31%	25%	23%
Remote visualization capability	32%	23%	34%	37%	31%	35%	26%
Data management tools	28%	23%	30%	25%	28%	29%	26%
Access to a system with GPUs specifically for data analysis & visualization	26%	20%	28%	30%	27%	25%	23%
General public access to your data over the web	25%	20%	27%	23%	25%	24%	18%
Access to databases at the OLCF	25%	25%	25%	23%	25%	20%	15%
The availability of utilizing Jupyter/IPython Notebook	24%	19%	25%	27%	24%	27%	30%
Dedicated workflow machines	24%	21%	24%	23%	23%	18%	21%
The availability of utilizing containers	23%	20%	23%	20%	24%	33%	25%
Analysis and visualization assistance from the OLCF	22%	18%	24%	22%	23%	22%	18%
Availability of a data, analytics, and visualization liaison	19%	14%	20%	19%	19%	16%	18%
Interactive analysis and/or integrated simulation including data from other experimental facilities	17%	12%	18%	18%	17%	14%	21%
The availability of Parallel Big Data R (pbdR) Deep Learning and Machine Learning Consultation for your science use-case	15%	13%	15%	16%	16%	7%	11%
Min	15%	12%	15%	16%	16%	7%	11%
Max	64%	62%	64%	73%	64%	57%	51%

Table 38. *Data Service/Feature Importance by PI Status and Total (rank ordered by Total %Imp, high to low)*

	PI Status (*N* = 84)			Non-PI Status (*N* = 301)			Total (*N* = 385)		
	M	*SD*	%Imp	*M*	*SD*	%Imp	*M*	*SD*	%Imp
Long-term data retention	3.6	1.12	62%	3.6	1.11	64%	3.6	1.11	64%
Access for your specific OLCF project members to your data over the web	2.9	1.28	36%	3.2	1.26	49%	3.2	1.27	46%
Long-term data curation	3.2	1.28	51%	3.2	1.17	44%	3.2	1.19	46%
Access for collaborators to your data over the web	2.8	1.30	31%	3.1	1.20	43%	3.0	1.23	41%
Access to a large shared-memory system for data analysis & visualization	2.7	1.22	23%	2.9	1.30	36%	2.9	1.28	33%
Out-of-the-box workflow tools/libraries	2.6	1.23	23%	2.9	1.26	34%	2.8	1.26	32%
Remote visualization capability	2.5	1.33	23%	2.8	1.33	34%	2.8	1.34	32%
Data management tools	2.7	1.13	23%	2.8	1.18	30%	2.8	1.17	28%
Access to a system with GPUs specifically for data analysis & visualization	2.6	1.33	20%	2.7	1.27	28%	2.7	1.28	26%
General public access to your data over the web	2.3	1.35	20%	2.6	1.26	27%	2.5	1.29	25%
Access to databases at the OLCF	2.6	1.23	25%	2.6	1.23	25%	2.6	1.23	25%
The availability of utilizing Jupyter/IPython Notebook	2.5	1.22	19%	2.7	1.31	25%	2.6	1.29	24%
Dedicated workflow machines	2.6	1.11	21%	2.6	1.24	24%	2.6	1.21	24%
The availability of utilizing containers	2.4	1.33	20%	2.6	1.22	23%	2.6	1.25	23%
Analysis and visualization assistance from the OLCF	2.4	1.23	18%	2.6	1.20	24%	2.5	1.21	22%
Availability of a data, analytics, and visualization liaison	2.3	1.17	14%	2.5	1.19	20%	2.4	1.19	19%
Interactive analysis and/or integrated simulation including data from other experimental facilities	2.1	1.19	12%	2.3	1.21	18%	2.2	1.20	17%
The availability of Parallel Big Data R (pbdR) Deep Learning and Machine Learning Consultation for your science use-case	2.1	1.14	13%	2.2	1.19	15%	2.2	1.18	15%

Table 39. *Data Service/Feature Importance by Project Allocation*

	INCITE (*N* = 141)			DD (*N* = 257)			ALCC (*N* = 83)			ECP (*N* = 73)		
	M	***SD***	**%Sat**	***M***	***SD***	**%Sat**	***M***	***SD***	**%Sat**	***M***	***SD***	**%Sat**
Long-term data retention	3.8	0.98	73%	3.7	1.12	64%	3.6	1.03	57%	3.2	1.21	51%
Access for your specific OLCF project members to your data over the web	3.2	1.26	48%	3.1	1.29	45%	3.3	1.26	49%	3.0	1.47	43%
Long-term data curation	3.3	1.20	46%	3.2	1.18	46%	3.0	1.20	37%	3.0	1.22	34%
Access for collaborators to your data over the web	3.1	1.22	45%	3.0	1.27	37%	3.3	1.25	49%	3.1	1.37	44%
Access to a large shared-memory system for data analysis & visualization	3.1	1.28	41%	2.9	1.30	34%	2.6	1.24	27%	2.6	1.34	25%
Out-of-the-box workflow tools/libraries	2.8	1.26	29%	2.8	1.26	31%	2.6	1.29	25%	2.5	1.32	23%
Remote visualization capability	2.9	1.40	37%	2.8	1.34	31%	2.7	1.36	35%	2.5	1.39	26%
Data management tools	2.8	1.06	25%	2.8	1.16	28%	2.8	1.18	29%	2.6	1.29	26%
Access to a system with GPUs specifically for data analysis & visualization	2.8	1.26	30%	2.7	1.30	27%	2.5	1.25	25%	2.4	1.39	23%
General public access to your data over the web	2.6	1.19	23%	2.5	1.31	25%	2.5	1.30	24%	2.3	1.28	18%
Access to databases at the OLCF	2.6	1.21	23%	2.6	1.24	25%	2.5	1.19	20%	2.3	1.20	15%
The availability of utilizing Jupyter/IPython Notebook	2.7	1.33	27%	2.7	1.26	24%	2.6	1.33	27%	2.6	1.47	30%
Dedicated workflow machines	2.6	1.24	23%	2.6	1.21	23%	2.5	1.14	18%	2.4	1.32	21%
The availability of utilizing containers	2.5	1.23	20%	2.6	1.26	24%	2.8	1.38	33%	2.5	1.40	25%
Analysis and visualization assistance from the OLCF	2.5	1.17	22%	2.5	1.22	23%	2.4	1.21	22%	2.2	1.20	18%
Availability of a data, analytics, and visualization liaison	2.5	1.20	19%	2.5	1.17	19%	2.2	1.16	16%	2.2	1.26	18%
Interactive analysis and/or integrated simulation including data from other experimental facilities	2.2	1.23	18%	2.3	1.21	17%	2.1	1.14	14%	2.1	1.34	21%
The availability of Parallel Big Data R (pbdR) Deep Learning and Machine Learning Consultation for your science use-case	2.2	1.23	16%	2.2	1.21	16%	1.9	0.99	7%	2.0	1.13	11%

User Suggestions for Improvement

This section summarizes the suggestions provided by respondents with respect to potential improvements in OLCF resources/services, which includes additions or changes.

OLCF Experience

When asked "What additional services, resources, and/or other improvements are needed to enhance your experience at the OLCF?" 158 respondents supplied comments; 14% indicated satisfaction, i.e., that *no additional services and/or resources* are needed to enhance their experience at the OLCF (Table 40). Among those expressing a need or preference, *tools/software/installations* were mentioned most frequently, followed by *filesystem, I/O, and data transfer* and *visualization and analysis*. See Appendix F: User Suggestions for Improvement for all responses by category. Select comments include:

Tools/Software/Installations	*"Tensorflow, Mathematica and machine learning tools."* *"If there are the same graphic tools in the various machines, it will be good."* *"Support for running software not in the OLCF list of software."*
Filesystem, I/O, and Data Transfer	*"Just more reliable file I/O, overall."* *"Post-processing via fast SSD drives or broader pipes for transferring data for analysis."*
Visualization and Analysis	*"As mentioned more work on interactive remote visualization and analysis is needed. I know through interaction with the visualization team that this is a work in progress."* *"In my view, the major areas where there are ongoing opportunities for improvement at OLCF (and most other HPC centers) relate to support for remote, interactive HPC visualization and analysis tasks that are historically an area that was not yet well served by existing hardware and software infrastructure, policy, and practice. I expect this to be an area of rapidly growing needs due to the continued increase in data sizes addressed by state-of-the-art simulation and experiment, not only in computational biology, but more broadly in other domains as well."*

Table 40. *Users' Suggestions for Additional Services and/or Resources Needed to Enhance Their Experience at the OLCF*

Category	*N* = 158	Percentage
Satisfaction	22	14%
Tools/software/installations	21	13%
Filesystem, I/O and data transfer	16	10%
Visualization and analysis	15	9%
Summit/SummitDev	13	8%
Performance/performance upgrade	12	8%
Training and tutorials	11	7%

Login and access	10	6%
Queue time and scheduling policy	10	6%
Wall/run time	7	4%
Tech support	7	4%
Documentation	6	4%
Data retention/purge policy and procedures	6	4%
Containers	6	4%
Data storage	5	3%
GPU resources	5	3%
Website	5	3%
Capacities	5	3%
Survey comments	5	3%
Libraries and updates	4	3%
Architecture	4	3%
Debugging tools	4	3%
Accounts and account status information	4	3%
User guides	3	2%
Allocations	3	2%
Compilers	3	2%
Workflow	3	2%
Stability/reliability	2	1%
Testing capabilities	2	1%
Miscellaneous/Other	11	7%

Note: Percentages total to more than 100% because responses could mention more than one type of improvement.

Compute or Data Resources

When asked, "Please describe how the OLCF can improve your experience using any of the HPC resources (i.e., Titan, Eos, Rhea, DTNs, HPSS, Lustre/Spider) and/or tell us if any additional resources are needed," 323 respondents provided comments. The largest proportion, 38%, indicated their *satisfaction* with OLCF compute and data resources. This was followed by *making more tools available* (9%), *software issues* (8%), *file systems and data transfer* (7%), *updates to purge policy and procedures* (7%), and *performance* (7%). See Appendix F: User Suggestions for Improvement for all responses by category. Refer to Table 41 for all themes identified.

Select comments include:

Make more tools available	*"More R language services, including access via RStudio Server."* *"My ability to do my research would be greatly improved if NICE DCV remote visualization were made available on Rhea."*
Software issues	*"It seems it is well known that there are some technical difficulty for running analysis software on Titan. And users will need Rhea instead. Not sure why it is the case. I guess users just need to be informed by this fact."*

	"Matlab licenses were short which was disruptive for our work. This may be prevented by limiting the number of licenses used by one user or buying more."
File systems and data transfer	*"Improved file system performance, particularly for interactive use, would be great."* *"Only improvement would be to facilitate the login procedure and ease of (small) data transfers between workstation and login node."*
Purge policy and procedures	*"Give some announcement on frequency at which SCRATCH will be purged."* *"Some of the scripts and files kept in Lustre got deleted automatically several times. This problem should be handled to avoid multiple builds of required software."*
Performance	*"Looking forward to increased resources with Summit."* *"Improve performance of Lustre."*

Table 41. *Users' Suggestions for Improvements to HPC Compute and Data Resources*

Category	***N* = 323**	**Percentage**
Satisfied	123	38%
Make more tools available	28	9%
Software issues	25	8%
File systems and data transfer	23	7%
Updates to data retention/purge policy and procedures	22	7%
Performance	21	7%
Maintenance	18	6%
Queuing policy/faster queues	16	5%
Reliability/stability/uptime of systems	16	5%
Administrative issues	16	5%
Training/instructional resources	15	5%
More documentation	14	4%
Allow more computing time/walltime	13	4%
Support issues	9	3%
GPU Resources	9	3%
Improve storage/memory	8	2%
Install better compilers	7	2%
Improve debugging	6	2%
More attention to small jobs	3	1%
Help with codes	2	1%
Miscellaneous/Other	41	13%

Note: Percentages total to more than 100% because responses could mention more than one type of improvement.

OLCF Website

When asked "What additional services or information would you like to have available on the OLCF website?" 20 users responded. 20% of these respondents indicated their *general satisfaction* with the existing resources. The largest proportion of users (30%) requested *additional tutorials, training, or guidance*; see Figure 14 for other requested services or information, excluding expressions of overall satisfaction.

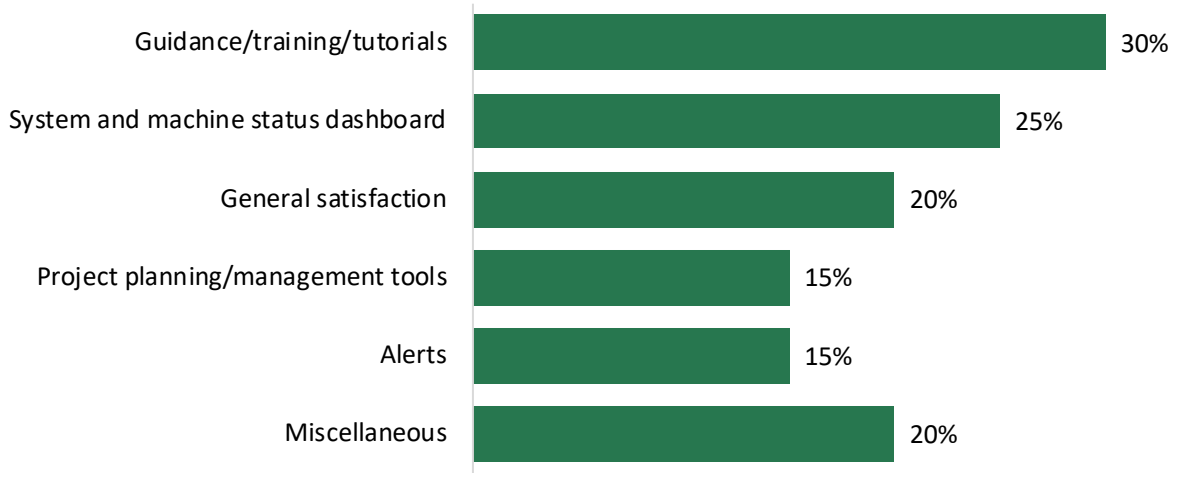

Figure 14. Users' suggestions for additional services or information on the OLCF website

Note: Some users provided more than one theme in their response.

Example replies in the most frequently reported categories are provided below (see Appendix F: User Suggestions for Improvement for all responses by category).

Additional Guidance/ Training/Tutorials	*"More training resources never hurt - also, keeping everything current is important. I don't think there is a problem with stale material right now, but for forward-planning, I think it is important to plan for documentation maintenance."* *"A general area of best practices or hint and tips from users."*
System and Machine Status Dashboard	*"The most important feature to me is to see machine uptime and when systems are down."* *"Current state of nodes and length of queue."*

Finally, survey respondents were introduced to OLCF's plans to develop a new MyOLCF portal with a new look and enhanced functionality. Users who were interested in an updated MyOLCF were asked to share what features should be added to that portal. The 43 users who responded suggested features in 9 different categories (Table 42), in addition to a number who requested more detail or who noted that they do not use MyOLCF functionality. The most frequently suggested features were dashboards and tools for monitoring *allocations, quotas, and usage* (51%), *job tracking/logging* (14%), *alerts and*

reminders (9%), and *project summary/dashboard* functionality (9%). See Appendix F: User Suggestions for Improvement for all topic suggestions organized by category.

Table 42. *Users' Suggestions for Updated Features for MyOLCF (N = 43)*

Topic	*n*	Percentage
Allocations, quotas, and usage	22	51%
Job tracking/logging	6	14%
Alerts/reminders	4	9%
Project summary/dashboard	4	9%
HPC resource/machine status	3	7%
Account summary/dashboard	2	5%
Reporting functions	2	5%
Ticketing/support functions	2	5%
Need more detail	1	2%
Other/Miscellaneous	6	14%
Don't use/not interested	8	19%

Note: Users add up to more than 100% because some provided more than one theme in their response.

Data analysis, visualization, and/or workflow

When asked, "What additional data analysis, visualization, and/or workflow services would you like the OLCF to provide?" 27 users responded, with 5 (19%) indicating they were satisfied with current services. Among those that expressed a need/preference, the largest proportions were interested in *assistance with setup and automation* for related services and *software/tools* for analysis and visualization. (Figure 15).

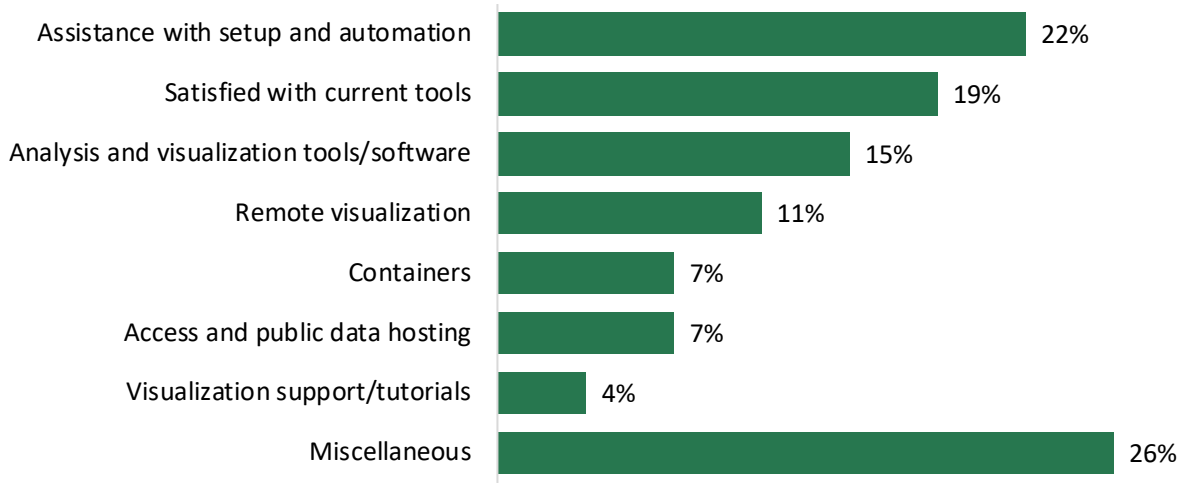

Figure 15. Users' suggestions for additional data analysis, visualization, and/or workflow services

Note: Percentages total to more than 100% because some provided more than one theme in their response.

Example replies in the most frequently reported categories are provided below (see Appendix F: User Suggestions for Improvement for all responses by category).

Assistance with Setup and Automation	*"Had some difficulty in setting up my analysis which has some setup/breakdown components that run on compilers different than Titan. Would be nice if there was a way to do all this work on batch nodes."* *"Better support for automated build/test workflows - not necessarily continuous integration, but automated periodic build/test."*
Analysis and Visualization Tools/Software	*"Tools for visualizing time varying volumetric data."* *"As a ParaView/Catalyst developer I am extremely biased of course, but I would like the ParaView and Catalyst installations to be updated more regularly, made web accessible, and promoted more widely to OLCF users."*

Other OLCF Issues

When asked to comment on any additional **important concerns not covered elsewhere** in the survey, 22 individuals replied. The largest proportion expressed *general satisfaction and appreciation* (23%), tied with *performance and support issues.* Other comments were distributed as seen in Table 43 (Appendix F: User Suggestions for Improvement for text of these comments).

Table 43. *Respondent Comments on Other Issues Not Addressed within the Survey*

Category	*N* = 22	Percentage
General satisfaction	5	23%
Performance and support issues	5	23%
Accounts, access and credentials	4	18%
Survey too long	3	14%
Communication and information sharing	2	9%
Allocations and small/development jobs	2	9%
Miscellaneous	2	9%

Note: Percentages total to more than 100% because responses could mention more than one issue.

Summary of Survey Observations

In most respects, users were highly satisfied with the OLCF resources/services. Table 44 summarizes satisfaction (*satisfied*, or *very satisfied*) ratings. The color scale indicates the relative magnitude of cell values: high-med-low = green-yellow-red. Examination of the table suggests that **satisfaction was highest** (across respondent types) for Rhea, Data Analysis and Visualization Support Services, Eos, and User Assistance. While the **lowest ratings** were found for Data Transfer Nodes, Communication, and Data Resources, these ratings still reflect a generally high satisfaction among users. Across user groups and project types, 16 out of 17 items were rated as either *satisfied* or *very satisfied* by 90% or more of users.

Table 44. *Summary Overall Satisfaction with Aspects of OLCF, by PI Status and Project Allocation*

		PI Status		Project Type			
	All	PI	Non-PI	INCITE	DD	ALCC	ECP
Max N responding:	*419*	*87*	*332*	*159*	*271*	*88*	*80*
OLCF	96%	97%	95%	94%	97%	95%	91%
Compute Resources	94%	94%	94%	95%	94%	94%	92%
Titan	95%	96%	95%	94%	95%	95%	93%
Eos	99%	100%	98%	100%	98%	100%	100%
Rhea	100%	100%	100%	100%	100%	100%	100%
Data Resources	91%	96%	89%	88%	91%	91%	80%
Data Transfer Nodes	87%	85%	88%	86%	88%	91%	92%
HPSS	95%	91%	96%	93%	94%	91%	100%
Lustre/Spider	92%	91%	92%	89%	92%	88%	84%
Support Services	94%	97%	93%	94%	95%	94%	91%
User Assistance	97%	98%	97%	98%	97%	95%	94%
Account Services	96%	100%	93%	97%	95%	100%	96%
Data Analysis and Visualization Support Services	100%	100%	100%	100%	100%	100%	100%
INCITE Liaison	95%	93%	96%	94%	97%	84%	78%
Communication	91%	98%	89%	89%	91%	91%	93%
Training and Technical Reference Documentation	94%	90%	95%	92%	93%	94%	92%
OLCF Website	93%	89%	94%	92%	93%	89%	90%
Min	**87%**	**85%**	**88%**	**86%**	**88%**	**84%**	**78%**
Max	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**

Longitudinal Comparisons of User Responses

This section reviews the results from the 2006 through 2018 OLCF User Surveys. In some cases, this effort has been complicated by changes to the survey items over time, and these are noted throughout.

OLCF Users

Figure 16 shows that **length of time using OLCF** (i.e. experience as an OLCF user) reported by most survey respondents has changed substantially between 2006 and 2018. Prior to 2009, about half of respondents reported using OLCF less than one year, and this category comprised the largest proportion of users. However, between 2009 and 2011, the largest proportion of users indicated having greater than two years of experience at OLCF. In 2012, user experience shifted back to the largest proportion of users reporting using OLCF less than one year. From 2013 to present, users who have been with OLCF for more than two years once again make up the greatest proportion of users, and slightly more than half of respondents are in that category.

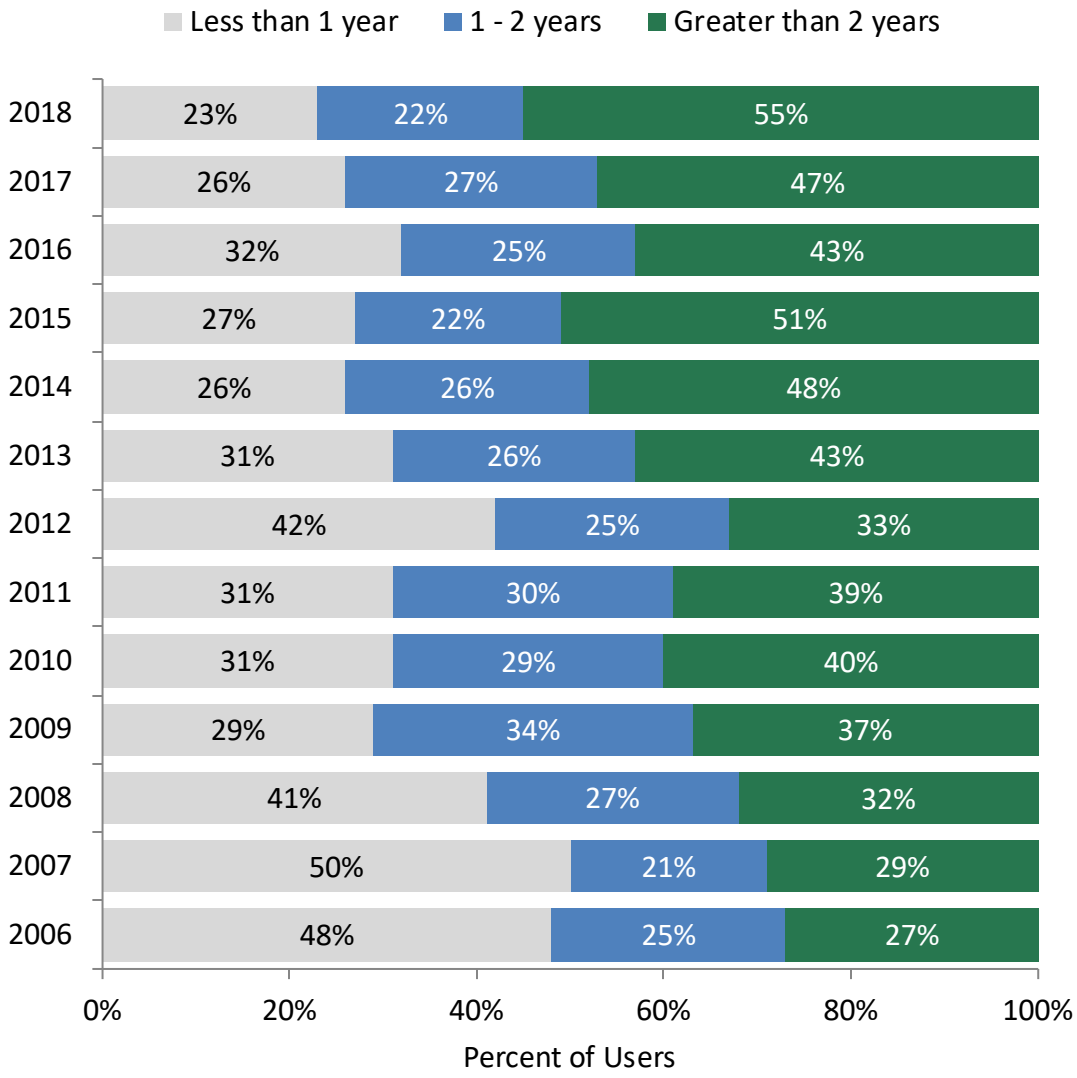

Figure 16. User years of experience with OLCF, 2006-2018.

With respect to **project classifications** (Figure 17), *survey respondent* data is available from 2007 to the present year, and OLCF data for *the entire pool of OLCF users* is available from 2014 to present. The figure shows these side-by-side and indicates that the distribution of respondents has tracked the overall potential sample well. In 2018, the ECP project allocation was added to the dataset.

Generally,

- Until 2017, INCITE projects have shown a downward trend in share of both the respondent and the user pool. In 2017 and 2018, INCITE projects have grown modestly.
- Director's Discretion projects remained relatively constant between 2007 and 2011 for respondents, and have generally trended upward since 2012.
- ASCR Leadership Computing Challenge (ALCC) supported projects began in 2010 and supported only 2% of respondents, but grew significantly from there.

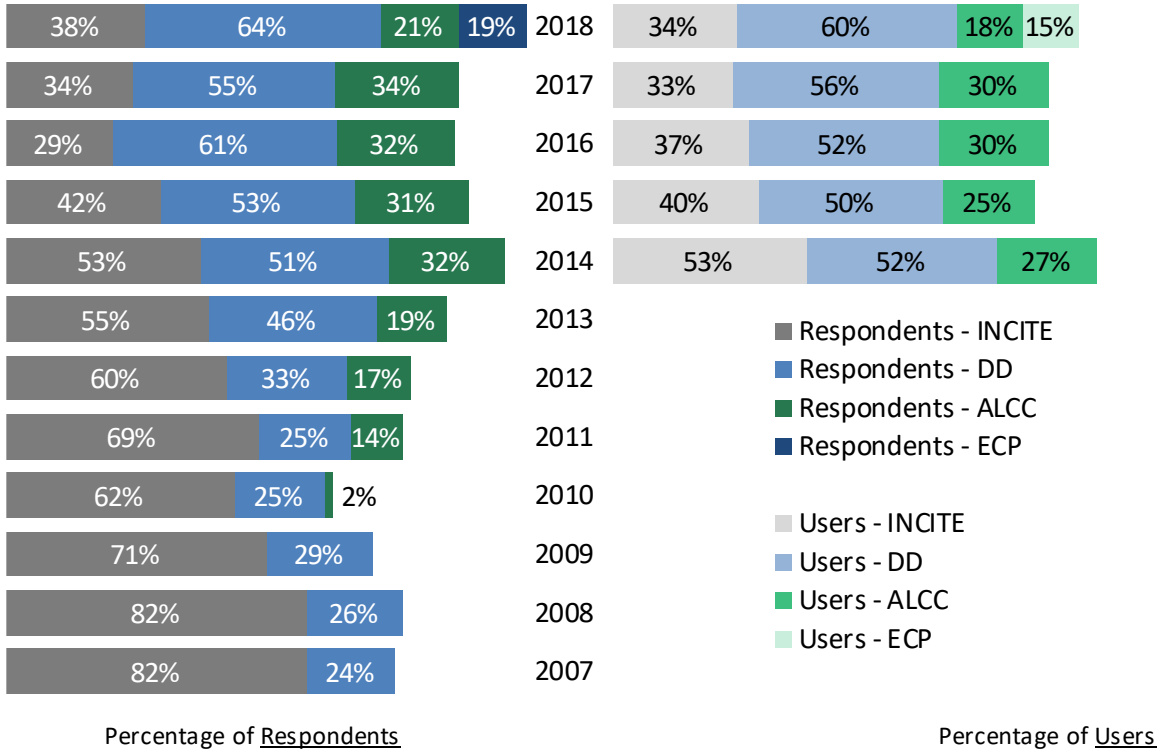

Figure 17. Respondent project allocations, 2007-2018, and OLCF user project allocations, 2014-2018
Note: Percentage total to more than 100% as users are often affiliated with multiple projects.

Computer Systems Utilization

A significant change from 2007 to 2008 was the removal of the Hawk system and the addition of the Lens system to the survey. In 2009, the IBM BlueGene/P (Eugene) and Development (Smoky) platforms were added. In 2011, the IBM BlueGene/P (Eugene) was decommissioned. On March 8th of 2011 the XT4 Jaguar was decommissioned, and on September 13th of 2012 the XT5 Jaguar was transitioned to Titan. On January 8th of 2014, the Rhea system was made available to users with accounts on INCITE- or ALCC-supported projects and to users with Director's Discretion projects upon request. Shortly thereafter, on March 3rd of 2014, the Eos system was made available to all OLCF projects and prioritized as a support resource for projects running or preparing to run production and leadership capability jobs on Titan.

A large majority of the respondents in the first seven years used a Jaguar system. The percentage of Jaguar users increased each year since 2006; however since its transition to Titan, usage of the system has dropped from 2012 (97% using XT5 Jaguar PF) to 2018 (72% using Titan).

The visualization system (Lens) increased its percentage of users by 10% from 2008-2010 (20% to 30%), decreased by 5% in 2011 (25%), remained relatively stable at 26% in 2013, and was removed from the survey in 2014. The percentage of users who accessed the HPSS data storage system remained stable from 2006 to 2008 (32-34%), spiked in 2009 (38%), remained stable from 2010-2013 (35-37%), dropped to (33%) in 2014, remained stable in 2015 (34%), and dropped substantially to 23% by 2018. The Lustre/Atlas storage platform (referred to in the 2017 survey forward as Lustre/Spider for better name recognition among users) has trended downward over the last several years, and that continued in 2017 with a drop to 56% from 2016's usage rate of 67%. In 2018, usage of this system rebounded to 65%.

Refer to Table 45 for systems usage over the past decade.

Table 45. *Utilization of Compute Systems, 2006-2018*

	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018
X1E Phoenix	54%	29%	14%	-	-	-	-	-	-	-	-	-	-
XT3	73%	-	-	-	-	-	-	-	-	-	-	-	-
XT4 Jaguar		86%	92%	75%	78%	55%						-	-
XT5 Jaguar PF	-	-	-	74%	80%	94%	97%	-	-	-	-	-	-
Titan	-	-	-	-	-	-	-	85%	83%	84%	82%	78%	72%
Eos	-	-	-	-	-	-	-	-	26%	21%	21%	20%	21%
Rhea	-	-	-	-	-	-	-	-	20%	17%	24%	22%	23%
Hawk	7%	7%	-	-	-	-	-	-	-	-	-		
Data Transfer Nodes	-	-	-	-	-	-	-	-	-	40%	33%	41%	38%
HPSS	34%	32%	34%	38%	36%	37%	35%	35%	33%	34%	26%	26%	23%
IBM BlueGene/P (Eugene	-	-	-	11%	-	-	-	-	-	-	-	-	-
Development (Smoky)	-	-	-	5%	7%	5%	-	-	-	-	-		-
Lustre/Atlas	-	-	-	-	-	-	45%	47%	77%	70%	67%	-	-
Lustre/Spider	-	-	-	-	-	-	-	-	-	-	-	56%	65%
Lens	-	-	20%	29%	30%	25%	27%	26%	-	-	-	-	-

Note: Percentages total to more than 100% because users often utilize more than one system.

Support Services Utilization

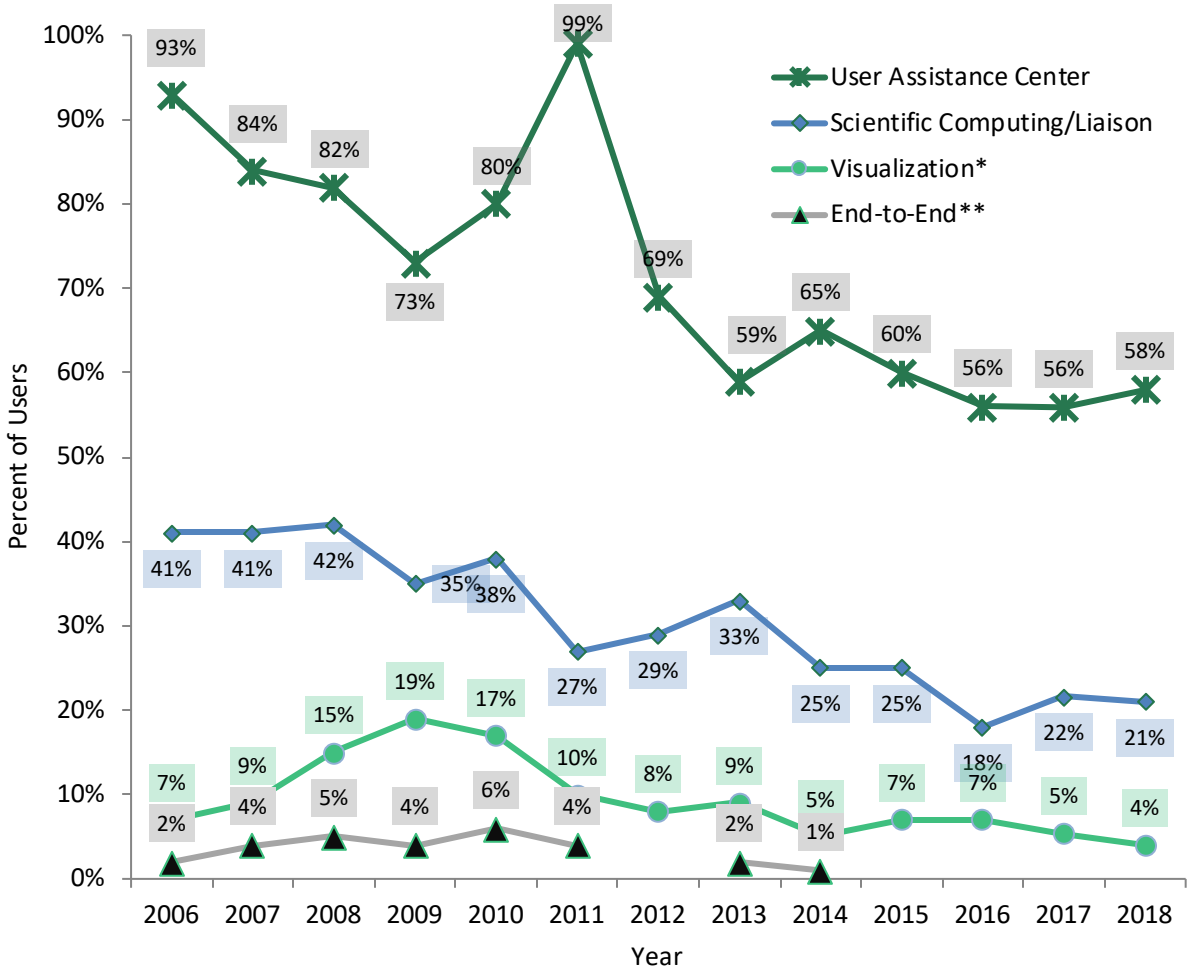

Figure 18. Utilization of support services, 2006-2018.

Notes:

*In 2017 forward, users reported their usage of "Data Analysis and Visualization Support Services."

** End-to-End was not included in the 2012, or 2015-2018 surveys.

Figure 18 shows a drop in use of the User Assistance Center from 93% to 73% between 2006 and 2009, a large increase from 2009 to 2011 (+26 percentage points to 99%), a substantial decrease back to well below the 2009 level in 2015 (60%) and a further drop in 2016 (56%), leveling back up to 58% in 2018. Use of the scientific computing/liaison has decreased from 42% in 2008 to 27% in 2011, slightly increased to 33% in 2013, and dropped again to 25% for both 2014 and 2015, and dropped again in 2016 (18%) before rising to 22% in 2017 and holding roughly steady in 2018. Use of data analysis and visualization support services saw a relatively large increase of 12 percentage points from 2006 to 2009 and a substantial decrease from 2009 to 2014 (-14 percentage points). Use of these services increased slightly in 2015 to 7%, and remained there in 2016, before dropping again to the 2014 level of 5% in 2017 and dropping slightly further in 2018.

Use of end-to-end remained relatively stable between 2006 and 2011, but was not included in the survey in 2012. In 2013, end-to-end was added back to the survey and reported to be used by 2% of users as it was originally in 2006. In 2014, end-to-end use decreased by 1 percentage point to 1%. End-to-end was dropped from the survey again from 2015 through 2018.

Satisfaction with OLCF Overall

With regard to **overall satisfaction with OLCF**, the percent of *very satisfied* respondents has shown a nearly uninterrupted trend upward since 2007 in which the proportion has more than doubled to 69% in 2017 and hovered just around that value in 2018 (Figure 19). The exceptions to this trend were moderate decreases in 2011 and 2012. The overall proportion of users indicating satisfaction (*satisfied* and *very satisfied* responses) has grown as well, from 91% in 2012 to 95-97% in each year from 2013 to 2018.

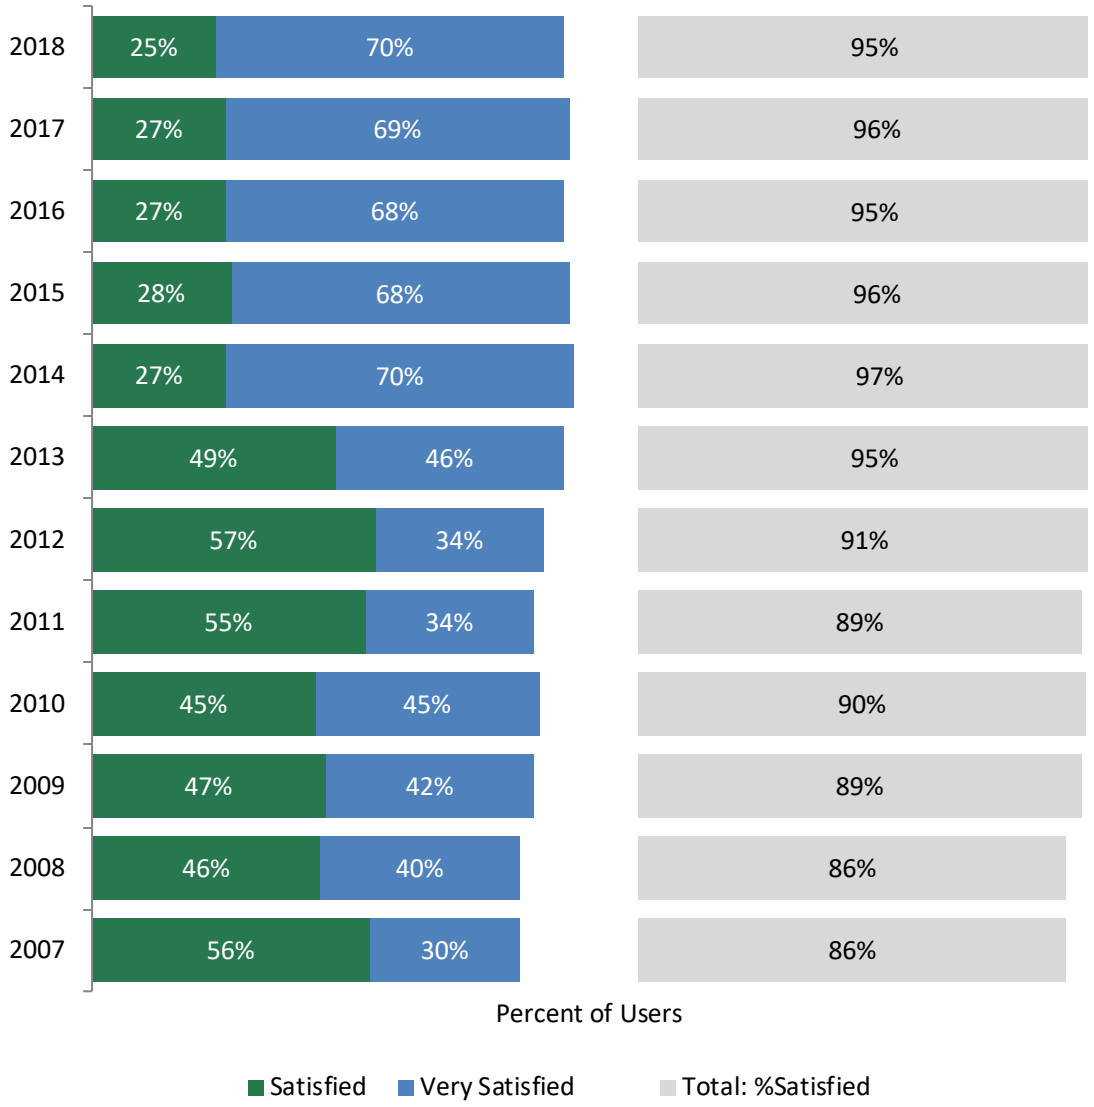

Figure 19. Proportion of respondents reporting being *satisfied* and *very satisfied* overall with OLCF, and the total of %Sat users.

Satisfaction with Compute Resources

Users were also asked to respond to a variety of questions about their opinions of the performance of the supercomputer systems of OLCF. First, respondents rated their satisfaction with the **ease of transferring data to/from OLCF**. The mean response to this question has grown since 2006 (3.8) and has been relatively stable at about 4.4 since 2017 (Figure 20).

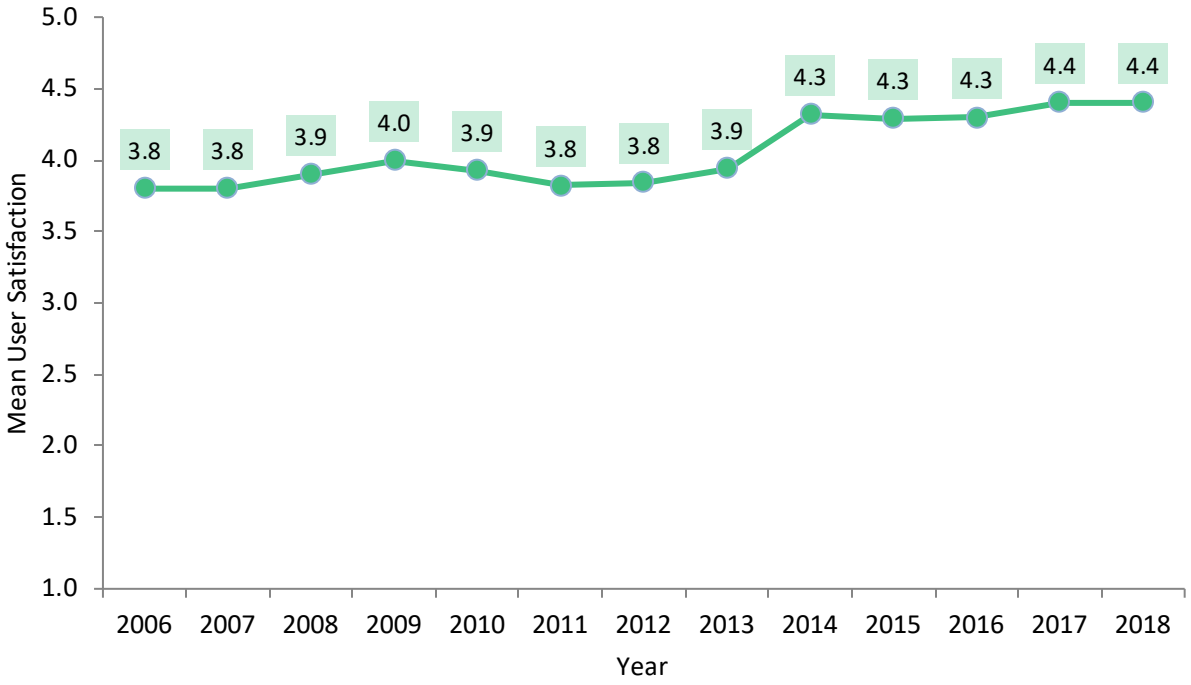

Figure 20. Ease of transferring data to and from OLCF, 2006 to 2018.
Note: Rating scale: 1 = Very Dissatisfied to 5 = Very Satisfied.

Users were also asked whether **sufficient notice was given prior to scheduled maintenance**. The responses in 2006 and 2007 were 97% and 98% "yes," respectively; however, this percentage dropped to 93% in 2009. In 2010 the survey changed, and users were asked to rate their satisfaction with the notice given prior to scheduled maintenance on a scale of 1 (*Very dissatisfied*) to 5 (*Very satisfied*) rather than a simple "yes" or "no." The mean response to this question has trended upward from 4.3 to 4.6 with a peak of 4.7 in 2014 (Figure 21).

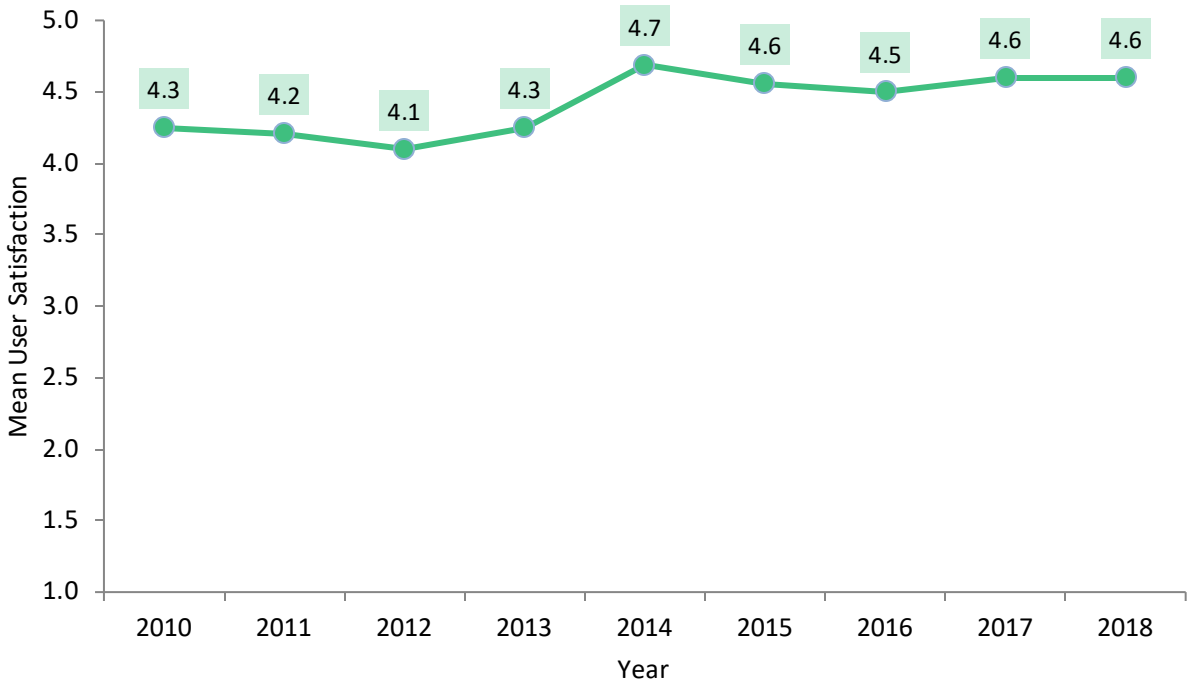

Figure 21. Sufficiency of notice given prior to scheduled maintenance, 2010 to 2018.
Note: Rating scale: 1 = Very Dissatisfied to 5 = Very Satisfied.

To gain further insight, each year through 2009, users were asked if **the level of scheduled and unanticipated outages** were acceptable (yes or no). The percentage of respondents indicating they felt the level of unanticipated outages was acceptable dropped from 68% in 2007 to 56% in 2008, but rose to 59% in 2009. Respondents who indicated they felt the frequency of scheduled outages was acceptable remained relatively stable from 2007 to 2008 (79% to 78%), but increased in 2009 (to 84%).

In 2011, the survey was changed to ask users to rate their satisfaction with the frequency of scheduled and unscheduled outages on a scale of 1 (*Very dissatisfied*) to 5 (*Very satisfied*) for each machine. User satisfaction with the frequency of scheduled Jaguar XT5 outages was unchanged from 2011 to 2012 (mean satisfaction = 3.6), while the mean satisfaction with the frequency of unscheduled (unanticipated) Jaguar XT5 outages was slightly better in 2012 (3.7) than in 2011 (3.5). Between 2013 and 2017, the mean satisfaction with the frequency of outages on Titan was relatively stable with modest increases (Figure 22), but these satisfaction scores were higher than the previous Jaguar system (see Table 45 for further history on use of systems over time). Since 2015, users have generally been just as satisfied with the frequency of *unscheduled* outages as with the frequency of anticipated, *scheduled* outages. In 2018, users were slightly less satisfied with the frequency of unanticipated, unplanned outages.

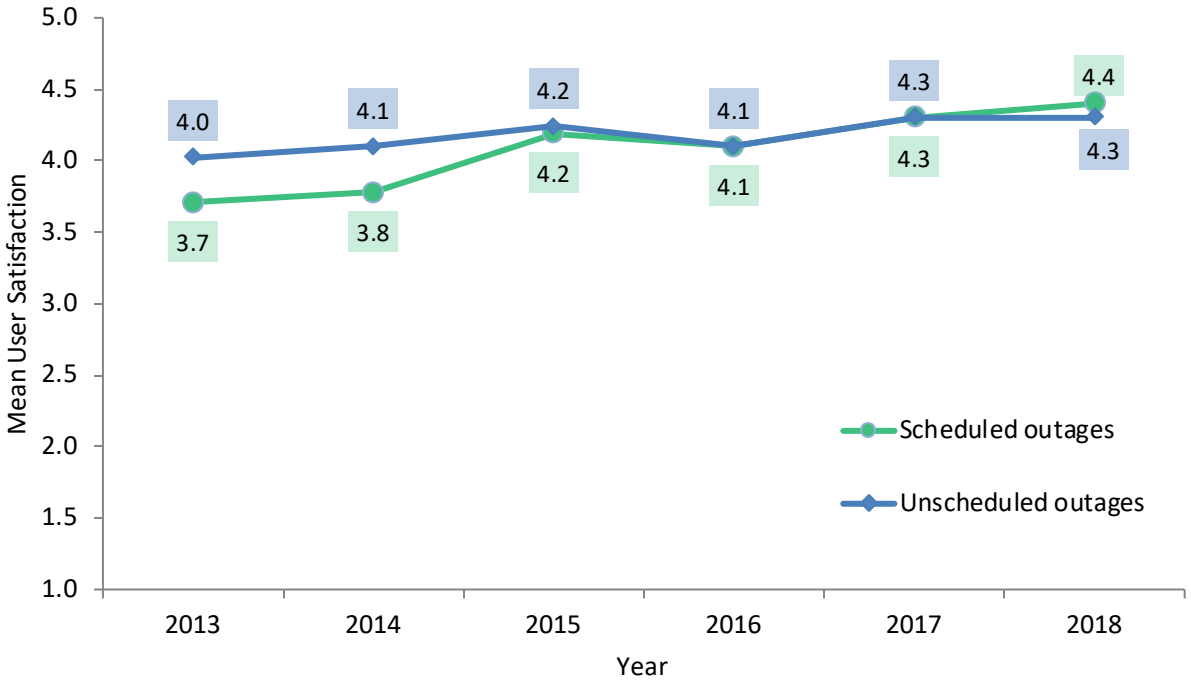

Figure 22. Mean satisfaction with the frequency of Titan scheduled and unscheduled outages, 2013-2018

Note: Rating scale: 1 = Very Dissatisfied to 5 = Very Satisfied.

Satisfaction regarding **sufficiency of project disk space** showed almost a 30-percentage point increase from 2007 to 2009 in the proportions indicating that their space was sufficient (Figure 23).

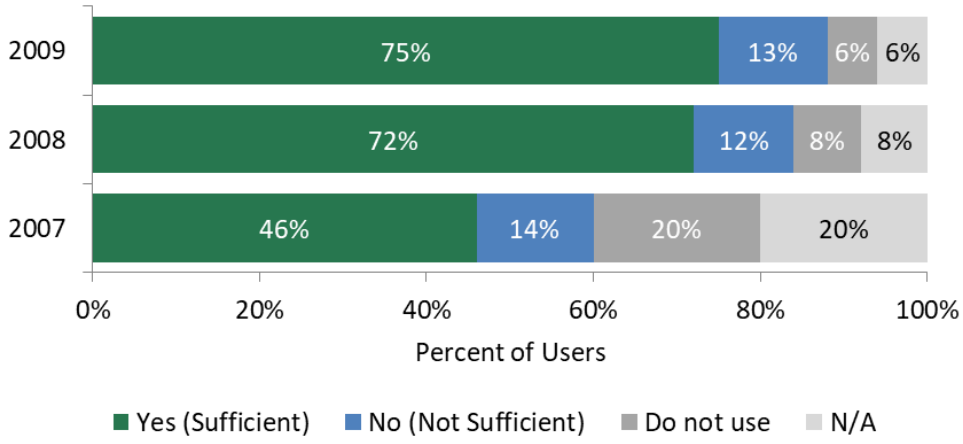

Figure 23. Reported sufficiency of the project disk space quota, 2007-2009

In 2010, this question was changed to utilize a point scale of 1 (*Very dissatisfied*) to 5 (*Very satisfied*). The mean rating remained stable from 2010 to 2013, and then increased in 2014; since then, the mean rating has been stable at this increased level (Figure 24).

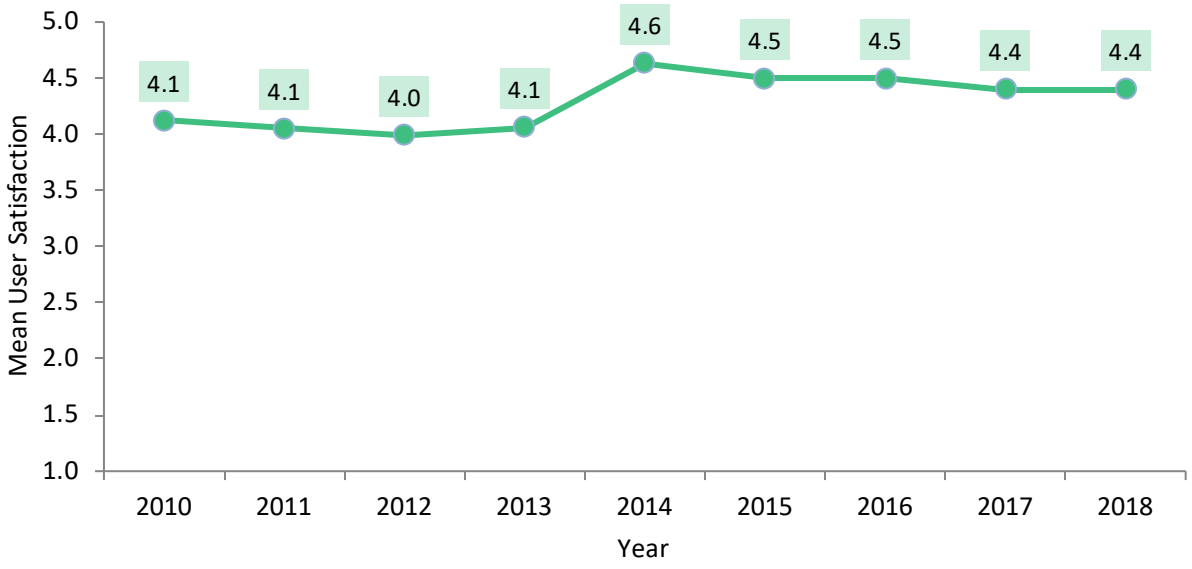

Figure 24. Mean satisfaction with the sufficiency of the project disk space quota, 2010-2018
Note: Rating scale: 1 = Very Dissatisfied to 5 = Very Satisfied.

User evaluation of the **XT3/XT4 platforms** revealed similar ratings after an overall increase from the 2006 survey (when the XT3 system was the latest platform) to the 2007 survey (XT4), but has decreased slightly since then (Table 46). The evaluation of the retired platform, the XT5, showed that users' satisfaction with the accessibility/usability of the batch queue system remained relatively stable from 2009 to 2011. In 2012, scratch disk size/performance and interface with HPSS were removed from this question (Table 47). Satisfaction ratings for accessibility/usability of the batch queue system and overall system performance remained relatively stable through the transition from XT5 to Titan.

The aspects of Titan evaluated on the survey received similar ratings in from 2014 to 2018 (Table 48). The proportions of respondents who were *satisfied* or *very satisfied* exceeded 80% for almost every rated dimension. The rare exceptions to this are highlighted in red text in Table 48, and there are some commonalities in these areas of concern from year to year. In 2018, the survey eliminated the options for *performance tools* and *debugging tools* and introduced a combined rating for *debugging and performance tools.* Ratings of *data analysis software* were also eliminated in the 2018 survey.

Table 46. *Comparison of Evaluation of XT3/XT4 Jaguar, 2006-2011*

	Mean Rating					
	2006	**2007**	**2008**	**2009**	**2010**	**2011**
Scratch disk size/performance	3.4	4.1	4.1	4.0	4.2	3.8/3.7
Interface with HPSS	4.0	4.0	3.9	4.0	3.8	3.7
Accessibility/usability of batch queue system	3.6	4.0	4.0	4.1	4.0	3.8/3.8
Throughput/turnaround time of batch queue system	3.1	3.7	3.6	3.8	3.7	3.6
Overall system performance	3.3	4.0	3.9	4.0	4.0	3.8

Note: Scratch disk size/performance and accessibility/usability of the batch queue system were each separated into two survey items in 2011.

Table 47. *Comparison of Evaluation of XT5 Jaguar PF/Titan, 2009-2013*

	Mean Rating				
	2009	**2010**	**2011**	**2012**	**2013***
Scratch disk size/performance	3.9	4.2	4.1/3.9	NA	NA
Interface with HPSS	4.0	3.8	3.8	NA	NA
Accessibility/usability of batch queue system	4.1	3.9	4.0/4.0	4.2/4.2	4.1/4.1
Throughput/turnaround time of batch queue system	3.9	3.7	3.7	NA	NA
Overall system performance	4.1	4.0	4.0	4.2	4.1

Note: Scratch disk size/performance and accessibility/usability of the batch queue system were each separated into two survey items in 2011.

*From 2013 forward users rated Titan.

Table 48. *Evaluation of Titan 2014-2018*

	2014				2015			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Batch wait time	254	4.1	0.78	83%	254	4	0.83	80%
Batch queue structure	254	4.2	0.8	86%	253	4.2	0.71	87%
Job success rate	255	4.4	0.68	91%	255	4.4	0.69	90%
Frequency of scheduled outages	258	4.1	0.78	80%	252	4.2	0.74	84%
Frequency of unscheduled (unanticipated) outages	249	4.4	0.74	86%	245	4.2	0.69	86%
Pre-2018: Performance tools	176	4.2	0.71	85%	177	4.3	0.72	87%
Pre-2018: Debugging tools	170	4.4	0.83	77%	162	4.2	0.74	82%
2018: Debugging and performance tools	--	--	--	--	--	--	--	--
Pre-2018: Data analysis software	141	4.2	0.78	79%	148	4	0.75	77%
Software/libraries	241	4.4	0.72	91%	237	4.3	0.74	89%
Programming environment	237	4.4	0.68	92%	232	4.3	0.74	88%
Scratch configuration	243	4.3	0.67	90%	239	4.3	0.68	88%
I/O performance	243	4.2	0.8	84%	242	4.2	0.71	86%
Overall satisfaction with Titan	257	4.5	0.58	96%	257	4.5	0.55	97%

Evaluation of Titan 2014-2018 (Continued)

	2016				2017				2018			
	N	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**	***N***	***M***	***SD***	**%Sat**
Batch wait time	290	3.9	0.89	72%	338	4.2	0.85	83%	287	4.2	0.77	87%
Batch queue structure	290	4	0.82	81%	336	4.2	0.74	87%	288	4.3	0.72	90%
Job success rate	293	4.2	0.88	85%	338	4.3	0.86	88%	287	4.4	0.73	92%
Frequency of scheduled outages	292	4.1	0.75	80%	335	4.3	0.7	90%	282	4.4	0.67	91%
Frequency of unscheduled (unanticipated) outages	282	4.1	0.78	81%	328	4.3	0.72	87%	281	4.3	0.71	89%
Pre-2018: Performance tools	205	4.2	0.73	86%	241	4.2	0.76	83%	--	--	--	--
Pre-2018: Debugging tools	191	4.2	0.71	85%	223	4.1	0.8	79%	--	--	--	--
2018: Debugging and performance tools	--	--	--	--	--	--	--	--	236	4.2	0.78	84%
Pre-2018: Data analysis software	171	4.1	0.79	80%	185	4.1	0.76	79%	--	--	--	--
Software/libraries	271	4.3	0.75	87%	328	4.3	0.78	87%	282	4.2	0.80	85%
Programming environment	263	4.3	0.7	90%	325	4.3	0.76	88%	278	4.3	0.72	91%
Scratch configuration	265	4.2	0.73	86%	315	4.3	0.74	88%	277	4.3	0.76	88%
I/O performance	269	4.2	0.79	84%	322	4.3	0.71	90%	273	4.3	0.79	87%
Overall satisfaction with Titan	299	4.4	0.69	93%	342	4.4	0.63	95%	298	4.5	0.63	95%

Satisfaction with Support Services

The proportion of respondents that reported making **no inquiries to the User Assistance Center** (UAC) has varied over the years, but the majority have always reported making from 1 to 5 inquiries in a year (from 50% to 63%; Figure 25). Users reporting no inquiries to the UAC have also made up a substantial proportion, ranging from a low of 22% in 2008 to a high of 34% in 2013 and 2017. The proportion making more than 20 inquiries has never been greater than 6%. UAC users have rated the service similarly over the years, with mean satisfaction ratings of various aspects ranging between 4.1 and 4.7 (Table 49). Ratings for all service dimensions in 2018 remained at levels similar to the 2014-2017 ratings

The most complete data with respect to satisfaction with the OLCF website(s) is available for timeliness of site information, the ease of finding information (i.e., site organization), the accuracy of information, and the OLCF system status information (note that in 2013, the Users' website was moved to a page within the main website). Mean satisfaction ratings have varied between 3.8 and 4.5, with slight increasing trends for all of these site dimensions (Table 50).

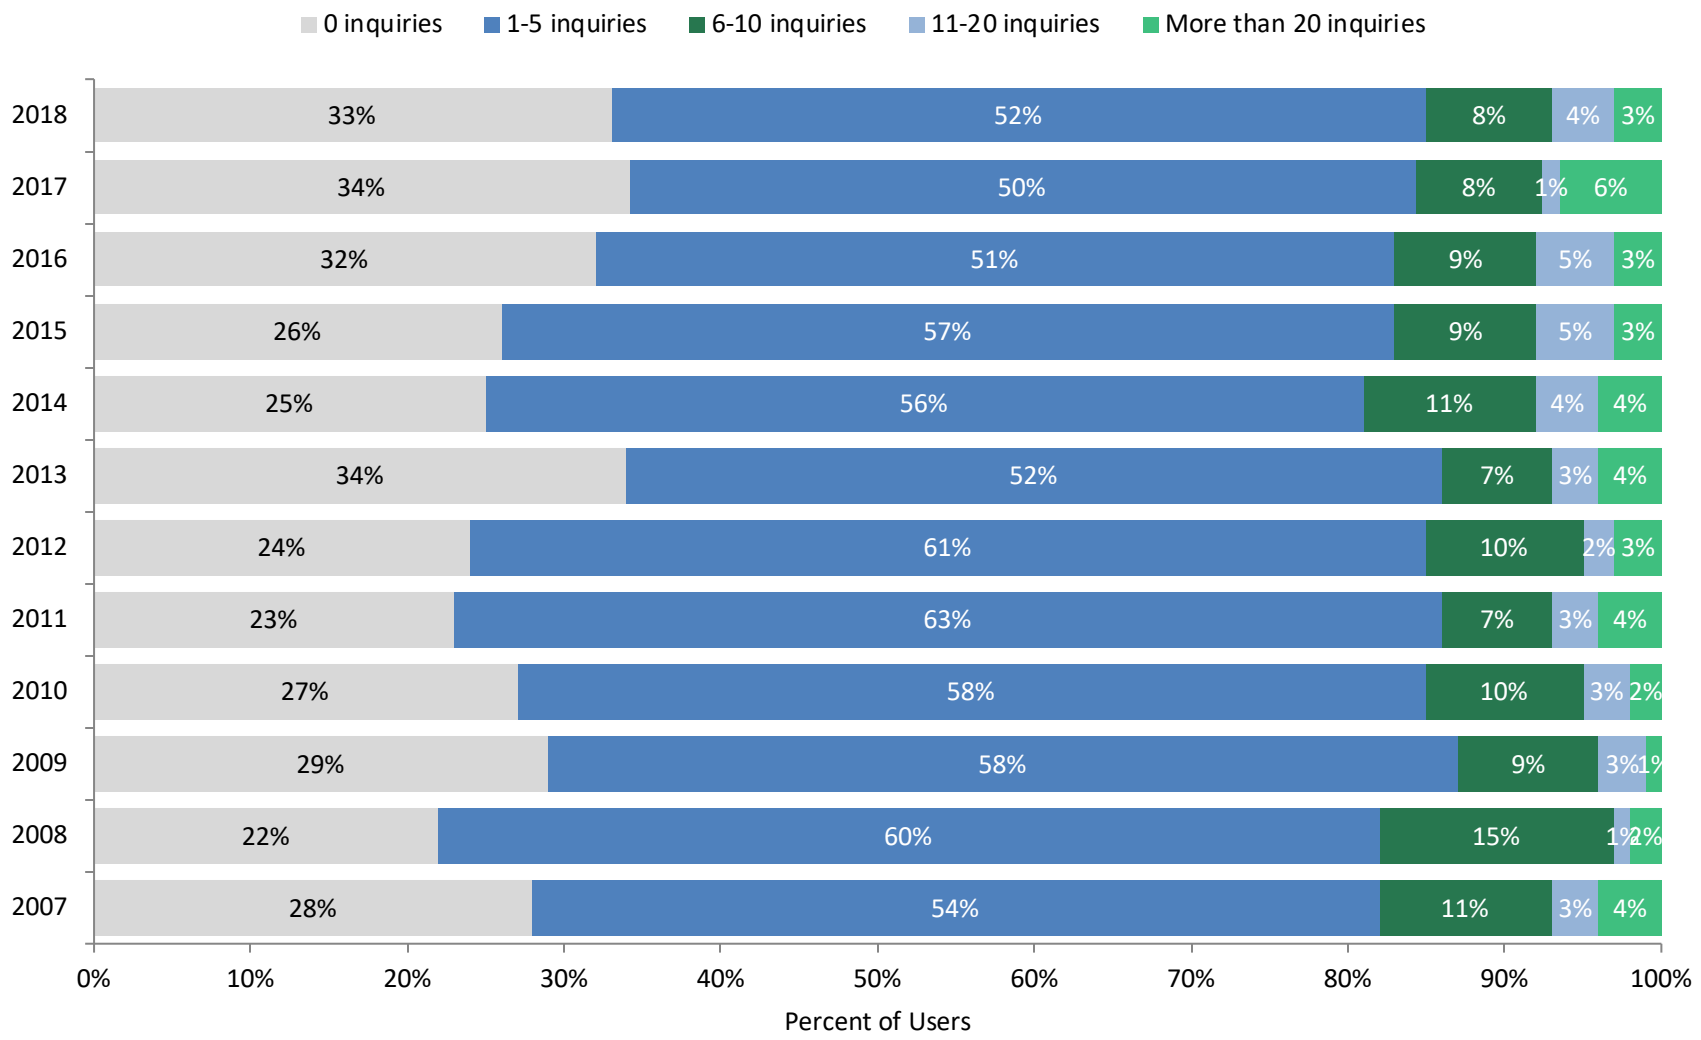

Figure 25. Proportions of respondents reporting various frequencies of User Assistance Center queries, 2007-2018

Table 49. *Mean Satisfaction Ratings of User Assistance Center (UAC) Aspects, 2007-2018*

	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018
The resolution of your queries by the UAC	4.3	4.2	4.4	4.2	4.2	4.3	-	-	-	-	-	-
Quality of technical advice/information	-	-	-	-	-	-	4.4	4.7	4.5	4.6	4.6	4.6
Response to special requests	-	-	-	-	-	-	4.4	4.6	4.6	4.6	4.5	4.5
The speed of initial response to your queries by the UAC	4.4	4.3	4.4	4.3	4.3	4.4	4.4	4.6	4.6	4.6	4.6	4.6
The speed of final resolution to your queries	4.2	4.1	4.3	4.2	4.2	4.4	4.4	4.5	4.5	4.5	4.6	4.5
Overall support from User Assistance	-	-	-	-	-	-	-	4.7	4.5	4.6	4.6	4.6

Note: In 2013, "the resolution of your queries by the UAC" was removed from the survey and replaced by "Quality of technical advice" and "Response to special requests (e.g., scheduling exceptions, software installation, etc.)" which each received average ratings of 4.4. In 2014, "Quality of technical advice" was removed from the survey and replaced by "Quality of technical information" and "Overall consulting services" was added to the survey; in 2016 forward, that was reworded to "Overall support from User Assistance."

Table 50. *Mean Satisfaction with Various Aspects of the OLCF Web Site, 2009-2017*

	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018
Timely information regarding system status	4.1	4.1	3.9	4.0[b]	4.0	4.4	4.3	4.3	4.4	4.3
Value of support information	3.9	4.9	-	-	-	-	-	-	-	-
Software inventory	4.0	3.9	-	-	-	-	-	-	-	-
Project information available on users.nccs.gov	-	3.9	-	-	-	-	-	-	-	-
OLCF system status information	-	-	4.1	4.2	4.1[c]	4.5	4.5	4.4	4.4	-
Accuracy of information	-	-	4.0/4.0[a]	4.1/4.0	4.1	4.4	4.3	4.4	4.4	4.4
Ease of finding information	4.0	3.8	3.9/3.9[a]	3.8/4.0[a]	3.9	4.2[d]	4.2	4.3	4.3	4.3

[a] In 2011 and 2012, users were asked to rate the OLCF Web site and the Users' OLCF Web site.

[b] In 2012, "timely information regarding system status" was changed to "timeliness of information."

[c] In 2013, OLCF system status information was moved to the user support section of the OLCF Web site.

[d] In 2014, "ease of finding information" was changed to "ease of navigation."

In response to the question regarding **information users would like to see on the OLCF Web site**, themes found in the open-ended responses differed quite a bit from 2007 to 2010. The most common theme in 2007, more documentation, was not found at all in the 2008 responses, but reemerged in 2009 further down the list of themes. This theme jumped back to the top of the list in 2010. The most common theme from 2008, *"tutorials and user guides,"* was not found in the 2007 responses, but was found at a lower frequency in 2009 and disappeared in 2010. Requests from the 2009 responses, but not the 2010 responses included: tips for maximizing machine usage, real-time information on job status, system specifications, libraries, bugs, other examples, FAQs, and a calendar.

Since 2011, users have been asked broadly to provide suggestions for the Web site. Sample suggestions from the most frequently occurring themes are provided below.

2011	*"Want to see my jobs and their projected start and finish times, in a separate list. Can be done via showstart..."* *"Provide predictions of whether down time is expected to last long or not."* *"Easier access on-line account to check the details of project quota usage. Maybe a login button clearly visible (e.g., next to the search button) in the home page of each system (e.g., http://www.nccs.gov/jaguar/) could help. Once I am logged in there is not much time before I have to enter my PIN code+token again. Increasing that time by a factor of 2-3 could also be helpful."*
2012	*"Documentation and FAQs provided can be improved a lot. It's often that a Google search for Cray system will take me to NERSC page where I can find some relevant information which is not available on OLCF."* *"1) Instructions for installing user specific packages/modules if they are not available by default. 2) View job information through smartphone apps or on-line"* *"Information about when machines are up or down requires you to login. It would be a lot easier to access this if I didn't have to do that. I'm not sure why that information can't be shared with everyone."* *"I don't think the existence of this site is well advertised - I think I found it by chance Google searching for some project information."*
2013	*"The most crucial information, such as the scheduling policy, should be made more prominent."* *"Please keep the online tutorials coming! These have been incredibly useful."* *"Allow a more versatile search with more specific search engines for various aspects of the system (software, hardware, scheduling, I/O, visualization, ...)"* *"Page with links to up-to-date OLCF internship opportunities and application procedures."*

	"Would be nice to see the load on the different resources from the website (maybe this is already possible)."
2014	*"Potentially an interactive chat facility, chat rooms to discuss issues with other user and admins"*
	"Recorded talks, Voice + slides would be sufficient."
2015	*"Fewer clicks to system status info."*
	"I'd like to be able to see a per-user breakdown of more job specifics, such as # of jobs run, usage by day, etc."
	"I found it hard to find what the issue was with a system and why it was down. Some MOTD updates online (or easier to find if they are there) would allow me to plan accordingly based on whether a system is expected to be down for a long time."
2016	*"If possible, add the details for running VASP."*
	"Easy to find info about system, number of nodes each system has etc."
	"More information of Lustre file system library."
2017	*"More details about software and sample batch scripts to run them."*
	"A lot of the documentation (especially for summit-dev) seems pretty sparse (although it is growing)."
	"Compute hours available on projects."
	"It would be nice to have more options on My OLCF, like restricting which user can use how many hours."
2018	*"The most important feature to me is to see machine uptime and when systems are down."*
	"cpu/gpu utilization averages of jobs."
	"I wish I could log in somewhere on the website and see a table of my projects, their current resource usage, and a list of deadlines for things like quarterly reports or quad charts. The communication of deadlines for reports and charts has generally been very sparse, and no one ever emails me to request or remind me of these requirements."
	"Just better system documentation."

Suggested Improvements for HPC Resources

Table 51 presents a summary of the types of suggestions made by users to **increase the quality of their experience using data and compute resources**. The categories with the highest relative proportion of users contributing to them are in green/bold type.

Note that because this question is open-ended, some users who do not have a suggestion may still use the available space to indicate overall satisfaction, while other satisfied users may skip the question or enter a minor suggestion. As a result, the percentage of users indicating satisfaction is not easily comparable across years and is not representative of overall satisfaction.

The table shows that just three suggestion categories have appeared (as the most often suggested) more than once, and that none of these has appeared for more than two consecutive years:

- Queuing policy/faster queues (2009, 2010, 2015),
- Reliability/stability/uptime of systems (2007, 2008), and
- Miscellaneous/Other (2012, 2013, 2016, 2018).

In short, this pattern suggests that OLCF staff members are highly responsive to user suggestions, as users' areas of concern do not persist through the next year. The *queuing policy* was consistently mentioned by a significant fraction of users until 2018, while concerns about *performance* ebb and flow from year to year. The *miscellaneous/other category* is expected to come up each year, as there are always users who provide unique feedback that does not fall easily into other categories. Other categories reflect topics that OLCF staff members should be able to address or that ongoing maintenance, upgrades, and resource additions will address over time.

Table 51. *Suggestions for How the OLCF Staff Can Improve Users' Computing Experience, 2007-2018*

	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016†	2017	2018
Satisfied	16%	12%	10%	8%	31%	15%	26%	21%	41%	33%	34%	38%
Queuing policy/faster queues	6%	12%	16%	22%	0%	9%	17%	9%*	11%	9%	11%	5%
Updates to data retention/purge policy and procedures	0%	0%	3%	6%	3%	15%	0%	5%	10%	6%	6%	7%
Make more tools available	0%	0%	0%	0%	0%	0%	4%	8%	8%	4%	18%	9%
File systems and data transfer	2%	2%	0%	3%	0%	0%	0%	13%	6%	4%	3%	7%
Performance	0%	0%	8%	5%	0%	21%	9%	10%	5%	5%	12%	7%
Miscellaneous/Other	17%	7%	5%	5%	8%	30%	26%	9%	4%	11%	4%	13%
Improve storage/memory	5%	0%	5%	8%	9%	0%	0%	3%	4%	4%	7%	2%
More documentation	4%	4%	0%	0%	3%	6%	0%	11%	4%	4%	7%	4%
Reliability/stability/uptime of systems	25%	17%	7%	16%	0%	15%	13%	3%	3%	6%	3%	5%
Training/instructional resources	0%	0%	8%	5%	0%	0%	0%	0%	2%	1%	2%	5%
Support issues	6%	6%	0%	5%	8%	15%	4%	7%	1%	1%	1%	3%
Software issues	16%	5%	6%	8%	10%	6%	0%	7%	0%	4%	5%	8%
Administrative issues	2%	5%	0%	2%	11%	0%	0%	0%	0%	7%	5%	5%
More attention to small jobs	0%	6%	0%	0%	0%	0%	0%	0%*	0%	3%	3%	1%
Improve debugging	0%	5%	5%	0%	0%	0%	0%	0%	0%	1%	n/a	2%
Allow more computing time/walltime	0%	4%	4%	10%	0%	0%	0%	0%*	0%	3%	11%	4%
Install better compilers	0%	2%	0%	0%	0%	0%	0%	0%	0%	1%	n/a	2%
GPU Resources	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	2%	4%	3%
Don't know	0%	0%	8%	0%	0%	0%	0%	0%	0%	n/a	n/a	n/a
Help with codes	0%	0%	8%	0%	0%	0%	0%	0%	0%	n/a	n/a	1%
Maintenance	0%	0%	6%	0%	0%	0%	0%	0%	0%	n/a	2%	6%
Not applicable**	0%	0%	7%	12%	13%	0%	0%	0%	n/a	n/a	n/a	n/a

Note: In 2012, this question was changed to "Please describe how the OLCF can improve your computing experience." In 2014, this question was changed to "Please describe how the OLCF can improve your experience using any of the HPC resources (i.e., Titan, Eos, Rhea, DTNs, HPSS, Lustre/Spider) and/or tell us if any additional resources are needed." *This theme was combined with 'queuing policy/faster queues' and 'allow more computing time/walltime' in 2014 as 'Review queue and walltime policies.' **"Not applicable" responses were not included in this table from 2014 forward.