

2015 Oak Ridge Leadership Computing Facility (OLCF) User Survey

Results and Analysis

PREPARED BY:

Erin M. Burr, Ph.D.
Lead Senior Evaluator
Scientific Assessment & Workforce Development

Robin T. Taylor, Ph.D.
Evaluation Specialist
Scientific Assessment & Workforce Development

PREPARED FOR:

Oak Ridge Leadership Computing Facility
Oak Ridge National Laboratory

January 2016

This document describes activities performed under contract number DE-AC05-06OR23100 between the U.S. Department of Energy and Oak Ridge Associated Universities (ORAU). ORAU manages the Oak Ridge Institute for Science and Education for the U.S. Department of Energy.

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

The Oak Ridge Institute for Science and Education (ORISE) is a U.S. Department of Energy institute focusing on scientific initiatives to research health risks from occupational hazards, assess environmental cleanup, respond to radiation medical emergencies, support national security and emergency preparedness, and educate the next generation of scientists. ORISE is managed by Oak Ridge Associated Universities (ORAU).

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.ornl.gov.

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. At the end of the eight-week survey period, 308 users completed the survey out of 873 possible respondents, giving an overall response rate of 35.3%. Findings of the survey are outlined as follows:

User Demographics

- 98% or 99% (302/306 of 308) of survey respondents reported using one or more of the following HPC resources systems: Titan (84%), Lustre/Atlas scratch file system (70%), Data Transfer Nodes (40%), HPSS archival storage system (34%), Eos (21%), and Rhea (17%).
- Survey respondents' projects were supported by Director's Discretion (53%), INCITE (42%), ALCC (31%), and Other (3%) sources such as NOAA, APOLLO, or general projects.

Overall Evaluation

- Overall ratings for the OLCF were positive, as 96% (291 of 303) reported being *satisfied* or *very satisfied* with OLCF overall. Only one user reported being *dissatisfied* and three users reported being *very dissatisfied*. On the scale of 1 (*Very dissatisfied*) to 5 (*Very satisfied*), the mean rating was 4.62, a slight decrease from 4.67 in 2014.
- With regard to overall satisfaction with OLCF over time, the percent of satisfied respondents has slowly, but steadily increased from 2007 (86%) to 2014 (97%) with a slight decrease in 2015 (96%).
- In response to an open-ended question about the best qualities of OLCF, thematic analysis of user responses identified user support and assistance (38% of responses), outstanding computing resources (33% of responses), and computing power/performance (29% of responses) as the respondents' most referenced best qualities.

User Assistance Evaluation

- For support services used, 60% of the 308 respondents reported using the User Assistance Center (UAC), followed by 25% using the Account Management services, 25% with assigned INCITE Scientific Computing/Liaisons, and 7% using data analysis and visualization support services.
- Overall satisfaction with the user support services provided by the OLCF was high with an average response of 4.54 ($SD = 0.68$) on a rating scale of 1 (*Very dissatisfied*) to 5 (*Very satisfied*). Mean ratings to questions of overall satisfaction with various aspects of user assistance ranged from 4.50 to 4.58.

Training and Education

- The majority of users who responded prefer online documentation ($n = 228$, 80%) or online training ($n = 179$, 63%).
- Out of the 286 who indicated the most convenient time to attend a training event, the majority of respondents indicated no preference ($n = 154$, 54%), and roughly one-third ($n = 84$, 29%) indicated the summer was the most convenient time.

- Users were asked to provide satisfaction ratings for their overall satisfaction with OLCF training and five specific training aspects. Satisfaction ratings for overall satisfaction was positive ($M = 4.31$, $SD = 0.69$, 89% satisfaction). The lowest satisfaction rating for training aspects was for monthly user conference calls ($M = 4.14$, $SD = 0.75$, 78% satisfaction).
- A total of 41 users indicated that they had participated in OLCF training during the 2015 calendar year. Of these users, 31 (76%) said they would recommend attending a future OLCF training event in person, 9 (22%) said maybe, and 1 (2%) person said no.
- When asked “What training topic(s) would you like to see offered in the future?” 16 topics were suggested by more than one user. The top three suggestions provided include: “GPUs” (28%), “Parallelization” (10%), and “Use of performance tools/performance monitoring” (10%).

OLCF Communications

- 94% of respondents who answered a question about their overall satisfaction with communications from the OLCF rated it as *satisfied* or *very satisfied*. The least positive ratings, which are still high, were for the aspect of communications, announcements on the OLCF website ($M = 4.32$, $SD = 0.72$, 88% satisfaction).
- Users responded that they feel adequately informed of OLCF changes (96%, 276 out of 287), events (97%, 278 out of 287), and current issues (95%, 273 out of 287).

OLCF Web Sites

- 98% (281 of 287) of respondents indicated that they had visited the <http://olcf.ornl.gov> Web site. Of these users 38% (106) indicated that they visit the site once a week or more, 7 of whom indicated that they visit the site every day.
- Overall, respondents indicated they were moderately satisfied with the OLCF Web site ($M = 4.27$, $SD = 0.64$). The user support aspects with the highest rating was system user guides ($M = 4.42$, $SD = 0.66$, 94% satisfaction).

OLCF Systems

- Overall ratings for the OLCF systems were positive, ranging from 91%-97% of users rating the systems with either *satisfied* or *very satisfied* on the scale of 1 (*Very dissatisfied*) to 5 (*Very satisfied*). Mean ratings for the systems ranged from 4.30 ($SD = 0.76$; Data Transfer Nodes) to 4.52 ($SD = 0.62$; Eos).
- Regarding maintenance and outages, 95% indicated sufficient notice is given prior to scheduled maintenance. The majority also indicated that they are *satisfied* or *very satisfied* with project disk space (93%), the bandwidth offered by the OLCF (92%), and the ease of transferring data to/from the OLCF (85%).

Data Analysis, Visualization, and Workflow

- When asked about the quality of technical assistance with data analysis and visualization, the average rating was 4.33 ($SD = 0.86$, 86% Satisfied).
- Users were moderately satisfied with the aspects of data analysis and visualization support services with overall means ranging from 4.24 to 4.33 with satisfaction all at 86%.
- Eighty-three percent of users ($n = 238$ out of 287) indicated that they had not taken advantage of the OLCF cross-platform submission capabilities in their workflow with 49 (17%) indicating that they had.

Introduction

A general online survey of all users of the Oak Ridge Leadership Computing Facility (OLCF) at Oak Ridge National Laboratory (ORNL) in 2015 was launched on October 6, 2015 and remained open for completion through December 2, 2015. Information was collected about the various users, the user experience with OLCF, and the OLCF support capabilities. Attitudes and opinions on the performance, availability, and possible improvements for OLCF and its staff were also solicited.

The survey was created with contributions from OLCF staff and evaluators with ORAU's Assessment and Evaluation team. The survey was hosted online by ORAU.

ORAU staff sent e-mails to an OLCF user distribution list provided by OLCF staff. Users on the list included all users who logged into an OLCF system between 01/01/2015 (the day OLCF staff pulled the last user list for the previous survey) through 10/01/2015. OLCF staff and vendors were removed from the list. Principal investigators (PIs) of all the projects that do not have an account were added as the OLCF is interested in capturing their feedback. A total of 308 users completed the survey (out of 873 possible respondents, a response rate of 35.3%).

The resulting data are discussed in the next two major sections (i.e., data analysis and findings, and longitudinal comparisons of user responses). The third section provides recommendations. The appendices at the end of the report include a copy of the survey, email invitations and reminders to complete the survey, a survey response log, and results from the open-ended survey questions.

Data Analysis and Findings

Data were analyzed using both quantitative and qualitative methods. The two fundamental goals that drove the collection and subsequent analysis were to understand the types of users and to understand their needs and preferences with the systems thus, where appropriate, data is broken out by users' PI status and/or project allocation. Descriptive statistics are used to understand users' ratings. In addition to means and standard deviations, the percentage of respondents indicating a 4 (*Satisfied*) or 5 (*Very satisfied*) were calculated to provide additional details for user satisfaction. Qualitative data was analyzed using grounded theory. Examples of the top themes are presented. Please note that percentages of response categories may add up to more than 100% of the number of users due to users providing multiple explanations within a comment.

User Demographics

A total of 877 names were indicated as users of the OLCF during 2015; however, four e-mails provided were deemed invalid resulting in a total user population of 873.

The overall response rate for survey completion is 35.3% ($n = 308$ out of 873). Half ($n = 157$, 51%) of survey respondents indicated using the OLCF for more than 2 years, while 22% ($n = 69$) indicated using the OLCF between 1 and 2 years, and 27% ($n = 82$) indicated having used OLCF for less than 1 year. The majority of survey respondents using OLCF were affiliated with a university or DOE/Laboratory/Government facility; see Table 1 for the occupational affiliation categories of survey respondents.

Table 1. *Occupational Affiliation of OLCF Survey Users*

Occupational affiliation	<i>n</i>	%
University	135	44%
DOE/Laboratory/Government	108	35%
Industry	29	9%
Other	19	6%
Foreign	17	6%

Users of the OLCF are categorized according to the following project allocations:

- 1) **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;
- 2) **DD**. The National Center for Computational Sciences' Director's Discretion program is designed to give new researchers an opportunity to carry out a program of scalability and productivity enhancements to their scientific codes;

- 3) **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; and
- 4) **Other.** Other programs includes the National Oceanic and Atmospheric Administration (NOAA), APOLLO, and General projects.

The counts and percentages of OLCF users by the different project types are shown in Table 2. Percentages add to more than 100% as users could have multiple project allocations. Figure 1 illustrates the breakdown of survey respondents according to the combination of OLCF projects indicated. Out of the 308 responses, a total of 228 (74%) users indicate having a single project-type allocation.

The Principal Investigator (PI) status was provided by OLCF staff. In 2015, 93 (10.7%) users were identified with PI status. Almost two-thirds of users with PI status completed the survey ($n = 61$), representing 19.8% of the overall survey respondents – see Table 3.

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

Project(s) allocations	Survey respondents = 308		OLCF 2015 users = 873	
	Count	Percentage	Count	Percentage
DD	164	53%	435	50%
INCITE	129	42%	351	40%
ALCC	97	31%	222	25%
Other	9	3%	31	4%

Note. Percentages add to more than 100% as users can have multiple projects.

User Classification Combinations

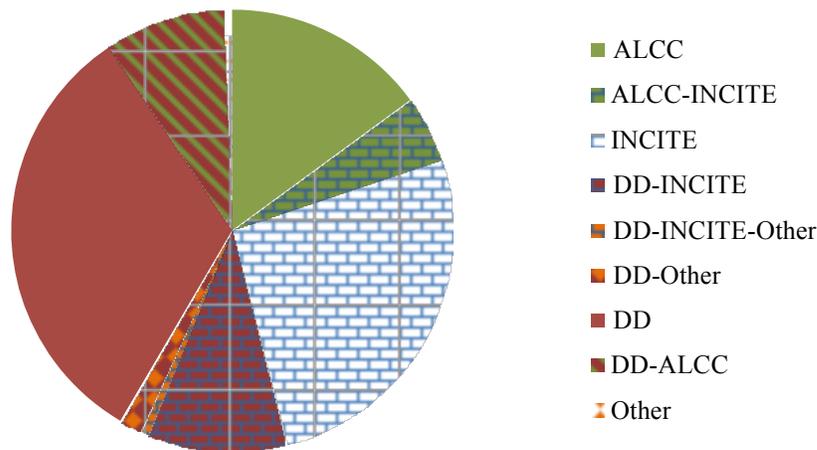


Figure 1. Breakdown of project allocations.

Table 3. Numbers of users with PI Status for Survey Respondents and OLCF 2015 Users

	Survey respondents = 308		OLCF 2015 Users = 873	
	<i>n</i>	%	<i>n</i>	%
INCITE PI only	8	3%	13	1%
DD PI only	44	14%	65	7%
ALCC PI only	7	2%	14	2%
Multiple PI allocations (ALCC+DD or DD+INCITE)	2	1%	2	>1%
No PI status	247	80%	779	89%

OLCF users were asked to indicate which OLCF HPC resources they utilized during the 2015 calendar year. The largest number of users indicated using Titan ($n = 257$, 84%), Lustre/Atlas ($n = 212$, 70%), and Data Transfer Nodes ($n = 120$, 40%). For 2015, the smallest number of respondents indicated utilizing Eos ($n = 64$, 21%) and Rhea ($n = 52$, 17%). Table 4 indicates the overall utilization of HPC resources as well as provides the breakdown of usage by project allocations and PI status.

The frequency in which survey respondents reported using OLCF support services during 2015 is presented in Table 5. The largest number of responses indicated utilizing the User Assistance Center ($n = 171$ of 287, 60%). Approximately one-quarter of respondents indicated using INCITE Scientific Computing Liaison ($n = 73$ or 25%) and/or Account management services ($n = 72$, 25%). Forty-nine respondents indicated they had taken advantage of the OLCF cross-platform (17%), 41 (14%) participated in an OLCF training event during 2015, and 21 (7%) utilized data analysis and visualization.

Table 4. HPC Resources Utilized During 2015 by Project Allocation

HPC Resources	PI Status (n = 61)		INCITE (n = 128 or 127)		DD (n = 163 or 160)		ALCC (n = 96)		Other (n = 9)		Total (N = 306 or 302)	
	n	%	n	%	n	n	n	%	n	%	n	%
Titan	57	93%	108	84%	142	257	257	86%	8	89%	257	84%
Eos	11	18%	28	22%	26	64	64	27%	1	11%	64	21%
Rhea	16	26%	26	20%	34	52	52	15%	1	11%	52	17%
Data Transfer Nodes	28	46%	55	43%*	70	120	120	39%	5	56%	120	40%*
HPSS	27	44%	53	42%*	53	103	103	33%	1	11%	102	34%*
Lustre/Atlas	43	70%	96	76%*	116	212	212	71%	7	78%	212	70%*

Note. Users add up to more than 100% because some use more than one system. * denotes the denominator is based on the lower sample size (n = 302 for Totals, n = 127 for INCITE, and n = 160 for DD).

Table 5. Support Services Used During 2015 by Project Allocation

Systems	PI Status (n = 59 or 58)		INCITE (n = 123)		DD (n = 151)		ALCC (n = 94 or 93)		Other (n = 7)		Total (N = 287 or 286)	
	n	%	n	%	n	n	%	%	n	%	n	%
User Assistance Center	41	69%	67	54%	107	71%	61	65%	7	100%	717	60%
Account Management	25	42%	32	26%	45	30%	22	23%	4	57%	72	25%
INCITE Scientific Computing Liaisons	11	19%	65	53%	18	12%	21	22%	2	29%	73	25%
OLCF Training	15	26%	17	14%	29	19%	13	14%	3	43%	41	14%
Data Analysis and Visualization Support Services	5	8%	8	7%	18	12%	4	4%	0	0%	21	7%
OLCF Cross-Platform Submission Capabilities	13	22%	20	16%	29	19%	13	14%	3	43%	49	17%

Note. Users add up to more than 100% because some use more than one service. * denotes the denominator is based on the lower sample size (n = 286 for Totals, n = 93 for ALCC, and n = 58 for PI status).

Overall Evaluation

Users were asked to rate their overall satisfaction with the OLCF, OLCF compute resources, OLCF data resources, and OLCF Support Services using a scale of 1 (*Very dissatisfied*) and 5 (*Very satisfied*). Tables 6 and 7 include descriptive statistics of ratings by PI status and overall and by project allocation, respectively. Overall satisfaction was generally very positive with means at or above 4.40: OLCF ($M = 4.62$, $SD = 0.65$, satisfaction = 96%), OLCF compute resources ($M = 4.59$, $SD = 0.65$, satisfaction = 96%), OLCF data resources ($M = 4.44$, $SD = 0.76$, satisfaction = 91%), and OLCF services (support, training, communications, website, etc. [$M = 4.45$, $SD = 0.79$, satisfaction = 89%]) – Table 6. Figure 2 provides the percentage of respondents indicating each of the satisfaction rating scales for OLCF, OLCF compute resources, OLCF data resources, and OLCF services.

When comparing the percentage of respondents indicating *satisfied* or *very satisfied*, PIs indicated similar overall levels of satisfaction as non-PIs. However, overall means were generally higher for PIs than for non-PIs. Overall satisfaction findings were also similar across project allocations with one exception – “Other” users were generally less satisfied with OLCF support services than INCITE, DD, and ALCC users.

Tables 6 and 7 also include specific ratings for overall satisfaction of specific compute resources (Titan, Eos, and Rhea), data resources (Data Transfer Nodes, HPSS, and Lustre/Atlas) and support services (consulting services, account services, INCITE liaison, communications, training, and website).

Table 6. *Descriptive Statistics for Overall Satisfaction by PI-Status*

	PI-Status				Non PI-Status				Total			
	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat
OLCF	61	4.69	0.70	97%	242	4.59	0.68	96%	303	4.62	0.65	96%
OLCF compute resources	60	4.62	0.69	98%	240	4.57	0.68	96%	300	4.59	0.65	96%
-Titan	57	4.53	0.57	96%	200	4.45	0.55	98%	257	4.47	0.55	97%
-Eos	11	4.73	0.47	100%	53	4.47	0.64	92%	64	4.52	0.62	94%
-Rhea	16	4.63	0.62	94%	36	4.42	0.60	94%	52	4.48	0.61	94%
OLCF data resources	51	4.53	0.81	92%	185	4.39	0.78	90%	236	4.44	0.76	91%
-Data Transfer Nodes	28	4.36	0.78	89%	94	4.28	0.75	90%	122	4.30	0.76	90%
-HPSS	27	4.41	0.84	85%	76	4.37	0.67	92%	103	4.38	0.72	90%
-Lustre/Atlas	43	4.49	0.70	93%	170	4.32	0.77	90%	213	4.36	0.76	91%
Support services	61	4.51	0.79	92%	237	4.41	0.82	88%	298	4.45	0.79	89%
-Consulting services	39	4.67	0.48	100%	126	4.50	0.72	90%	165	4.54	0.68	93%
-Account services	25	4.60	0.65	92%	47	4.60	0.58	96%	72	4.60	0.60	94%
-INCITE liaison	10	4.80	0.42	100%	48	4.60	0.71	92%	58	4.64	0.67	93%
-Communications	57	4.53	0.60	95%	225	4.42	0.66	93%	282	4.44	0.65	94%
-Training	50	4.44	0.64	92%	197	4.27	0.70	89%	247	4.31	0.69	89%
-Website	55	4.44	0.57	96%	217	4.34	0.65	94%	272	4.36	0.63	94%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 7. Descriptive Statistics for Overall Satisfaction by Project Allocation

	INCITE				DD				ALCC				Other			
	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat*												
OLCF	127	4.71	0.52	98%	162	4.57	0.69	94%	96	4.64	0.62	98%	9	4.44	0.73	89%
Compute resources	126	4.63	0.63	95%	160	4.56	0.64	96%	96	4.51	0.70	96%	9	4.33	1.00	89%
-Titan	108	4.48	0.56	97%	142	4.45	0.55	97%	83	4.43	0.52	99%	8	4.13	0.64	88%
-Eos	28	4.50	0.64	93%	26	4.54	0.65	92%	26	4.46	0.65	92%	---	---	---	---
-Rhea	26	4.50	0.71	88%	34	4.44	0.61	94%	14	4.50	0.52	100%	---	---	---	---
Data resources	95	4.51	0.73	91%	132	4.45	0.70	92%	76	4.38	0.78	91%	9	4.22	0.67	89%
-Data Transfer Nodes	57	4.44	0.63	96%	71	4.35	0.78	90%	37	4.24	0.89	86%	5	4.40	0.55	100%
-HPSS	53	4.32	0.73	89%	53	4.43	0.67	91%	33	4.58	0.71	94%	---	---	---	---
-Lustre/Atlas	97	4.33	0.80	92%	116	4.36	0.70	91%	69	4.38	0.79	90%	7	4.14	0.69	86%
Support services	125	4.52	0.74	89%	159	4.43	0.81	87%	95	4.39	0.75	89%	9	4.22	0.83	78%
-Consulting services	64	4.55	0.75	92%	103	4.53	0.64	92%	58	4.60	0.56	97%	7	4.29	0.76	86%
-Accounts team	32	4.56	0.56	97%	45	4.64	0.61	93%	22	4.77	0.43	100%	4	4.25	0.50	100%
-INCITE liaison	51	4.71	0.54	96%	15	4.53	0.83	93%	19	4.68	0.58	95%	---	---	---	---
-Communications	122	4.49	0.65	96%	149	4.41	0.66	92%	92	4.49	0.60	95%	7	4.29	0.49	100%
-Training	103	4.33	0.65	92%	134	4.33	0.68	88%	82	4.33	0.74	90%	7	4.43	0.53	100%
-Website	121	4.36	0.59	96%	139	4.32	0.62	92%	89	4.44	0.67	96%	7	4.29	0.49	100%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

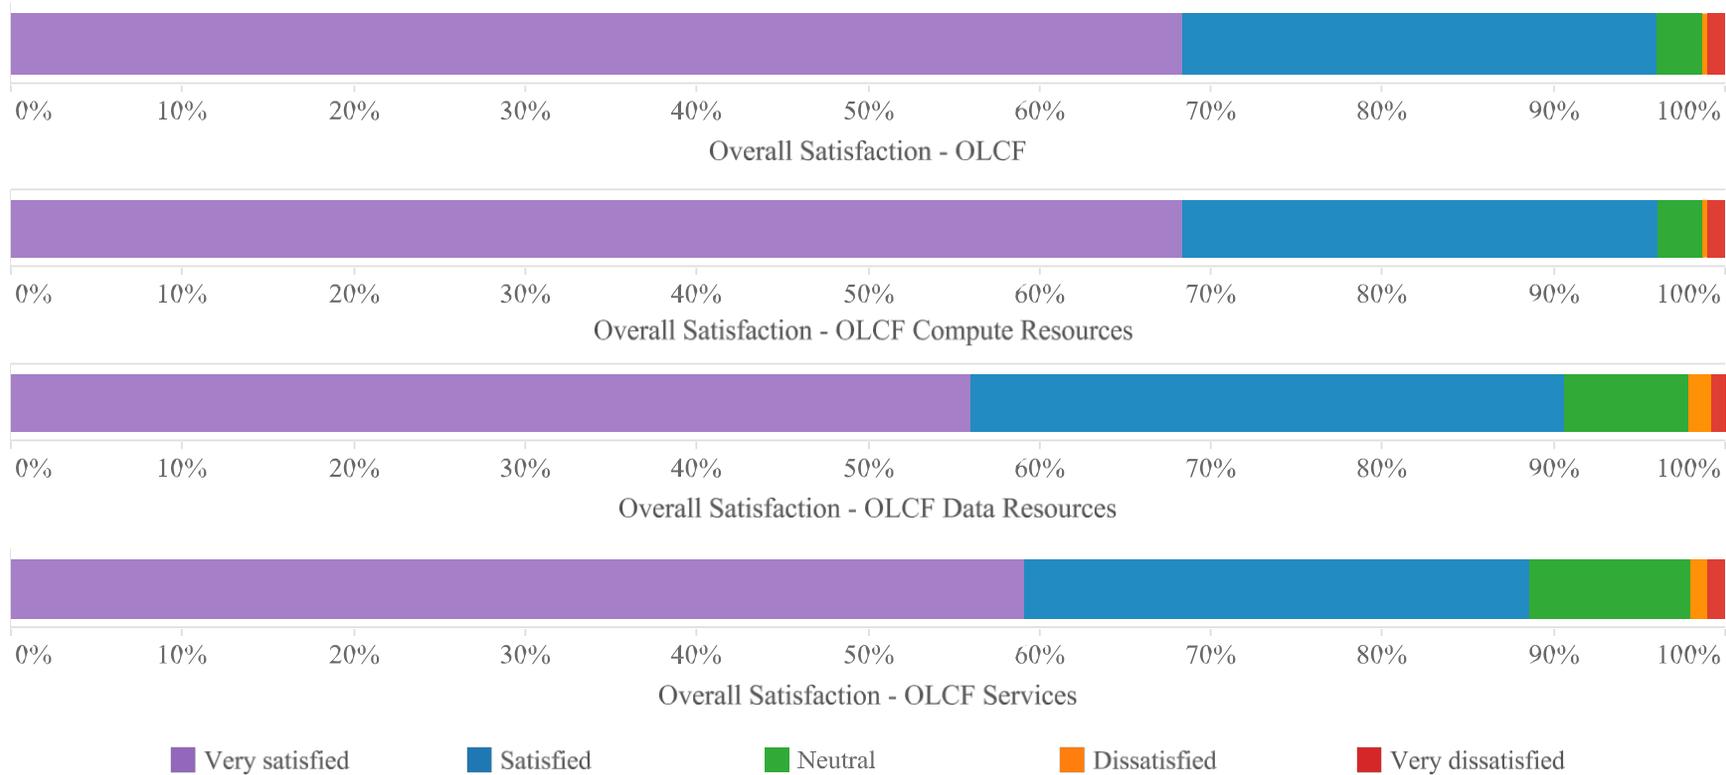


Figure 2. Overall satisfaction ratings for OLCF, OLCF compute resources, OLCF Data Resources, and OLCF Support Services by percentages indicating each satisfaction response.

If a user rated any of the aspects of their overall satisfaction with the OLCF, its resources and support services with *very dissatisfied* or *dissatisfied*, they were asked to explain their rating. Out of 13 users who gave a dissatisfied rating on at least one of the four satisfaction items, 7 users provided explanations for their dissatisfaction; however no themes were present in the feedback received.

Additionally, users were asked to provide feedback of what they believed to be the best qualities of the OLCF. Thematic analysis of user responses identified user support and assistance (found in 39% of responses), the outstanding quality of the computational resource (found in 32% of responses), and computing power/performance (found in 29% of responses) as the top reported best qualities of the OLCF (Table 8). Some respondent comments about these qualities included:

User assistance staff/outstanding support

“A very active and engaged support staff - fast, knowledgeable staff who are willing to work with special requests.”

“OLCF seems to be exceptionally well organized in terms of communication with its users and technical support.”

“The staff. They always try to make it work to get our science done! And super friendly!”

Outstanding computing resources

“The unique computing resources of the OLCF are outstanding.”

“The computing resources allow completion of my work.”

“The OLCF provides a valuable computational resource and need expertise for using the resource.”

Table 8. *Best Qualities of OLCF*

	<i>N</i> = 260	%
User assistance staff/outstanding support	100	38%
Outstanding computing resource	85	33%
Computing power/performance	75	29%
Stability/Reliability	26	10%
Documentation/website	13	5%
Low wait time for system use/good scheduling policy	12	5%
Availability/Ease of access and use	46	18%
Training/Training materials	2	1%
Software libraries/software development	4	2%
Data management and transfer	2	1%
Overall satisfaction	6	2%

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

OLCF HPC Resources

Users who used OLCF HPC Resources were asked to provide satisfaction ratings for the following aspects: a) notice given prior to scheduled maintenance, b) project disk space, c) ease of transferring data to/from the OLCF, and bandwidth offered by the OLCF. Table 9 includes descriptive ratings by PI status and overall, and Table 10 includes descriptive ratings by project allocation. The highest mean satisfaction rating was for notice given prior to scheduled maintenance ($M = 4.56$, $SD = .67$, 95% satisfied), and the lowest overall mean rating was for “ease of transferring data to/from the OLCF” ($M = 4.29$, $SD = .87$, 85% satisfied).

If a user rated any of the aspects of OLCF compute and data resources with *very dissatisfied* or *dissatisfied* they were asked to explain their rating (Table 11). Of 21 who were *very dissatisfied* or *dissatisfied* with one or more aspects of the OLCF compute and data resources, 16 provided explanations for their dissatisfaction. The most common feedback dealt with “Security features make transferring data to/from OLCF difficult” (38%) and “Transferring data to/from OLCF is time-consuming” (31%). Sample comments for these themes include:

Security features make transferring data to/from OLCF difficult

“Involving grid credentials within Globus make the transfers more complicated. Not having Globus from the file system to HPSS is a negative.”

“The use of OSG certificates for Globus transfer introduces a high barrier to entry, making the process very user *un*friendly. It is also highly non-standard for large compute facilities.”

Transferring data to/from OLCF is time-consuming

“There are issues with transferring large data files to/from OLCF. bbcp not always works, sometimes the speed is lower than regular sftp. Globus online is just unusable....”

“GridFTP through the openscience grid project was tedious and took quite a while to be approved in order to transfer my data. Additionally, simply rsyncing a few files between my local machine and a transfer node frequently failed without any clear cause (from the log reports that I generated at the time)”

Users were asked to indicate their opinion regarding the performance of OLCF compute and data resources compared to the previous year. Of users with more than 2 years as a user, 60 (38%) noted overall improvements of the OLCF compute and data systems, and 96 (61%) of users noted the performance of OLCF compute and data systems is about the same as it was the last year. Only one user with greater than 2 years as a user (0.6%) noted an overall decrease in performance of the OLCF compute and data systems indicating, “There is consistent planned downtime during the workweek of compute and storage resources.” For users with 1 to 2 years of time as a user ($n = 69$), no one indicated overall decreases, 17 (25%) indicated improvements, and 52 (75%) indicated performance was about the same.

Table 9. Descriptive Statistics for Aspects of the OLCF HPC Compute and Data Resources by PI Status and Overall Totals

Overall satisfaction with...	PI Status				Non-PI Status				Total			
	N	M	SD	%Sat*	N	M	SD	%Sat*	N	M	SD	%Sat*
Notice given prior to scheduled maintenance	61	4.56	0.67	93%	243	4.56	0.67	95%	304	4.56	0.67	95%
Project disk space	60	4.58	0.62	93%	241	4.48	0.70	93%	301	4.50	0.68	93%
Ease of transferring data to/from the OLCF	58	4.33	0.80	86%	235	4.28	0.89	85%	293	4.29	0.87	85%
Bandwidth offered by the OLCF	59	4.54	0.65	95%	235	4.46	0.68	92%	294	4.47	0.67	93%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 10. Descriptive Statistics for Aspects of the OLCF HPC Compute and Data Resources by Project Allocation

Overall satisfaction with...	INCITE				DD				ALCC				Other			
	N	M	SD	%Sat*	N	M	SD	%Sat*	N	M	SD	%Sat*	N	M	SD	%Sat*
Notice given prior to scheduled maintenance	129	4.59	0.71	96%	161	4.50	0.70	94%	96	4.58	0.66	96%	9	4.44	0.53	100%
Project disk space	128	4.47	0.71	93%	159	4.50	0.65	93%	95	4.53	0.70	93%	8	4.13	0.64	88%
Ease of transferring data to/from the OLCF	121	4.31	0.86	88%	155	4.27	0.86	86%	94	4.27	0.94	81%	8	3.63	1.30	63%
Bandwidth offered by the OLCF	122	4.56	0.63	94%	155	4.45	0.67	93%	94	4.43	0.74	89%	9	3.89	0.78	67%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 11. *Users' Explanations for Dissatisfaction with OLCF Compute and Data Resources*

	<i>N</i> = 16	%
Security features make transferring data to/from OLCF difficult	6	38%
Transferring data to/from OLCF is time-consuming	5	31%
Scheduled maintenance occurs with little to no notice	2	13%
Errors experienced while transferring data to/from OLCF	2	13%
Miscellaneous	5	31%

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

Titan

A total of 257 (84% of survey respondents) indicated using Titan during the 2015 calendar year. For users who utilized Titan during this time, they were asked to provide satisfaction ratings to multiple aspects of the Titan system. Descriptive statistics of ratings are shown in Table 12 by PI status and overall total, and in Table 13 by project allocations. Users' overall satisfaction with Titan was high ($M = 4.47$, $SD = 0.55$, 97% satisfaction). Aspects of Titan that appear to have lower satisfaction ratings (as evidenced by their overall mean scores and percent satisfied) include debugging tools, data analysis software, frequency of scheduled outages, and batch wait time. Means for these aspects were between 4.08 and 4.19 with percentages ranging from 77% to 84%.

If a user rated any of the aspects of Titan with *very dissatisfied* or *dissatisfied* they were asked to explain their rating. Of 27 users who indicated at least one dissatisfied response, 22 provided explanations for their dissatisfaction (Table 14). The most common explanations dealt with "Batch wait time is too long" (36%) and "Performance too slow" (18%). Sample comments for these themes include:

Batch wait time is too long

"Batch queue structure should be modified to give access to different types of projects with equal opportunity. The batch wait time is too long, sometime days. It has improved in 2015, but lots can be done to enhance performance. Many of the codes do not get updated although the related modules are updated. Therefore, the old code doesn't run on Titan. The only solution then is to compile in own directory... User needs to compile in their home directory to solve this problem."

"It often has long waiting period. Sometimes a job does not run after 10 days in the queue. A substantial fraction of my jobs were terminated accidentally."

Performance too slow

"We have noticed huge difference in timings when reading file using Fortran default access type with the Cray compiler only. Using access stream solved this problem."

"Scratch disk performance is often very slow."

Table 12. Descriptive Statistics for Satisfaction Ratings of Titan 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	57	4.30	0.65	93%	197	4.03	0.87	77%	254	4.09	0.83	80%
Batch queue structure	57	4.30	0.60	93%	196	4.17	0.74	85%	253	4.20	0.71	87%
Job success rate	57	4.35	0.69	88%	198	4.40	0.70	91%	255	4.39	0.69	90%
Frequency of scheduled outages	57	4.18	0.71	82%	195	4.19	0.75	85%	252	4.19	0.74	84%
Frequency of unscheduled (unanticipated) outages	55	4.27	0.65	89%	190	4.23	0.70	85%	245	4.24	0.69	86%
Performance tools	38	4.39	0.64	92%	139	4.24	0.74	86%	177	4.27	0.72	87%
Debugging tools	36	4.36	0.64	92%	126	4.10	0.76	79%	162	4.15	0.74	82%
Data analysis software	36	4.22	0.72	83%	112	4.04	0.76	75%	148	4.08	0.75	77%
Software/libraries	49	4.41	0.61	94%	188	4.21	0.77	87%	237	4.25	0.74	89%
Programming environment	49	4.22	0.69	86%	183	4.26	0.75	89%	232	4.25	0.74	88%
Scratch configuration	51	4.20	0.78	82%	188	4.29	0.66	90%	239	4.27	0.68	88%
I/O performance	55	4.20	0.78	82%	187	4.24	0.69	88%	242	4.23	0.71	86%
Overall satisfaction with Titan	57	4.53	0.57	96%	200	4.45	0.55	98%	257	4.47	0.55	97%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 13. Descriptive Statistics for Satisfaction Ratings of Titan by Project Allocation

	INCITE				DD				ALCC				Other			
	N	M	SD	%Sat	N	M	SD	%Sat	N	M	SD	%Sat	N	M	SD	%Sat
Batch wait time	107	4.10	0.81	82%	140	4.08	0.81	80%	82	4.15	0.83	82%	8	3.88	0.35	88%
Batch queue structure	107	4.20	0.72	86%	139	4.17	0.71	86%	81	4.31	0.65	90%	8	4.13	0.64	88%
Job success rate	108	4.41	0.68	91%	141	4.39	0.67	91%	81	4.40	0.70	90%	8	4.38	0.52	100%
Frequency of scheduled outages	107	4.21	0.75	85%	139	4.13	0.80	81%	81	4.23	0.69	88%	8	3.63	1.30	63%
Frequency of unscheduled (unanticipated) outages	104	4.29	0.66	88%	135	4.17	0.73	82%	80	4.34	0.62	93%	7	4.14	0.90	71%
Performance tools	71	4.21	0.79	85%	102	4.25	0.68	86%	55	4.36	0.59	95%	6	4.00	0.89	67%
Debugging tools	69	4.13	0.82	80%	90	4.16	0.70	82%	51	4.20	0.66	86%	7	4.00	0.82	71%
Data analysis software	61	4.00	0.82	70%	86	4.10	0.72	79%	40	4.13	0.72	80%	5	4.20	0.45	100%
Software/libraries	99	4.36	0.61	93%	132	4.19	0.78	87%	79	4.22	0.76	87%	8	4.13	0.64	88%
Programming environment	100	4.29	0.66	91%	129	4.17	0.77	84%	72	4.24	0.74	88%	8	4.13	0.64	88%
Scratch configuration	103	4.25	0.68	88%	132	4.23	0.72	86%	78	4.29	0.70	91%	7	4.14	0.69	86%
I/O performance	100	4.18	0.67	87%	134	4.19	0.74	84%	76	4.25	0.66	88%	7	4.00	0.58	86%
Overall satisfaction with Titan	108	4.48	0.56	97%	142	4.45	0.55	97%	83	4.43	0.52	99%	8	4.13	0.64	88%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 14. *Users' Explanations for Dissatisfaction with Titan*

	<i>N</i> = 22	%
Batch wait time is too long	8	36%
Performance too slow	4	18%
Not enough wall time	2	9%
Scheduled outages occur too frequently	2	9%
Need for updated versions of CUDA to be installed	2	9%
Jobs failing	2	9%
Miscellaneous	8	36%

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

Eos

A total of 64 (21%) respondents indicated using Eos during the 2015 calendar year. For users who utilized Eos during this time, they were asked to provide satisfaction ratings to multiple aspects of the Eos system, Table 15 and Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 16. Users' overall satisfaction with Eos was high ($M = 4.52$, $SD = 0.62$, 94% satisfaction). The highest rated aspect of the Eos system was for Job success rate ($M = 4.56$, $SD = 0.64$, 92% satisfied). The lowest rated aspects of the Eos were debugging ($M = 4.15$, $SD = 0.73$, 80% satisfied), data analysis software ($M = 4.15$, $SD = 0.77$, 78% satisfied), and batch wait time ($M = 4.16$, $SD = 0.89$, 80% satisfied).

If a user rated any of the aspects of Eos with *very dissatisfied* or *dissatisfied* they were asked to explain their rating. Of four users who gave at least one dissatisfaction rating with Eos, two provided explanations:

"I found it very hard to get jobs of greater than 4,000 cores through the queue in any reasonable amount of time."

"The job scheduling policy makes it difficult for me to utilize my time and the processors efficiently. Long wait times in the queue (on the order of 3 days). Weekly outages seems excessive; it seems like a monthly outage would be more typical.)"

Table 15. *Descriptive Statistics for Satisfaction Ratings of Eos by PI Status and Overall Totals*

	PI-Status				Non-PI Status				Total			
	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat
Batch wait time	11	4.55	0.69	91%	53	4.08	0.92	77%	64	4.16	0.89	80%
Batch queue structure	11	4.45	0.82	82%	53	4.30	0.70	87%	64	4.33	0.71	86%
Job success rate	11	4.82	0.40	100%	53	4.51	0.67	91%	64	4.56	0.64	92%
Frequency of scheduled outages	11	4.73	0.47	100%	52	4.21	0.72	87%	63	4.30	0.71	89%
Frequency of unscheduled (unanticipated) outages	10	4.70	0.48	100%	52	4.23	0.73	83%	62	4.31	0.71	85%
Performance tools	7	4.43	0.79	86%	42	4.19	0.71	83%	49	4.22	0.71	84%
Debugging tools	7	4.43	0.79	86%	39	4.10	0.72	79%	46	4.15	0.73	80%
Data analysis software	7	4.43	0.79	86%	33	4.09	0.77	76%	40	4.15	0.77	78%
Software/libraries	8	4.38	0.74	88%	51	4.22	0.70	84%	59	4.24	0.70	85%
Programming environment	7	4.43	0.79	86%	48	4.23	0.69	85%	55	4.25	0.70	85%
Scratch configuration	10	4.50	0.71	90%	49	4.31	0.65	90%	59	4.34	0.66	90%
I/O performance	11	4.18	0.60	91%	50	4.32	0.65	90%	61	4.30	0.64	90%
Overall satisfaction with Eos	11	4.73	0.47	100%	53	4.47	0.64	92%	64	4.52	0.62	94%

Note. %Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 16. *Descriptive Statistics for Satisfaction Ratings of Eos by Project Allocation*

	INCITE				DD				ALCC			
	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat
Batch wait time	28	4.21	0.83	82%	26	4.23	0.95	81%	26	4.00	0.98	77%
Batch queue structure	28	4.39	0.69	89%	26	4.27	0.83	77%	26	4.31	0.62	92%
Job success rate	28	4.50	0.69	89%	26	4.58	0.64	92%	26	4.58	0.64	92%
Frequency of scheduled outages	27	4.33	0.62	93%	26	4.35	0.85	85%	26	4.19	0.80	85%
Frequency of unscheduled (unanticipated) outages	27	4.30	0.72	85%	26	4.42	0.70	88%	25	4.32	0.69	88%
Performance tools	23	4.30	0.70	87%	18	4.33	0.77	83%	20	4.15	0.67	85%
Debugging tools	22	4.27	0.70	86%	17	4.24	0.75	82%	18	4.06	0.73	78%
Data analysis software	17	4.29	0.77	82%	15	4.20	0.86	73%	15	4.13	0.74	80%
Software/libraries	26	4.31	0.68	88%	24	4.17	0.82	75%	23	4.22	0.67	87%
Programming environment	25	4.28	0.74	84%	22	4.32	0.72	86%	21	4.19	0.68	86%
Scratch configuration	27	4.30	0.67	89%	24	4.50	0.59	96%	23	4.26	0.69	87%
I/O performance	26	4.27	0.72	85%	25	4.32	0.63	92%	25	4.24	0.60	92%
Overall satisfaction with Eos	28	4.50	0.64	93%	26	4.54	0.65	92%	26	4.46	0.65	92%

Note. %Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

When asked why they have made little use or not utilized Eos, the most common reason provided was “No need for Eos” (28%; Table 17). This was followed by “Titan fits better/allocation is for Titan” (22%) and “No access/allocation of hours” (16%). Sample comments for these themes include:

Not need for Eos

“My project is a data-only project. I utilize the HPSS systems and the data transfer nodes.”

“It has not been pertinent for my work.”

“Not applicable to my work.”

Titan fits better/allocation is for Titan

“I like Titan too much and don't have a chance to try. Titan is too good.”

“I was specifically tasked with getting our code to build on Titan.”

“The Titan resource was deemed ideal for my project.”

No access/allocation of hours

“Not included in allocation (as far as I know).”

“I do not believe we have an allocation on Eos.”

Table 17. *Users’ Explanations for Having Made Little or No Use of Eos*

	<i>N</i> = 115	%
No need for Eos	29	28%
Titan fits better/allocation is for Titan	23	22%
No access/allocation of hours	17	16%
No	14	13%
Other systems are more appropriate	12	11%
Something I need isn't installed/limitations	7	7%
Not familiar with Eos	7	7%
Have not had time	4	4%
Plan to use Eos/just starting to use	2	2%
Rhea fits better/allocation is for Rhea	2	2%

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

Rhea

A total of 52 (17%) respondents indicated using Rhea during the 2015 calendar year. Rhea users during were asked to provide satisfaction ratings to multiple aspects of the Rhea system, see Table 18 and *Table 19*. Users' overall satisfaction with Rhea was high ($M = 4.48$, $SD = 0.61$, 94% satisfied). The highest rated aspects of the Rhea system were for batch wait time, job success rate, and programming environment, and batch queue structure. The lowest rated aspects of Rhea were for frequency of scheduled outages, data analysis software, and scratch configuration.

If a user rated any of the aspects of Rhea with *very dissatisfied* or *dissatisfied* they were asked to explain their rating. Of eight users indicating dissatisfaction for at least one aspect of Rhea, seven users provided explanations for their dissatisfaction. Three (43%) responses referenced "slow performance". The rest of the feedback varied. Comments on slow performance include:

"Only use Rhea to run back-end for VisIT. Any wait beyond a few minutes for a single node greatly diminishes usefulness for on-demand examination of simulation data."

"Rhea is very slow, but its point is not performance."

When asked why they have made little use or not utilized Rhea, the most common reason provided was "No need for Rhea" (32%; Table 20). This was followed by "Titan fits better/allocation is for Titan" (18%) and "No reason" (15%). Sample comments for these themes include:

No need for Rhea

"The pre- and post-processing requirements of our jobs are rather minimal."

"We do most of our processing and visualization locally."

"In the past year I used most of the resource for computation and did not have suitable data for post-processing on Rhea."

Titan fits better/allocation is for Titan

"Titan/EOS is most important for my work."

"Was primarily coding for GPUs on Titan"

"Use for visualization only. Ran analysis jobs directly on Titan, no point in reconfiguring makefiles for analysis for a different programming environment"

No reason

"No specific reason, I use Eos."

"There is no specific reason."

"None"

Table 18. *Descriptive Statistics for Satisfaction Ratings of Rhea by PI Status and Overall Totals*

	PI-Status				Non-PI Status				Total			
	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat
Batch wait time	16	4.50	0.82	94%	34	4.47	0.83	94%	50	4.48	0.81	94%
Batch queue structure	16	4.56	0.63	94%	34	4.35	0.77	97%	50	4.42	0.73	96%
Job success rate	16	4.56	0.81	81%	36	4.42	0.84	92%	52	4.46	0.83	88%
Frequency of scheduled outages	15	4.20	0.86	73%	35	4.11	0.93	83%	50	4.14	0.90	80%
Frequency of unscheduled (unanticipated) outages	15	4.27	0.80	80%	35	4.26	0.85	89%	50	4.26	0.83	86%
Performance tools	10	4.60	0.70	90%	23	4.13	0.81	83%	33	4.27	0.80	85%
Debugging tools	10	4.50	0.71	90%	21	4.19	0.81	86%	31	4.29	0.78	87%
Data analysis software	14	4.57	0.65	93%	24	3.92	0.88	75%	38	4.16	0.86	82%
Software/libraries	15	4.73	0.46	100%	33	4.12	0.78	82%	48	4.31	0.75	88%
Programming environment	15	4.67	0.49	100%	31	4.13	0.88	90%	46	4.30	0.81	93%
Scratch configuration	15	4.40	0.91	87%	34	4.06	0.98	82%	49	4.16	0.96	84%
I/O performance	15	4.40	0.74	87%	36	4.14	0.90	83%	51	4.22	0.86	84%
Overall satisfaction with Rhea	16	4.63	0.62	94%	36	4.42	0.60	94%	52	4.48	0.61	94%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 19. *Descriptive Statistics for Satisfaction Ratings of Rhea by Project Allocation*

	INCITE				DD				ALCC			
	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat
Batch wait time	24	4.46	1.06	88%	33	4.42	0.71	94%	14	4.57	0.51	100%
Batch queue structure	24	4.42	0.93	92%	33	4.42	0.56	97%	14	4.50	0.52	100%
Job success rate	26	4.50	0.95	88%	34	4.47	0.71	88%	14	4.57	0.65	93%
Frequency of scheduled outages	24	4.17	1.01	79%	34	4.15	0.82	79%	14	4.29	0.91	86%
Frequency of unscheduled (unanticipated) outages	24	4.25	0.99	83%	34	4.26	0.71	85%	14	4.50	0.65	93%
Performance tools	14	4.14	0.95	79%	21	4.33	0.73	86%	10	4.60	0.52	100%
Debugging tools	12	4.25	0.97	83%	22	4.27	0.70	86%	9	4.67	0.50	100%
Data analysis software	17	4.06	0.97	82%	25	4.28	0.74	84%	10	4.40	0.70	90%
Software/libraries	23	4.35	0.83	87%	32	4.31	0.64	91%	13	4.46	0.52	100%
Programming environment	23	4.17	1.03	87%	30	4.40	0.50	100%	11	4.55	0.52	100%
Scratch configuration	25	4.04	1.14	76%	32	4.19	0.82	88%	12	4.25	1.14	83%
I/O performance	25	4.16	1.03	80%	33	4.27	0.76	88%	14	4.29	0.73	86%
Overall satisfaction with Rhea	26	4.50	0.71	88%	34	4.44	0.61	94%	14	4.50	0.52	100%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 20. *Users' Explanations for Having Made Little or No Use of Rhea*

	N = 119	%
No need for Rhea	38	32%
Titan fits better/allocation is for Titan	21	18%
No reason	18	15%
No access/allocation of hours	10	8%
Something I need isn't installed/limitations	9	8%
Not familiar with Rhea	8	7%
Other systems are more appropriate	7	6%
Plan to use Rhea/just starting to use	5	4%
Eos fits better/allocation is for Eos	4	3%
Have not had time	3	3%
Other responses	1	1%

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

Data Transfer Nodes (DTNs)

Users were asked to indicate their satisfaction with using Data Transfer Nodes during the 2015 calendar year. The overall mean satisfaction rating for 122 users providing a rating was $M = 4.30$ ($SD = 0.76$, %Sat = 90%). Means varied slightly by project allocations ($M_{DD} = 4.35$, $M_{INCITE} = 4.44$, and $M_{ALCC} = 4.24$).

If a user rated any of the aspects of the Data Transfer Nodes with *very dissatisfied* or *dissatisfied* they were asked to explain their rating. Three of the five users (60%) who were dissatisfied with the Data Transfer Nodes provided explanations for their dissatisfaction. The explanation provided by all three respondents was “Difficulty experienced in transferring data”. Feedback included:

“The dtn nodes are great when it comes to data sets below a TB. Unfortunately, our raw data is in the size of PB and our funding agencies (EU, national) require us to archive scientific data for about 10 years at home. If we could establish a tape exchange or similar to transfer selected large scale data sets this would be very helpful”.

“Frequent issues trying to transfer large amount of data from Titan to local storage across the data transfer nodes (gridFTP would not function across ssh as it was specified).”

“Difficulty in transferring data across different DOE supercomputing facilities.”

HPSS

A total of 102 (34%) indicated using HPSS during the 2015 calendar year. For users who utilized HPSS during this time, they were asked to provide satisfaction ratings to multiple aspects of the HPSS system, see Table 21 for PI status and total ratings and Table 22 for ratings by project allocation. Users' overall satisfaction with HPSS was moderately high ($M = 4.38$, $SD = 0.72$, 91% satisfaction). The highest rated aspects of the HPSS system were reliability (data integrity [$M = 4.53$, $SD = 0.68$, 95% satisfied]) and ability to store files ($M = 4.49$, $SD = 0.70$, 95% satisfied). The lowest rated aspects of HPSS were for time to retrieve files, time to store files, and hsi interface. Means for these aspects were between 4.23 and 4.27 with percentages ranging from 80% to 86%.

If a user rated any of the aspects of the HPSS with *very dissatisfied* or *dissatisfied* they were asked to explain their rating (Table 23). All eight users who provided a dissatisfaction rating for at least one aspect of HPSS provided explanations for their dissatisfaction. The top two explanations provided were "hsi/htar is problematic" (63%) and "File access issues" (25%). Sample comments for these themes include:

hsi/htar is problematic

"The hsi interface produces stderr and stdout to the same channel and does not produce an error return code when there is an issue (like file not uploaded). The combination of these two items means that there is no way using the supplied tool to tell if your data storage or retrieval task has completed successfully. The htar tool is limited in the file size that it can store. It is already unacceptably small limit, and will continue to get worse."

"The hsi interface is so dated and doesn't include so many of the basics folks expect when at a terminal prompt. On the plus side, it's highly scriptable and lightweight (unlike Globus in my experience... it seems HPSS interfaces are a quite difficult problem)."

File access issues

"Again, as indicated in my previous responses, I would like ongoing access to my HPSS data in between projects."

"Retrieving files has been really slow at times."

Table 21. *Descriptive Statistics for Satisfaction Ratings of HPSS by PI Status and Overall Totals*

	PI-Status				Non-PI Status				Total			
	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat
hsi interface	24	4.25	0.90	79%	74	4.23	0.87	80%	98	4.23	0.87	80%
htar interface	20	4.30	0.80	90%	63	4.32	0.82	84%	83	4.31	0.81	86%
Ability to store files	27	4.59	0.57	96%	76	4.45	0.74	95%	103	4.49	0.70	95%
Ability to retrieve files	27	4.41	0.69	89%	74	4.38	0.73	92%	101	4.39	0.72	91%
Reliability (data integrity)	27	4.56	0.64	93%	70	4.51	0.70	96%	97	4.53	0.68	95%
Time to store files	27	4.33	0.73	85%	75	4.25	0.81	87%	102	4.27	0.79	86%
Time to retrieve files	27	4.19	0.92	74%	73	4.27	0.79	86%	100	4.25	0.82	83%
Frequency of scheduled outages	26	4.19	0.69	85%	70	4.37	0.64	91%	96	4.32	0.66	90%
Frequency of unscheduled (unanticipated) outages	26	4.27	0.72	85%	69	4.42	0.63	93%	95	4.38	0.66	91%
Overall satisfaction with HPSS	27	4.41	0.84	85%	76	4.37	0.67	92%	103	4.38	0.72	90%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 22. *Descriptive Statistics for Satisfaction Ratings of HPSS by Project Allocation*

	INCITE				DD				ALCC			
	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat
hsi interface	53	4.13	0.94	77%	51	4.33	0.82	82%	29	4.24	0.87	79%
htar interface	46	4.26	0.91	83%	42	4.36	0.76	88%	20	4.30	0.86	85%
Ability to store files	53	4.43	0.67	94%	53	4.53	0.58	96%	33	4.64	0.78	97%
Ability to retrieve files	51	4.37	0.66	90%	53	4.42	0.63	92%	32	4.59	0.84	94%
Reliability (data integrity)	48	4.58	0.58	96%	49	4.47	0.65	92%	32	4.63	0.79	97%
Time to store files	52	4.21	0.78	83%	53	4.28	0.72	85%	33	4.61	0.79	97%
Time to retrieve files	51	4.20	0.80	80%	52	4.25	0.76	81%	32	4.59	0.80	97%
Frequency of scheduled outages	48	4.35	0.60	94%	50	4.22	0.68	86%	33	4.39	0.66	91%
Frequency of unscheduled (unanticipated) outages	47	4.38	0.61	94%	50	4.30	0.71	86%	33	4.52	0.62	94%
Overall satisfaction with HPSS	53	4.32	0.73	89%	53	4.43	0.67	91%	33	4.58	0.71	94%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 23. *Users' Explanations for Dissatisfaction with HPSS*

	<i>N</i> = 8	%
hsi/htar is problematic	5	63%
File access issues	2	25%
Miscellaneous	1	13%

Lustre/Atlas Scratch Filesystem

Seventy percent (212 of 302) users indicated using Lustre/Atlas Scratch Filesystem during the 2015 calendar year. For users who utilized the system during this time, they were asked to provide satisfaction ratings to multiple aspects of the system – see Table 24 for ratings by PI status and overall total and Table 25 for ratings by project allocation. Users' overall satisfaction with Lustre/Atlas was moderately high ($M = 4.36$, $SD = 0.76$, 91% satisfaction). The highest rated aspect of the system was size ($M = 4.45$, $SD = 0.71$, 92% satisfied). The lowest rated aspect of Lustre/Atlas was for file and directory operations ($M = 4.27$, $SD = 0.83$, 86% satisfied).

If a user rated any of the aspects of the Lustre/Atlas with *very dissatisfied* or *dissatisfied* they were asked to explain their rating (Table 26). Of 15 users who indicated dissatisfaction with at least one aspect of the Lustre/Atlas Scratch Filesystem, 14 provided explanations. The two most common explanations provided were “Slow performance” (64%) and “Issues with data purging” (29%). Sample comments for these themes include:

Slow performance

“Compiling on these file systems is very slow. I believe that paths are tied to projects, but it is a pain as project IDs change even though the science project remains the same.”

“Part of this is how vis software does I/O, but performance variability on Lustre kills us. If performance was 50% slower but reliably fast on such supercomputers, the design of visualization software may slowly adapt to have a more interactive model.”

Issues with data purging

“Lustre purge policy and performance are unsatisfactory.”

“All data was deleted automatically after a very short period of time which is very inconvenient. I have to transfer files back and forth frequently.”

Table 24. Descriptive Statistics for Satisfaction Ratings of Lustre/Atlas Scratch Filesystem by PI Status and Overall Totals

	PI Status				Non-PI Status				Total			
	N	M	SD	%Sat	N	M	SD	%Sat	N	M	SD	%Sat
Size	43	4.56	0.67	95%	170	4.42	0.72	91%	213	4.45	0.71	92%
I/O bandwidth	43	4.44	0.63	93%	166	4.31	0.75	89%	209	4.33	0.73	89%
File and directory operations	42	4.33	0.87	83%	169	4.25	0.82	86%	211	4.27	0.83	86%
Reliability	43	4.60	0.58	95%	169	4.38	0.75	89%	212	4.42	0.72	91%
Frequency of scheduled outages	43	4.44	0.70	88%	164	4.32	0.74	89%	207	4.34	0.73	89%
Frequency of unscheduled (unanticipated) outages	43	4.44	0.70	88%	163	4.29	0.79	85%	206	4.32	0.77	85%
Overall satisfaction with the Lustre/Atlas filesystem	43	4.49	0.70	93%	170	4.32	0.77	90%	213	4.36	0.76	91%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 25. Descriptive Statistics for Satisfaction Ratings of Lustre/Atlas Scratch Filesystem by Project Allocation

	INCITE				DD				ALCC				Other			
	N	M	SD	%Sat	N	M	SD	%Sat	N	M	SD	%Sat	N	M	SD	%Sat
Size	97	4.48	0.69	94%	116	4.34	0.75	87%	69	4.64	0.51	99%	7	4.14	0.69	86%
I/O bandwidth	96	4.40	0.69	93%	113	4.27	0.73	87%	68	4.35	0.73	91%	7	4.14	0.69	86%
File and directory operations	95	4.28	0.82	88%	116	4.24	0.82	84%	69	4.28	0.84	84%	7	4.14	0.69	86%
Reliability	96	4.44	0.72	92%	116	4.42	0.67	91%	68	4.40	0.83	84%	7	4.29	0.49	100%
Frequency of scheduled outages	96	4.36	0.73	93%	113	4.28	0.77	84%	66	4.39	0.70	91%	7	4.00	1.00	86%
Frequency of unscheduled (unanticipated) outages	95	4.38	0.76	91%	112	4.27	0.77	82%	66	4.35	0.77	85%	7	4.14	0.90	71%
Overall satisfaction with the Lustre/Atlas filesystem	97	4.33	0.80	92%	116	4.36	0.70	91%	69	4.38	0.79	90%	7	4.14	0.69	86%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 26. *Users' Explanations for Dissatisfaction with Lustre/Atlas*

	N = 14	%
Slow performance	9	64%
Issues with data purging	4	29%
Miscellaneous	2	14%

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

In response to a question regarding how OLCF could improve their experience using any of the HPC resources (i.e., Titan, Eos, Rhea, DTNs, HPSS, Lustre/Atlas) and/or tell if any additional resources are needed, the most common response theme was “Satisfied/no suggestions” (41%). This was followed by “Review queue and walltime policies” (11%), and “Re-evaluate purge policy/Provide more advanced notification before purging” (10%). Refer to Table 27 for all themes identified. Select comments include:

Satisfied/no suggestions

“Overall the OLCF resources served the projects on which I collaborate quite well. Of particular note is the collaborative nature of OLCF staff in addressing complex issues that arise when running at scale. Don Maxwell and Jack Wells and many others went above and beyond, upgrading the system software stack to improve support for huge pages, something we identified during our at scale testing of S3D on Legion, part of the ASCR Exact co-design project.”

“My projects have been able to make effective use of titan over the past year, and we consider it our go-to system for both testing and production runs of our climate model and model components”

“We're honored to used OLCF's resources to continue our research and high-impact publications.”

Review queue and walltime policies

“For me the only issue is queue time and run time. If I ask for more cpu usage and wall time, my queue time usually goes up--sometimes by a lot; but if I don't ask for the extra cpu usage, I can hit the wall time. so I have to find the right balance to get the job done most efficiently--but I think this is a general issue with everyone who's jobs are large.”

“The max job run times are problematic for jobs that can't scale to large numbers of processors. This is true in the case of some electronic structure calculations. Either the job doesn't finish on a small number of cores or, many cores are wasted so that the job can have the necessary wall time to finish.”

“It would be helpful if for large jobs, the maximum time limit was 24 hours instead of 12.”

Re-evaluate purge policy/Provide more advanced notification before purging

“The policy of removing all files older than a certain date periodically came as a surprise to me, but I may have simply missed that line in the documentation and/or had files in a different place than I thought I did.”

“Provide continuous access to data as users transition from one project to another. If there are costs to doing this, you should be requesting appropriate funding from DOE. Their new data policies for scientific work require this service be provided.”

“Increase the scratch directory and allow data to store for a longer time, and only delete large files!”

Table 27. *Users’ Descriptions of How the OLCF can Improve Their Experience Using Any of the HPC Resources*

	<i>N</i> = 294	%
Satisfied/no suggestions/NA	121	41%
Review queue and walltime policies	31	11%
Re-evaluate purge policy/Provide more advanced notification before purging	29	10%
Get more software/libraries/update software/libraries	23	8%
Improve data transfer/Offer other methods	19	6%
Improve system performance	14	5%
Improve accessibility	14	5%
Improve documentation and resources/website	13	4%
Improve storage/memory	12	4%
Increase stability/reliability	9	3%
More training	7	2%
Improve workflow tools/compatibility	7	2%
Assistance with GPU usage/more GPU resources	5	2%
More/quicker user assistance/support	3	1%
Improve HSI/HPSS system interface	3	1%
Advance notices of outages	2	1%
Leaving OLCF for resources with a better fit	2	1%
Miscellaneous	13	4%

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

User Opinions of OLCF Services

Users were asked to approximate the number of queries they submitted during 2015. Roughly one-quarter ($n = 75$, 24%) responded that they did not submit a query and 21 (7%) did not provide a response. The largest number of users indicated submitting between 1 and 5 queries ($n = 164$, 53%), while 27 (9%) submitted between 6 and 10 queries, 7 (2%) submitted between 11 and 20 queries, and 14 (5%) submitted more than 20 queries.

User Assistance

Users who indicated they used the OLCF User Assistance were generally satisfied. The percentages of satisfaction across assistance aspects were all 90% or above with means ranging from 4.50 to 4.58 (Table 28 and Table 29). If a user rated any of the aspects of User Assistance with *very dissatisfied* or *dissatisfied* they were asked to explain their rating. Of three users who provided a dissatisfaction rating, two provided explanations for their dissatisfaction. Comments included: “*Poor response to my request for extended data access*” and “*Lack of prompt response.*”

Account Management

Users were asked to rate their satisfaction with two aspects of accounts and allocations as well as provide an overall account services rating. Descriptive statistics for ratings PI status and overall and by project allocation are shown in Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 30 and Table 31. Overall mean ratings for speed of responses to account management queries, effectiveness of response to account management queries, and overall account services were similar (means of 4.53 and 4.63 with satisfaction percentages of 92% to 94%).

If a user rated any of the aspects of Accounts and Allocations with *very dissatisfied* or *dissatisfied* they were asked to explain their rating. Only one user was *dissatisfied* with one aspect of Accounts and Allocations. The *dissatisfied* user provided the following explanation for their dissatisfaction:

“People on my project team needed accounts for certain systems. In most cases, the accounts took a long time to create. In some cases, we were never informed that the accounts were created after waiting for at least a month. I just believe communication with the users about anticipated time to service and service completion would be helpful.”

INCITE Scientific Computing Liaisons

25% ($n = 73$) of users indicated that their project has an assigned INCITE scientific computing liaison. INCITE scientific computing liaison satisfaction ratings are favorable, with an overall mean of 4.64 and a satisfaction percentage of 93% (Table 32). See Table 33 for satisfaction ratings by project allocation. For these aspects, only one user provided a *dissatisfied* rating to aspects of their INCITE scientific computing liaison. The explanation they provided was “*My liaison was too busy to help.*”

Table 28. Descriptive Statistics for Satisfaction Ratings of the User Assistance by PI Status and Overall Totals

	PI-Status				Non-PI Status				Total			
	N	M	SD	%Sat	N	M	SD	%Sat	N	M	SD	%Sat
User assistance - Speed of initial response to queries	41	4.66	0.48	100%	131	4.55	0.71	92%	172	4.58	0.67	94%
User assistance - Speed of final resolution to queries	41	4.63	0.54	98%	131	4.46	0.76	88%	172	4.50	0.71	90%
User assistance - Quality of technical information	41	4.59	0.55	98%	128	4.52	0.72	91%	169	4.53	0.68	92%
User assistance - Response to special requests (i.e., scheduling exceptions, quota increases, software installations, etc.)	34	4.68	0.47	100%	109	4.53	0.79	88%	143	4.57	0.73	91%
User assistance - Overall consulting services	39	4.67	0.48	100%	126	4.50	0.72	90%	165	4.54	0.68	93%

Note. *%Sat is the percentage of respondents indicating 4 (Satisfied) or 5 (Very satisfied).

Table 29. Descriptive Statistics for Satisfaction Ratings of the User Assistance by Project Allocation

	INCITE				DD				ALCC				Other			
	N	M	SD	%Sat	N	M	SD	%Sat	N	M	SD	%Sat	N	M	SD	%Sat
User assistance - Speed of initial response to queries	67	4.58	0.70	93%	108	4.56	0.65	94%	61	4.62	0.58	95%	7	4.43	0.53	100%
User assistance - Speed of final resolution to queries	67	4.51	0.75	93%	108	4.51	0.66	91%	61	4.52	0.67	90%	7	4.29	0.76	86%
User assistance - Quality of technical information	65	4.57	0.73	94%	107	4.51	0.63	93%	61	4.57	0.62	93%	7	4.29	0.76	86%
User assistance - Response to special requests (i.e., scheduling exceptions, quota increases, software installations, etc.)	58	4.62	0.77	91%	94	4.55	0.67	90%	50	4.60	0.73	90%	5	4.80	0.45	100%
User assistance - Overall consulting services	64	4.55	0.75	92%	103	4.53	0.64	92%	58	4.60	0.56	97%	7	4.29	0.76	86%

Note. *%Sat is the percentage of respondents indicating 4 (Satisfied) or 5 (Very satisfied).

Table 30. Descriptive Statistics for Satisfaction Ratings of Account Management 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 responses to account management queries	25	4.56	0.65	92%	47	4.51	0.72	91%	72	4.53	0.69	92%
Effectiveness of response to account management queries	25	4.64	0.64	92%	46	4.63	0.57	96%	71	4.63	0.59	94%
Overall account services	25	4.60	0.65	92%	47	4.60	0.58	96%	72	4.60	0.60	94%

Note. *%Sat is the percentage of respondents indicating 4 (Satisfied) or 5 (Very satisfied).

Table 31. Descriptive Statistics for Satisfaction Ratings of Account Management by Project Allocation

	INCITE				DD				ALCC				Other			
	N	M	SD	%Sat	N	M	SD	%Sat	N	M	SD	%Sat	N	M	SD	%Sat
Speed of responses to account management queries	32	4.47	0.67	91%	45	4.58	0.72	91%	22	4.77	0.43	100%	4	4.25	0.50	100%
Effectiveness of response to account management queries	32	4.53	0.62	94%	44	4.68	0.60	93%	22	4.86	0.35	100%	4	4.25	0.50	100%
Overall account services	32	4.56	0.56	97%	45	4.64	0.61	93%	22	4.77	0.43	100%	4	4.25	0.50	100%

Note. *%Sat is the percentage of respondents indicating 4 (Satisfied) or 5 (Very satisfied).

Table 32. Descriptive Statistics for Satisfaction Ratings of INCITE Scientific Computing Liaison by INCITE 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	9	4.89	0.33	100%	49	4.63	0.70	92%	58	4.67	0.66	93%
Speed of final resolution to queries	9	4.89	0.33	100%	49	4.55	0.71	92%	58	4.60	0.67	93%
Quality of technical support	9	4.78	0.44	100%	48	4.58	0.71	92%	57	4.61	0.67	93%
Response to special requests	8	4.88	0.35	100%	46	4.54	0.75	89%	54	4.59	0.71	91%
Overall support from your INCITE Scientific Computing Liaison	10	4.80	0.42	100%	48	4.60	0.71	92%	58	4.64	0.67	93%

Note. *%Sat is the percentage of respondents indicating 4 (Satisfied) or 5 (Very satisfied).

Table 33. *Descriptive Statistics for Satisfaction Ratings of INCITE Scientific Computing Liaison by INCITE Project Allocation*

	INCITE				No INCITE Project			
	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat
Speed of initial response to queries	51	4.75	0.52	96%	7	4.14	1.21	71%
Speed of final resolution to queries	51	4.67	0.55	96%	7	4.14	1.21	71%
Quality of technical support	50	4.68	0.55	96%	7	4.14	1.21	71%
Response to special requests	47	4.66	0.60	94%	7	4.14	1.21	71%
Overall support from your INCITE Scientific Computing Liaison	51	4.71	0.54	96%	7	4.14	1.21	71%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Communications

Users were mostly satisfied (94% satisfied overall) with how the OLCF keeps them informed of changes, events, and current issues. Table 34 and Table 35 include the descriptive statistics for satisfaction ratings of communications by PI and overall, and by project allocation. Three users provided a *very dissatisfied* or *dissatisfied* rating of OLCF communications. Of the three users, two users provided explanations:

“Sometimes downtimes are not announced there. It would be really cool to have some sort of visualization that showed how loaded the machine was currently and how many requests are pending.”

“Need for quicker response to query submitted in user assistance.”

When asked “Do you feel adequately informed about OLCF changes? Please explain.”, 96% (276 out of 287) of users responded that they feel adequately informed of changes while 4% (or 11 users) indicated no, they did not feel adequately informed of changes. Users were asked to explain their responses and of 14 users who provided explanations, 9 had indicated “yes” and 5 had indicated “no”, see Table 36 for qualitative themes. Among those who said “yes”, the most common explanation was “Emails are adequate” (67%). Among those who said “no”, the most common explanations were “Not enough information provided” (40%) and “Issues related to upgrades” (40%).

When asked “Do you feel adequately informed about OLCF events? Please explain”, 97% (278 out of 287) of users responded that they feel adequately informed of events while 3% (or 9 users) indicated no. Table 37 includes the breakdown summary for 10 users who explained their response (9 said “yes” and 1 said “no”). Among those who said “yes”, the most common explanation was “Emails contain needed information” (60%). The one who said “no”, said “*I can only find information by visiting the OLCF website.*”

When asked “Do you feel adequately informed about current issues? Please explain”, roughly 95% (273 out of 287) of users responded that they feel adequately informed of current issues while 5% (14 users) indicated no. Of the 10 users providing explanations for their ratings, 7 users had responded with yes and 3 users had responded with no, see Table 38. Among those who said “yes”, the most common explanation was “Emails contain needed information” (57%). Among those who said “no”, the explanation provided by all was “Information about issues is not all communicated to users” (100%).

Table 34. *Descriptive Statistics for Satisfaction Ratings of Communications by PI Status and Overall Totals*

	PI-Status				Non-PI Status				Total			
	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat
E-mail announcements	57	4.54	0.60	95%	226	4.47	0.63	95%	283	4.48	0.62	95%
Announcements on the OLCF website	53	4.45	0.61	94%	201	4.28	0.74	86%	254	4.32	0.72	88%
Overall communications	57	4.53	0.60	95%	225	4.42	0.66	93%	282	4.44	0.65	94%

Note. %Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 35. *Descriptive Statistics for Satisfaction Ratings of Communications by Project Allocation*

	INCITE				DD				ALCC				Other			
	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat												
E-mail announcements	122	4.56	0.62	98%	150	4.45	0.61	94%	92	4.57	0.56	97%	7	4.29	0.49	100%
Announcements on the OLCF website	108	4.35	0.75	90%	131	4.30	0.69	87%	83	4.39	0.66	90%	5	4.20	0.45	100%
Overall communications	122	4.49	0.65	96%	149	4.41	0.66	92%	92	4.49	0.60	95%	7	4.29	0.49	100%

Note. %Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 36. *Users' Explanations for Why They Do or Don't Feel Adequately Informed about OLCF Changes*

	<i>N</i> = 14	%
Explanations for those who said "yes" (N = 9)		
Emails are adequate	6	67%
Miscellaneous	3	33%
Explanations for those who said "no" (N = 5)		
Not enough information provided	2	40%
Issues related to upgrades	2	40%
Miscellaneous	1	20%

Table 37. *Users' Explanations for Why They Do or Don't Feel Adequately Informed about OLCF Events*

	N = 10	%
Explanations for those who said "yes" (= 9)		
Emails contain needed information	3	60%
Miscellaneous	2	40%
Explanations for those who said "no" (n = 1)		
I can only find information by visiting the OLCF website.	1	100%

Table 38. *Users' Explanations for Why They Do or Don't Feel Adequately Informed about Current Issues*

	N = 10	%
Explanations for those who said "yes" (n = 7)		
Emails contain needed information	4	57%
Miscellaneous	4	57%
Explanations for those who said "no" (n = 3)		
Information about issues is not all communicated to users	3	100%

Training

A total of 286 (93%) of survey respondents provided a response to the question, "How do you prefer to receive training?" (see Figure 3). The majority of users who responded prefer online documentation ($n = 228$, 80%) or online training ($n = 179$, 63%). Two users indicated other with one user providing the following comment: "Please add native Linux support to web sessions/calls!"

Out of the 286 who indicated the most convenient time to attend a training event, the majority of respondents indicated no preference ($n = 154$, 54%), and roughly one-third ($n = 84$, 29%) indicated the summer was the most convenient time. Small numbers of users indicated the spring, fall, or winter, see Figure 4 for the distribution of users' responses.

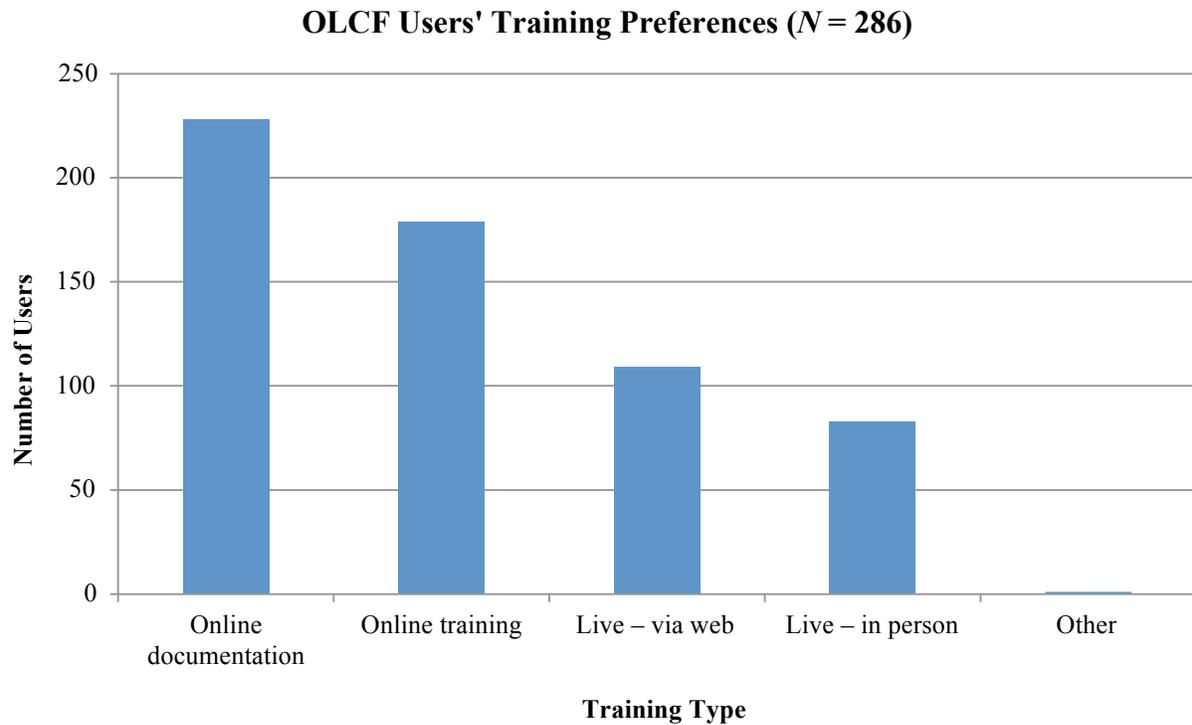


Figure 3. Training preferences of OLCF users.

Preference in Time of Year for Training (N = 286)

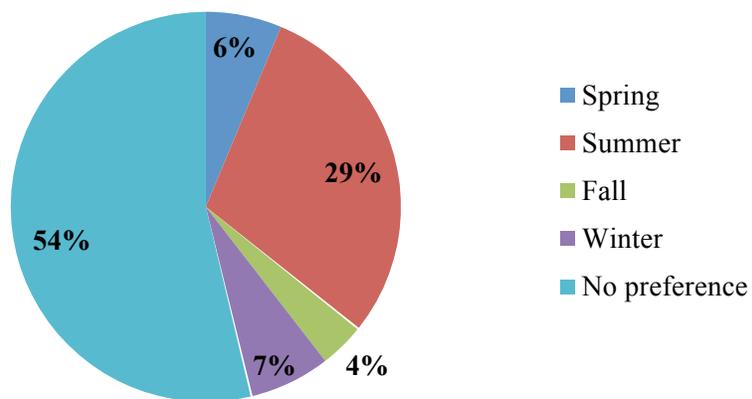


Figure 4. Most convenient time to attend a training event.

A total of 41 users indicated that they had participated in OLCF training during the 2015 calendar year. Of these users, 31 (76%) said they would recommend attending a future OLCF training event in person, 9 (22%) said maybe, and 1 person (2%) said no. When asked “What training topic(s) would you like to see offered in the future?” 16 topics were suggested by more than one user (Table 39). The top three suggestions provided include: “GPUs” (28%), “Parallelization/parallel profiling” (10%), and “Use of performance tools/performance monitoring” (10%). Sample responses for these suggestions include:

GPUs

- “GPU programming - currently being offered but haven't taken advantage”
- “Programming on graphical processing units for non-computer-science majors.”
- “More GPU events”

Parallelization/parallel profiling

- “Parallel programming with OpenACC and optimizing codes for Summit”
- “Code writing, Parallel computing and parallelization tips”

Use of performance tools/performance monitoring

- “Performance tools.”
- “How to properly use the latest performance and debugging tools on Titan”

Table 39. *Users’ Suggestions for Training Topics They Would Like to See Offered by the OLCF in the Future*

	<i>N = 72</i>	<i>%</i>
GPUs	20	28%
Parallelization/parallel profiling	7	10%
Use of performance tools/performance monitoring	7	10%
Debugging	6	8%
Data visualization	5	7%
Info and updates on OLCF systems, libraries, and tools	5	7%
OpenACC	5	7%
CUDA	5	7%
Coding/Code optimization	5	7%
Compiling	4	6%
Data analysis and management	4	6%
HPC Basics	4	6%
Don't know	3	4%
Advanced topics	3	4%
No preference	2	3%
I/O	2	3%
Data transfer	2	3%
Satisfied with current offerings	1	1%
Miscellaneous	6	8%

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

Users were asked to provide satisfaction ratings for their overall satisfaction with OLCF and five specific training aspects (Table 40 and Table 41). Satisfaction ratings for overall satisfaction were positive ($M = 4.31$, $SD = 0.69$, 89% satisfaction). The lowest satisfaction rating for training aspects was for monthly user conference calls ($M = 4.14$, $SD = 0.75$, 78% satisfaction).

OLCF Website

Users provided satisfaction ratings for User Support aspects of the OLCF website. Overall, users were generally satisfied with the user support information provided on the OLCF website.

If a user rated any of the aspects of User Assistance (help@olcf.ornl.gov or (865) 241-6536) with *very dissatisfied* or *dissatisfied* they were asked to explain their rating.

Of 11 users who indicated *very dissatisfied* or *dissatisfied* with one or more aspects of User Assistance (help@olcf.ornl.gov or (865)241-6536), 9 provided explanations for their dissatisfaction. The two most common responses themes included “Poor system status information” (44%) and “Poor organization” (33%). Sample comments for these themes include:

System status information

“Had trouble getting reliable up-to-date account status at My OLCF at times in the past”

“The system status color coding is poor, since the red and green are indistinguishable for some color-blind people.”

“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.”

Organization

“MyOLCF is looking much better. The website has improved greatly in the last year. The user guides continue to be a resource I regularly consult. I don't know that the searchable knowledge base is bad, but I rated it dissatisfied just because I don't find myself consulting it or regularly finding it helpful.”

“System status page is hidden under support. There should be a direct link on the main page.”

Table 40. Descriptive Statistics for Satisfaction Ratings of Training Aspects by PI Status and Overall Totals

	PI-Status				Non-PI Status				Total			
	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat
Getting Started Guide	50	4.46	0.68	90%	195	4.35	0.68	90%	245	4.38	0.68	90%
Web Tutorials	47	4.23	0.73	87%	161	4.23	0.76	84%	208	4.23	0.75	85%
Training Events	33	4.30	0.68	88%	127	4.15	0.79	81%	160	4.18	0.77	83%
Archived Training Event Slides	37	4.24	0.72	89%	136	4.20	0.78	83%	173	4.21	0.76	84%
Monthly User Conference Calls	32	4.25	0.72	84%	106	4.11	0.76	76%	138	4.14	0.75	78%
Overall satisfaction with OLCF training	50	4.44	0.64	92%	197	4.27	0.70	89%	247	4.31	0.69	89%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 41. Descriptive Statistics for Satisfaction Ratings of Training Aspects by Project Allocation

	INCITE				DD				ALCC				Other			
	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat												
Getting Started Guide	99	4.34	0.72	90%	131	4.37	0.73	87%	82	4.45	0.63	95%	6	4.50	0.55	100%
Web Tutorials	87	4.23	0.71	86%	112	4.21	0.75	82%	65	4.26	0.76	89%	5	4.60	0.55	100%
Training Events	71	4.21	0.67	89%	86	4.17	0.74	80%	49	4.24	0.83	86%	6	4.17	0.75	83%
Archived Training Event Slides	75	4.20	0.68	88%	101	4.21	0.74	83%	54	4.24	0.82	85%	6	4.17	0.75	83%
Monthly User Conference Calls	59	4.07	0.72	78%	75	4.13	0.74	79%	37	4.24	0.80	78%	4	4.25	0.50	100%
Overall satisfaction with OLCF training	103	4.33	0.65	92%	134	4.33	0.68	88%	82	4.33	0.74	90%	7	4.43	0.53	100%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 42. *Descriptive Statistics for Satisfaction Ratings of Aspects of User Support on 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
System user guides	55	4.45	0.63	93%	215	4.41	0.67	95%	270	4.42	0.66	94%
Searchable knowledge base:	48	4.25	0.67	88%	195	4.20	0.78	85%	243	4.21	0.76	85%
Software pages	51	4.18	0.71	82%	198	4.24	0.75	87%	249	4.22	0.74	86%
OLCF system status	55	4.45	0.66	95%	209	4.33	0.71	91%	264	4.36	0.70	92%
My OLCF	48	4.33	0.78	85%	168	4.32	0.68	90%	216	4.32	0.71	89%
Overall rating of User Support info on the OLCF website	55	4.44	0.57	96%	217	4.34	0.65	94%	272	4.36	0.63	94%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 43. *Descriptive Statistics for Satisfaction Ratings of Aspects of User Support on the OLCF Website by Project Allocation*

	INCITE				DD				ALCC				Other			
	N	M	SD	%Sat	N	M	SD	%Sat	N	M	SD	%Sat	N	M	SD	%Sat
System user guides	119	4.40	0.64	95%	140	4.40	0.68	92%	87	4.47	0.68	98%	7	4.71	0.49	100%
Searchable knowledge base	107	4.17	0.73	86%	126	4.22	0.75	84%	79	4.28	0.80	87%	7	4.14	0.90	71%
Software pages	111	4.27	0.66	90%	127	4.18	0.76	82%	81	4.28	0.78	86%	7	4.14	0.69	86%
OLCF system status	119	4.34	0.68	93%	138	4.42	0.64	92%	82	4.39	0.70	96%	7	4.43	0.79	86%
My OLCF	95	4.35	0.60	94%	118	4.26	0.73	85%	66	4.41	0.72	94%	6	4.17	0.75	83%
Overall rating of User Support info on the OLCF website	121	4.36	0.59	96%	139	4.32	0.62	92%	89	4.44	0.67	96%	7	4.29	0.49	100%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 44. *Users' Explanations for Dissatisfaction with User Assistance Website*

	N = 9	%
Poor system status information	4	44%
Poor organization	3	33%
Outdated information	2	22%

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

Additionally, users were asked to provide satisfaction ratings for the OLCF Website overall as well as aspects of a) ease of navigation, b) accuracy of information, and c) timeliness of information. Descriptive statistics for ratings are provided in Table 45 for all users as well as in Table 46 by project allocation. Overall users were generally satisfied with the website, with less satisfied ratings provided for ease of navigation ($M = 4.19$, $SD = 0.67$, 88% satisfied).

If a user rated any of the aspects of the OLCF Website, <http://olcf.ornl.gov>, with *very dissatisfied* or *dissatisfied* they were asked to explain their rating (Table 47). Of 8 users who were *very dissatisfied* or *dissatisfied* with one or more aspects of the OLCF Website, <http://olcf.ornl.gov>, 6 provided explanations for their dissatisfaction. The two most common responses themes included “Outdated information” (50%) and “Poor organization” (33%). Sample comments for these themes include:

Outdated information

“Some software pages are outdated and inaccurate with respect to the current versions available on HPC resources.”

Poor organization

“Ideally navigation gives comprehensive overview of all topics available and "where you are" in the content tree. On the OLCF website, the sidebar shows the top level divisions, but there is no persistent display, otherwise, of where you are browsing in the documentation.”

Table 45. *Descriptive Statistics for Satisfaction Ratings of 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	58	4.21	0.64	91%	218	4.18	0.67	88%	276	4.19	0.67	88%
Accuracy of information	58	4.40	0.59	95%	219	4.26	0.70	89%	277	4.29	0.68	90%
Timeliness of information	58	4.31	0.60	93%	215	4.23	0.69	88%	273	4.25	0.67	89%
Overall satisfaction with the OLCF website	58	4.33	0.60	93%	218	4.26	0.65	90%	276	4.27	0.64	91%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 46. *Descriptive Statistics for Satisfaction Ratings of OLCF Website by Project Allocation*

	INCITE				DD				ALCC				Other			
	N	M	SD	%Sat	N	M	SD	%Sat	N	M	SD	%Sat	N	M	SD	%Sat
Ease of navigation	119	4.16	0.69	88%	144	4.17	0.66	87%	91	4.21	0.71	88%	6	4.00	0.63	83%
Accuracy of information	119	4.31	0.63	94%	144	4.24	0.71	88%	92	4.35	0.69	90%	6	4.17	0.75	83%
Timeliness of information	119	4.24	0.68	90%	143	4.20	0.69	87%	88	4.32	0.69	90%	6	4.00	0.63	83%
Overall satisfaction with the OLCF website	119	4.27	0.65	92%	144	4.24	0.65	90%	91	4.31	0.68	90%	6	4.17	0.41	100%

Note. *%Sat is the percentage of respondents indicating 4 (*Satisfied*) or 5 (*Very satisfied*).

Table 47. *Users' Explanations for Dissatisfaction with the OLCF Website, <http://olcf.ornl.gov>*

	N = 6	%
Outdated information	3	50%
Poor organization	2	33%
Poor system status information	1	17%

When asked “What additional services or information would you like to have available on the OLCF website?” the majority of respondents indicated that nothing additional was needed (60%); however, five topics were suggested by more than one user (Table 48). The top two suggestions provided were “Better website organization” (11%) and “More information about allocations and tools” (9%) Sample comments for these suggestions include:

Better website organization

“Higher information density on the home page would be nice. This probably goes against best practices for webpages, but if system status and big announcements could be listed on the home page that would be awesome.”

“It is surprisingly difficult to find information about resetting account passwords, separately from contacting the OLCF User Assistance Center directly.”

“fewer clicks to "system status" info”

More information about allocations and tools

“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.”

“GPU utilization by our project through MyOLCF”

“Make sure that disk storage and file system information is up-to-date at all times.”

Table 48. *Users’ Suggestions for Additional Services or Information They Would Like to Have Available on the OLCF Website*

	<i>N = 55</i>	<i>%</i>
Nothing additional needed	33	60%
Better website organization	6	11%
More information about allocations and tools	5	9%
Better communication about issues	3	5%
Tutorials	2	4%
Machine downtime status indicators needed	2	4%
Miscellaneous	4	7%

Data Analysis, Visualization, and Workflow

Users were asked to indicate where they analyze data produced by OLCF jobs. Of 287 users providing a response to the question, 45 (16%) indicated that they do not need data analysis. Of the users indicating need (*n* = 242), 22 (9%) indicated producing all analyses at OLCF, 30 (12%) indicated producing most at OLCF, and 42 (17%) indicating producing half at OLCF. The majority of users indicated that they produce all analysis elsewhere (*n* = 83, 34%) or most elsewhere (*n* = 65, 27%).

Eighty-three percent of users (*n* = 238 out of 287) indicated that they had not taken advantage of the OLCF cross-platform submission capabilities in their workflow with 49 (17%) indicating that

they had. Those who responded “yes” were then asked, “Do you have any suggestions for improving the OLCF cross-platform submission?” Eleven users responded, but only one provided a suggestion: “*Allow submission of jobs from outside OLCF enclave, or allow SSH connections to be allowed open - in a similar way to NERSC does.*”

When asked “Why haven't you taken advantage of the OLCF cross-platform submission capabilities in your workflow?” 84% (200 of the 238) of users who indicated not having taken advantage of the OLCF cross-platform submission capabilities in their workflow chose to provide an explanation (Table 52). The most two common responses included “Not aware of them, need more information” (45%) and “No need” (37%). Sample responses for each of these themes include:

Not aware of them, need more information

“Need to learn more about this”

“I don't know what it is.”

“Need more information and assistance with this”

No need

“Need has not arisen. Logging into Titan / Eos is easy.”

“Mostly focused on projects where it is not relevant”

“Work hasn't required it.”

Table 51. *Descriptive Statistics for Satisfaction Ratings for Aspects of Data Analysis and Visualization Support Services by PI Status and Overall*

	PI-Status				Non-PI Status				Total			
	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat	<i>N</i>	<i>M</i>	<i>SD</i>	%Sat
Ability to perform data analysis	4	4.25	0.50	100%	15	4.27	0.88	87%	19	4.26	0.81	89%
Ability to perform project workflows	4	4.25	0.50	100%	14	4.29	0.73	86%	18	4.28	0.67	89%
Sufficiency of the OLCF hardware for your data analysis, visualization, and workflow needs	5	4.20	0.45	100%	16	4.31	0.87	88%	21	4.29	0.78	90%
Sufficiency of tools for your data analysis, visualization, and workflow needs	5	4.00	0.71	80%	16	4.13	0.96	75%	21	4.10	0.89	76%

Table 52. *Users' Explanations for Why They Haven't Taken Advantage of the OLCF Cross-Platform Submission Capabilities in Their Workflow*

	<i>N</i> = 200	%
Not aware of them, need more information	89	45%
No need/not applicable	73	37%
Have not had a chance/have not tried yet	15	8%
Only use/need to use Titan	7	4%
Not sure of the benefit	6	3%
Not compatible with current workflow/working on integrating	5	3%
Plan to use/will consider using in the future	5	3%
No specific reason	2	1%
Miscellaneous	8	4%

Users were asked to rate the importance of 15 different data aspects. Descriptive statistics of ratings by PI status and overall totals are included in Table 53 and by project allocation in Table 54. Figure 5 demonstrates data aspects by percentage of users indicating slightly important or more important. Long-term data retention, I/O bandwidth to local disk, long-term data curation, and archival storage space received the highest percentage of users indicating the aspect of being at least slightly important or more. The least important data aspect was for general public access to your data over the web.

Table 53. Importance Ratings by PI-Status and Total.

	PI-Status				Non-PI Status				Total			
	<i>N</i>	<i>M</i>	<i>SD</i>	%Imp	<i>N</i>	<i>M</i>	<i>SD</i>	%Imp	<i>N</i>	<i>M</i>	<i>SD</i>	%Imp
General public access to your data over the web	59	2.42	1.29	64%	228	2.31	1.19	64%	287	2.33	1.21	64%
Access for collaborators to your data over the web	59	2.83	1.30	75%	228	2.80	1.21	80%	287	2.81	1.22	79%
Access for your specific OLCF project members to your data over the web	59	2.85	1.30	76%	228	3.03	1.24	82%	287	2.99	1.26	80%
Long-term data retention	59	3.63	1.20	93%	228	3.62	1.05	94%	287	3.62	1.08	94%
Long-term data curation	59	3.25	1.25	86%	228	3.32	1.15	90%	287	3.30	1.17	89%
Access to databases at the OLCF	59	2.66	1.38	68%	228	2.66	1.22	75%	287	2.66	1.26	74%
Workflow tools/libraries	59	3.05	1.31	81%	228	3.07	1.16	87%	287	3.07	1.19	86%
Dedicated workflow machines	59	2.78	1.38	71%	228	2.66	1.14	79%	287	2.69	1.19	78%
Data management tools	59	2.95	1.24	80%	228	2.87	1.18	83%	287	2.89	1.19	83%
Analysis and visualization assistance from the OLCF	59	2.73	1.24	75%	228	2.68	1.15	81%	287	2.69	1.16	79%
Access to a large shared-memory system for data analysis & visualization	59	3.00	1.33	81%	228	2.80	1.22	78%	287	2.84	1.24	79%
Access to a system with GPUs specifically for data analysis & visualization	59	2.63	1.35	69%	228	2.68	1.29	74%	287	2.67	1.30	73%
Filesystem metadata performance	59	2.86	1.31	78%	228	2.92	1.11	86%	287	2.91	1.15	85%
Archival storage space	59	3.51	1.32	88%	228	3.23	1.22	88%	287	3.29	1.24	88%
I/O bandwidth to local disk	59	3.73	1.08	93%	228	3.44	1.09	93%	287	3.50	1.09	93%

Note. %Imp represents the percentage of respondents indicating 2 (*Slightly important*) to 5 (*Extremely important*).

Table 54. *Importance Ratings by Project Allocation.*

	INCITE				DD				ALCC				Other			
	N	M	SD	%Imp	N	M	SD	%Imp	N	M	SD	%Imp	N	M	SD	%Imp
General public access to your data over the web	123	2.49	1.23	71%	151	2.30	1.19	63%	94	2.26	1.24	59%	7	3.00	1.41	86%
Access for collaborators to your data over the web	123	2.93	1.25	80%	151	2.83	1.22	81%	94	2.68	1.25	78%	7	2.71	1.60	71%
Access for your specific OLCF project members to your data over the web	123	3.04	1.30	80%	151	3.05	1.22	83%	94	2.84	1.30	78%	7	3.00	1.73	71%
Long-term data retention	123	3.76	1.06	95%	151	3.52	1.09	93%	94	3.69	1.02	98%	7	2.86	1.57	71%
Long-term data curation	123	3.43	1.17	92%	151	3.26	1.14	89%	94	3.24	1.18	88%	7	2.71	1.60	71%
Access to databases at the OLCF	123	2.69	1.29	72%	151	2.70	1.24	76%	94	2.55	1.26	71%	7	2.43	1.13	71%
Workflow tools/libraries	123	3.04	1.24	85%	151	3.23	1.18	89%	94	3.01	1.13	88%	7	3.14	1.35	86%
Dedicated workflow machines	123	2.72	1.24	76%	151	2.85	1.18	81%	94	2.60	1.21	77%	7	2.86	1.57	71%
Data management tools	123	2.84	1.22	80%	151	3.08	1.16	87%	94	2.66	1.22	77%	7	3.14	1.57	86%
Analysis and visualization assistance from the OLCF	123	2.66	1.21	78%	151	2.82	1.18	81%	94	2.59	1.16	78%	7	2.71	1.50	71%
Access to a large shared-memory system for data analysis & visualization	123	2.72	1.31	76%	151	3.00	1.24	82%	94	2.82	1.28	78%	7	3.43	1.81	71%
Access to a system with GPUs specifically for data analysis & visualization	123	2.57	1.36	69%	151	2.80	1.28	76%	94	2.54	1.27	71%	7	3.29	1.70	71%
Filesystem metadata performance	123	2.93	1.17	85%	151	2.95	1.17	83%	94	2.88	1.20	84%	7	3.14	1.35	86%
Archival storage space	123	3.47	1.25	90%	151	3.30	1.20	91%	94	3.16	1.27	86%	7	2.86	1.57	71%
I/O bandwidth to local disk	123	3.47	1.16	90%	151	3.56	1.06	94%	94	3.43	1.13	91%	7	3.71	1.38	86%

Note. %Imp represents the percentage of respondents indicating 2 (*Slightly important*) to 5 (*Extremely important*).

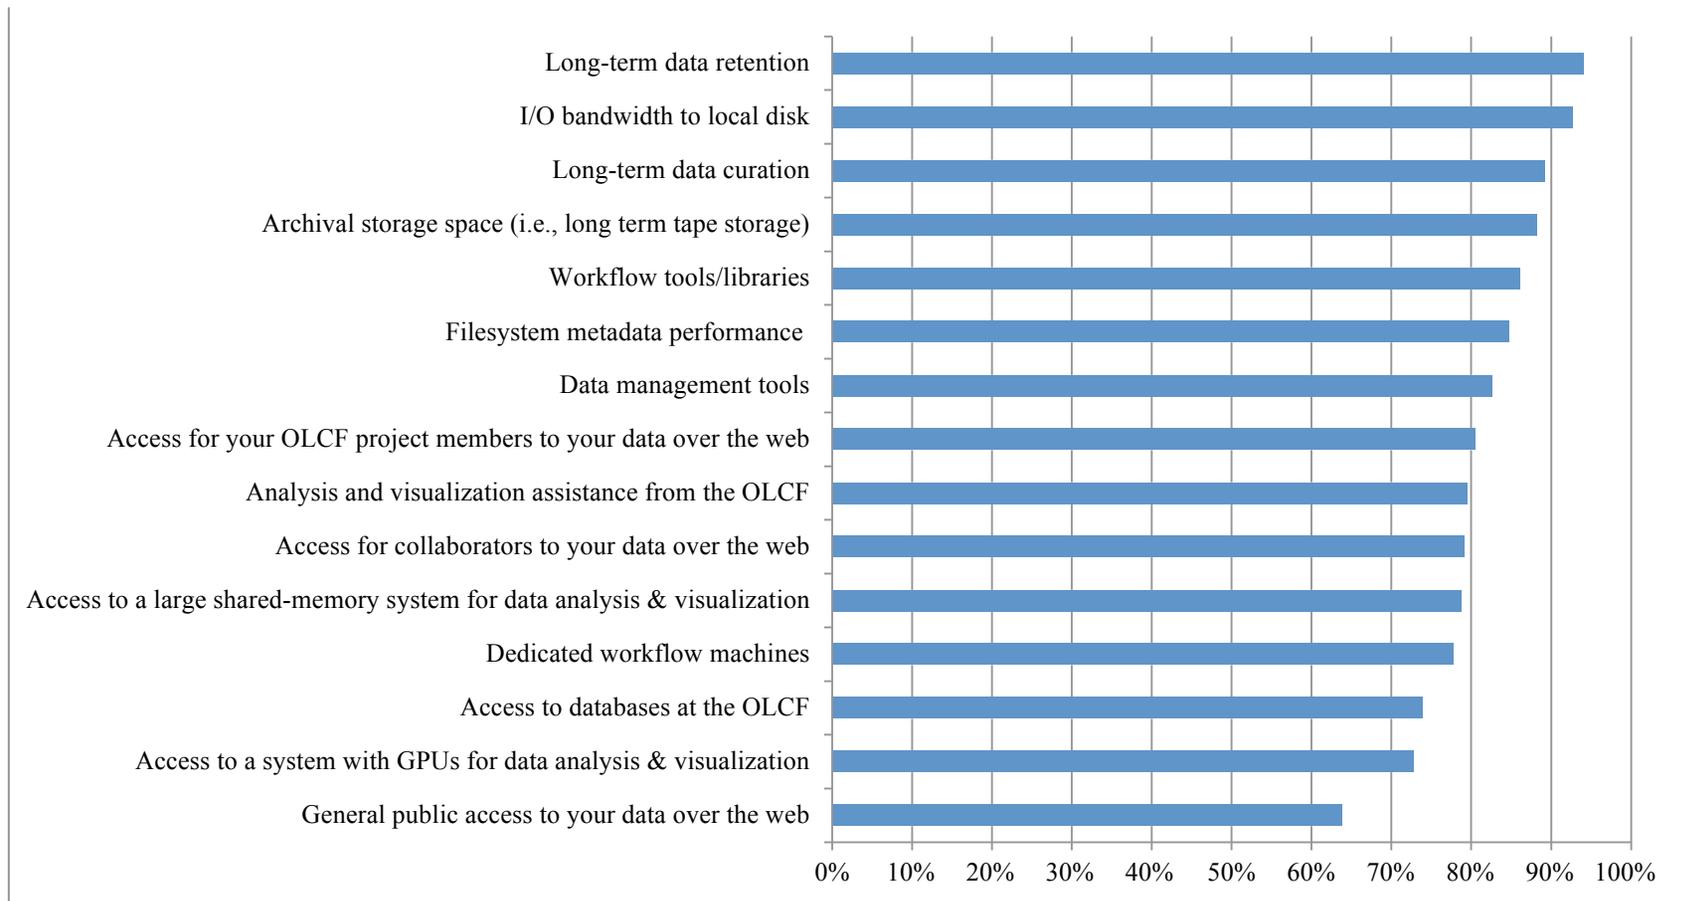


Figure 5. Percentage of total respondents indicating importance.

When asked “What additional data analysis, visualization, and/or workflow services would you like the OLCF to provide?”, 63 users responded. While the largest proportion of users (68%) indicated that no additional data analysis, visualization, and/or workflow services are needed, 8% of respondents to this question ($n = 5$) responded with “miscellaneous visualization services” and 8% responded with “miscellaneous workflow services” (See Table 55). Example responses for each of these categories of responses include:

Miscellaneous visualization services

“It would be helpful for the visualization support scientists to proactively contact users for potential collaboration.”

“more convenient visualization tools on EVEREST”

“more insight licenses”

Miscellaneous workflow services

“Support for Pegasus and other similar workflow systems.”

“Some pilot demonstration projects to illustrate the end-to-end workflow customized for domain scientists, and how to customize the generic end-to-end workflow based on user requirements.”

Table 55. *Users’ Suggestions for Additional Data Analysis, Visualization, and/or Workflow Services they Would Like the OLCF to Provide*

	<i>N = 63</i>	<i>%</i>
No additional services needed/NA/not sure	43	68%
Miscellaneous visualization services/tools	5	8%
Miscellaneous workflow services/tools	5	8%
Miscellaneous data analysis services/tools	3	5%
Satisfied	3	5%
Miscellaneous other	6	10%

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

Additional Services, Resources, and/or Other Improvements Needed

Users were asked “What additional services, resources, and/or other improvements are needed to enhance your experience at the OLCF?” While the largest proportion of users (38%) indicated that no additional services and/or resources are needed to enhance their experience at the OLCF, 7% ($n = 19$) of respondents to this question responded with “Improved accessibility” and 7% ($n = 18$) responded with “Better performance/capacity.” Response themes are provided in

Table 56. Example responses for each of these three categories of miscellaneous responses include:

No additional needs/don't know

“I have none right now. I do not know about the future though.”

“I can't think of anything. I've had a great experience and questions answered quickly throughout my time.”

Improved accessibility

“I'd like to get rid of these darn physical OTP tokens and start using soft tokens. Even both. I grow tired of having to keep track of 5 different tokens.”

“user access -- We have several team members from politically "sensitive" countries, including a recent Gordon Bell prize winner. There is currently no way for these team members to access OLCF systems. I would like to see some avenue for them to get user accounts.”

Better performance/capacity

“I could always use more machines with faster cores and larger memory, but that is the eternal struggle. Otherwise my experiences have been highly positive and I believe for the given resources that my needs have been thoroughly met.”

“More computing power”

Table 56. *Users' Suggestions for Additional Services and/or Resources Needed to Enhance Their Experience at the OLCF*

	<i>N</i> = 273	%
No additional needs/don't know	103	38%
Improved accessibility	19	7%
Better performance/capacity	18	7%
More memory/storage	17	6%
Review queue policy/queue time	13	5%
Longer data retention	11	4%
More tools/improved tools	10	4%
Increase walltime/core hours	9	3%
Workflow improvements	9	3%
More documentation/improve website	8	3%
Improve data transfer	8	3%
Training requests/suggestions	7	3%
Testing improvements/test clusters	7	3%
Enhanced user support/list of contacts	6	2%
Updates to software and libraries	6	2%
What to know more about visualization/improve visualization	6	2%
Stability/reliability of systems	5	2%
More development of GPU resources	4	1%
Increased bandwidth	3	1%
More advanced notifications	2	1%
Miscellaneous	19	

