2020 Oak Ridge Leadership Computing Facility User Survey

Findings and Recommendations

PREPARED BY:

Ann M. Martin, Ph.D.
Senior Evaluator
ORISE STEM Workforce Development

Kelly P. Townsend, MSBA
Evaluation Specialist
ORISE STEM Workforce Development

Kate Miller-Bains, Ph.D.
Evaluation Specialist
ORISE STEM Workforce Development

Laura A. Davenport, Ph.D.
Evaluation Specialist
ORISE STEM Workforce Development

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

PREPARED FOR:

Oak Ridge Leadership Computing Facility
Oak Ridge National Laboratory
February 2021

This document describes activities performed under contract number DE-SC0014664 between the U.S. Department of Energy and Oak Ridge Associated Universities. 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 (DOE) asset that is dedicated to enabling critical scientific, research, and health initiatives of the department and its laboratory system by providing world class expertise in STEM workforce development, scientific and technical reviews, and the evaluation of radiation exposure and environmental contamination.

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

ORAU provides innovative scientific and technical solutions to advance national priorities in science, education, security and health. Through specialized teams of experts, unique laboratory capabilities and access to a consortium of more than 120 major Ph.D.-granting institutions, ORAU works with federal, state, local and commercial customers to advance national priorities and serve the public interest. A 501(c)(3) nonprofit corporation and federal contractor, ORAU manages the Oak Ridge Institute for Science and Education (ORISE) for the U.S. Department of Energy (DOE). Learn more about ORAU at www.orau.org.

Table of Contents

List of Tables

List of Figures

Executive Summary

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

Respondents

At the end of the five-week survey period, 688 users completed the survey out of 1260 possible respondents, giving an overall response rate of 54.6%. Respondents' projects were supported ECP (41%), INCITE (39%), Director's Discretion (29%) and ALCC (18%) allocations.

Findings Highlights

Overall Evaluation

The proportions of all respondents *satisfied* or *very satisfied* with OLCF resources/services ranged from 92% to 97% for the five "overall" evaluation items. Specifically, ratings for these five major categories of resources/services were a) OLCF (97%), b) Compute Resources (96%), c) Data Resources (93%), d) Support Staff (93%), and e) Support Services (92%). Thematic analysis of open-ended comments identified *computing power/hardware/HPC resources/performance* (39% of respondents) and *user support/staff* (31% of respondents) as the most valued OLCF qualities.

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

	High rating			Medium rating		Low rating	
	All	**PI**	**Non-PI**	**INCITE**	**DD**	**ALCC**	**ECP**
Max N responding:	688	93	595	267	200	126	284
OLCF	**97%**	**98%**	**96%**	**97%**	**98%**	**97%**	**96%**
Compute Resources	**96%**	**97%**	**96%**	**96%**	**96%**	**97%**	**95%**
Rhea	94%	89%	96%	94%	95%	92%	100%
Summit	96%	98%	96%	97%	97%	97%	96%
Data Resources	**93%**	**98%**	**93%**	**94%**	**95%**	**95%**	**91%**
Data Transfer Nodes	91%	91%	91%	91%	87%	92%	85%
HPSS	98%	100%	98%	99%	97%	97%	96%
Alpine GPFS Scratch Filesystem	95%	100%	94%	95%	95%	96%	93%
Support Staff	**93%**	**92%**	**93%**	**94%**	**95%**	**93%**	**92%**
Support Services	**92%**	**90%**	**92%**	**92%**	**93%**	**92%**	**90%**
Support received*	94%	94%	95%	95%	95%	91%	93%
OLCF website	94%	94%	94%	92%	92%	93%	95%
Communications	96%	98%	95%	97%	97%	93%	96%
Support/training documentation	93%	94%	93%	94%	93%	91%	93%
Problem resolution	94%	94%	94%	96%	91%	91%	95%

*Support received lists out: User Assistance, Accounts, INCITE Scientific Computing Liaisons, and Advanced Data/Workflow Liaisons.

OLCF Systems, Data Resources, and Compute Resources

Summit, which became available to users in 2019, was utilized by 93% of users, and Rhea by 21% of users. Titan and Eos were eliminated from service and therefore from this year's survey. The majority of the 480 users (76.7%) noted no changes in overall OLCF computing performance over the last year, while 22.9% cited improved performance; a small number of users (0.4%) noted a decrease in performance compared to FY 2019. Overall satisfaction across the compute resources and data resources ranged from 91% (Data Transfer Nodes) to 98% (HPSS) of users either *satisfied* or *very satisfied*. 93% of users were *satisfied* or *very satisfied* with notice for scheduled downtimes, and 92% of users were that satisfied with both project disk space and bandwidth offered by the OLCF.

Support Services

Users were asked to provide ratings of their overall satisfaction with support received from the wide variety of OLCF services available. Most respondents (94%) were either *satisfied* or *very satisfied* with support received from user assistance, accounts, INCITE Scientific Computing Liaisons, and Advanced Data/Workflow Liaisons.

Communication with Users

96% of respondents were overall *satisfied* or *very satisfied* with how OLCF keeps them informed of changes, events, downtimes, and current issues.

Problem Resolution

More than half (57%) of respondents submitted between one and five queries to OLCF (via phone or email) in 2020. 94% of respondents were *satisfied* or *very satisfied* with OLCF's problem resolution overall. The highest rated aspect of OLCF's problem resolution was the *quality of OLCF response to reported issues* (94% satisfaction) followed closely by *timeliness of OLCF responses to reported issues* (93% satisfaction) and *usefulness of support and training documentation* (93% satisfaction).

Website

Half (49%) of respondents indicated that they had visited the OLCF website during 2020. 39% of respondents indicated they visit the OLCF website (http://olcf.ornl.gov) once a week or more frequently. 94% of respondents indicated they were *satisfied* or *very satisfied* with the OLCF Website. The highest rated aspect of the OLCF website was the *usefulness of content* (95% satisfaction). *Search capabilities* were the lowest rated aspect of website usability (84% satisfaction).

Data Analysis and Visualization

22% of respondents indicated they analyze *most* or *all* of their data at OLCF while 53% analyze most or all of their data elsewhere. 25% of respondents analyze about half of their data at OLCF and the other half elsewhere. When asked about the source of users' data, the largest proportion of users is working with data that is primarily (*most* or *all*) sourced from OLCF jobs (66%). Fifty-two (8%, *N* = 681) respondents indicated they were interested in scheduling one-on-one conversations with OLCF analysis and visualization specialists in order to consult on needs and approaches, and their contact information was referred to the OLCF to arrange consultations.

Introduction

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

This report first briefly describes the data collection and analysis procedures. It then presents findings with respect to user characteristics, patterns of OLCF resource use, and satisfaction ratings of OLCF resources/services. The report also provides some basic longitudinal comparisons of user responses from 2006 through 2020. Finally, recommendations for possible improvements are offered.

Data Collection and Analysis

Survey Revisions

In collaboration with OLCF, the Assessment and Evaluation team at ORAU made substantial revisions to the user survey in FY 2019. The revisions in 2020 reduced the total number of possible items in the survey by approximately 56%, from 181 to 80 items. These revisions were largely carried over to the FY2020 survey. Additional revisions were made in FY 2020 to account for systems that had been sunset and were no longer applicable to the user experience, and to add an overall satisfaction item to measure satisfaction with OLCF support staff.

Prior to the FY 2019 survey revisions, the average survey completion time as measured and reported by the survey software interface was 18.5 minutes and the average user responded to 90 items. In FY 2019, the average response time was 9.8 minutes and the typical respondent answered 38 items. For the current FY 2020 survey, the average response time was 6.74 minutes and the typical respondent once again answered 38 items.

Data Collection

The survey sampling frame was constituted by first collecting the names of individuals who had logged into an OLCF system between January 1, 2020, through September 30, 2020. OLCF staff and vendors, as well as individuals with invalid email addresses, were then removed from the list. Any users who did not have at least one project allocation categorized as INCITE, DD, ALCC, or ECP were also removed from the list, per guidance from OLCF indicating that additional project allocations were not intended for the annual user survey.

OLCF staff invited all OLCF users from this list to participate in the survey, which was hosted online beginning on October 12, 2020, and remained open for completion through November 16, 2020 (see Appendix B: Survey Administration Timeline and Appendix F: Survey). Since visitors to the OLCF website and others on OLCF distribution lists could access the survey, an additional four users were identified and added to the user group after they had responded. Seventeen users were removed from the user group because their email addresses were unreachable at the time the survey was administered and

one additional user was removed from the user group after discovering their user data (including their email) was incorrect and therefore they never received an invitation to complete the user survey.

Overall, this process resulted in a sampling frame with 1260 OLCF users. A total of 688 users completed or partially completed the survey, resulting in a response rate of 54.6%. Figure 15, within Appendix B: Survey Administration Timeline, highlights the value of each reminder email in increasing the response rate. Appendix A: Survey Invitations and Reminders provides the text of each reminder email. The reminders sent by Sheila Moore on November 3, 2020, were particularly effective, resulting in a two-day burst in responses comprising almost 40% of the total responses received. These reminders were specific to each project allocation and included user lists, so that OLCF was able to leverage the influence of PIs in encouraging their colleagues to respond. Based on this successful implementation, this reminder approach is recommended for future iterations of the survey.

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

OLCF (Overall)

OLCF Computing Resources

- Summit
- Rhea

OLCF Data Resources

- Data Transfer Nodes (DTNs)
- HPSS
- Alpine GPFS Scratch Filesystem

OLCF Support Services

- Overall: Support staff
- Overall: Support services
- Support received (User Assistance, Accounts, INCITE Scientific Computing Liaisons, Advanced Data/Workflow Liaisons)
- OLCF website
- Communication with users
- Support and training documentation
- Problem resolution
- Data analysis and visualization

Data Analysis

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

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

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

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

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

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

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

Note that, in recent years, an "Other" category has been reported which combined General, Vendor, and other smaller programs, and the Summit Early Science ("ES") program was reported as a separate allocation in FY 2019. Only the four allocations listed above are reported in FY 2020.

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

- *N* = Total number of respondents who answered the question
- *n* = Total number of respondents who answered the specific item in the question or who provided a specific response
- *M* = the arithmetic average of respondents' scores from 1 (*very dissatisfied)* to 5 (*very satisfied)*
- *SD* = Standard deviation (indicating average deviation from the mean)
- *Var* = Variance, the square of the standard deviation, or the deviation from the mean in squared units; this statistic is included only in the overall summary tables, because it is reported by OLCF to the Department of Energy.
- *%Sat* = percentage of respondents indicating 4 (*satisfied*) or 5 (*very satisfied*) on satisfaction scales

Color coding has been used in the report tables as below, unless otherwise noted:

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

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

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

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

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

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

Access
Accessibility and authentication
SSH
Accounts and credentials
Hardware Computing Resources
Performance
Systems (Summit, Andes, etc.)
Management, infrastructure, and maintenance
GPU resources
Stability/reliability
Running Jobs
Containers
Workflow
Scheduling policy
Queues
Wall/run time
Data Management
Data retention/purge policy and procedures
Data storage
Data transfer and I/O
File systems
Software
Software/tools/modules
Libraries and updates
Compilers
Visualization
Development, test, and debugging tools
Monitoring/profiling
User Support
Documentation
Tutorials
Training
Support and ticketing
Communication
Example Additional Categories
Satisfaction
Miscellaneous
Survey suggestions
Support for scientific research

Findings

Respondents

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

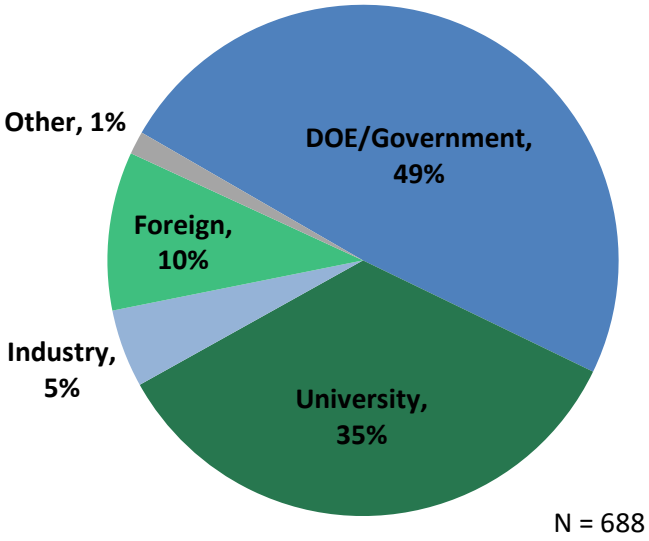

Figure 1. Respondent occupational affiliation

Note: Percentages may not add up precisely to 100% due to rounding in each category.

The distribution of OLCF users across **project allocations** is shown in Figure 2 and in detail in Table 2. There are statistically significant differences between the respondent pool and the user pool for each allocation; results of chi-square testing were statistically significant at the $p < 0.001$ level for ECP, INCITE, and DD, and at the $p < 0.05$ level for INCITE. The pool of survey respondents over-represents ECP, INCITE, and ALCC users, and under-represents DD users. It is therefore possible that bias exists in the findings due to the self-selection of survey respondents, and throughout this report tables separately report the findings from each allocation.

Note that the table categories are not exclusive (e.g., the INCITE category includes individuals assigned to INCITE, but who may also have been assigned to other projects). Note that 76% of respondents were assigned by OLCF to a single project allocation (i.e., assignment to only INCITE, only DD, only ALCC, or only ECP).

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

	OLCF Users (*N* = 1260)		Survey Respondents (*N* = 688)	
	Percentage	***n***	**Percentage**	***n***
ECP	33%	420	41%	284
INCITE	32%	399	39%	267
DD	44%	549	29%	200
ALCC	16%	200	18%	126

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

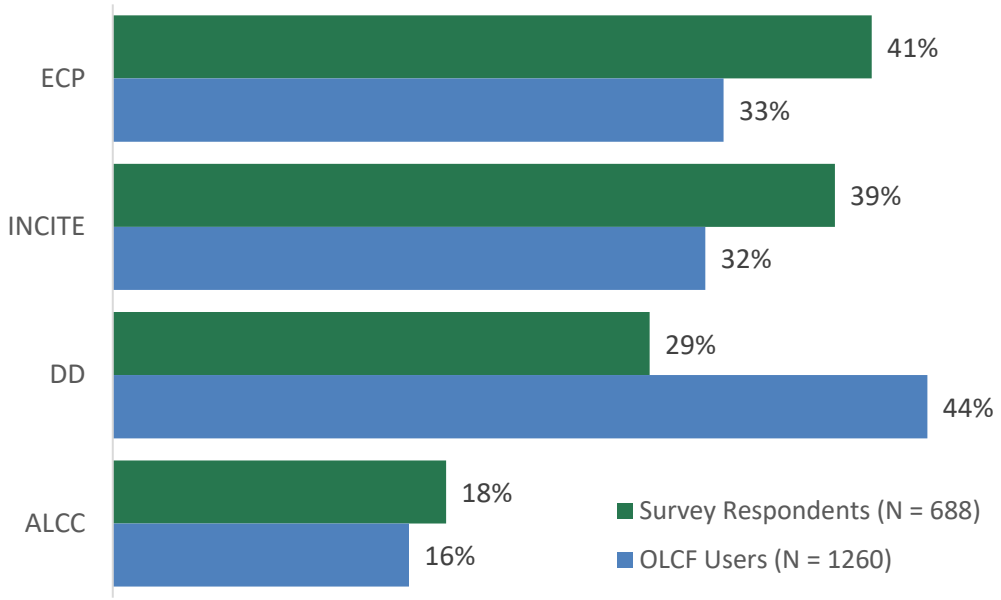

Figure 2. Project allocations for OLCF users and survey respondents

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

The proportions of OLCF users and of 2020 survey respondents with PI status on at least one project are displayed in Figure 3. The survey respondent pool slightly over-represents PIs. Throughout this report, tables separately report findings from respondents with PI status from those without PI status.

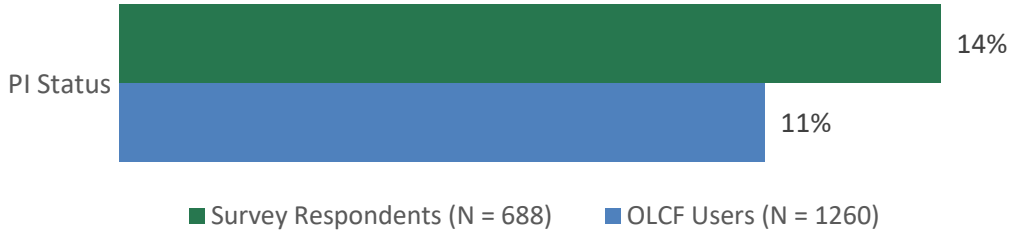

Figure 3. PI Status for OLCF users and survey respondents

Resource Utilization

Overall experience using the OLCF was approximately evenly split between three categories, with the largest proportion (39%) having more than 2 years of experience using the facility (Figure 4).

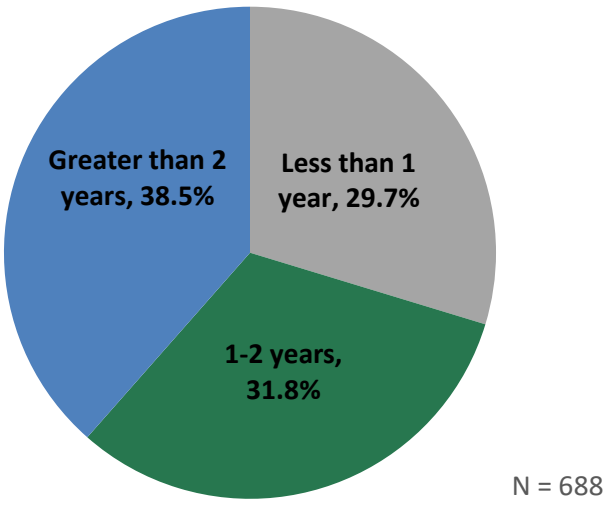

Figure 4. Experience using the OLCF

Note: Percentages may not add up precisely to 100% due to rounding in each category.

Respondents were asked to indicate **which OLCF HPC resources they utilized** during the 2020 calendar year; utilization of the OLCF website was the only specific **OLCF support service** about which users were asked. For all categories, the largest proportion of respondents indicated using Summit; HPSS was utilized by the smallest proportion (Table 3).

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

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

	PI Status		INCITE		DD		ALCC		ECP		Total	
	n	**% Users**	***n***	**% Users**	***n***	**% Users**	***n***	**% Users**	***n***	**% Users**	***n***	**% Users**
Summit	88	95%	248	93%	183	92%	121	96%	274	96%	641	93%
Rhea	29	31%	85	32%	62	31%	40	32%	27	10%	146	21%
Data Transfer Nodes	33	35%	90	34%	78	39%	50	40%	46	16%	193	28%
HPSS	26	28%	73	27%	41	21%	34	27%	28	10%	123	18%
Alpine GPFS	39	42%	106	40%	88	44%	48	38%	85	30%	231	34%
OLCF Website	52	56%	136	51%	108	54%	73	58%	127	45%	334	49%
I have not used any of the listed resources	2	2%	4	1%	2	1%	1	1%	5	2%	13	2%

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

Overall Satisfaction

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

Table 4 summarizes descriptive statistics for these overall satisfaction ratings for all respondents and broken down by PI status, while Table 5 reports satisfaction statistics across project allocations. The tables also include ratings of **specific compute resources** (i.e., Rhea and Summit), data resources (i.e., Data Transfer Nodes, HPSS, and Alpine GPFS), and both **support staff** and **support services** (i.e., support received via user assistance, accounts, INCITE Liaisons, Advanced Data/Workflow Liaisons, as well as the OLCF website, communications, support and training documentation, and problem resolution). Across these 15 key items, which include the five "overall" ratings, and considering the entire group of "All" respondents, the tables show that:

- *%Sat* ranged from 91% to 98%,
- Mean satisfaction ratings ranged from 4.4 to 4.6, and
- SDs ranged from 0.52 to 0.82.

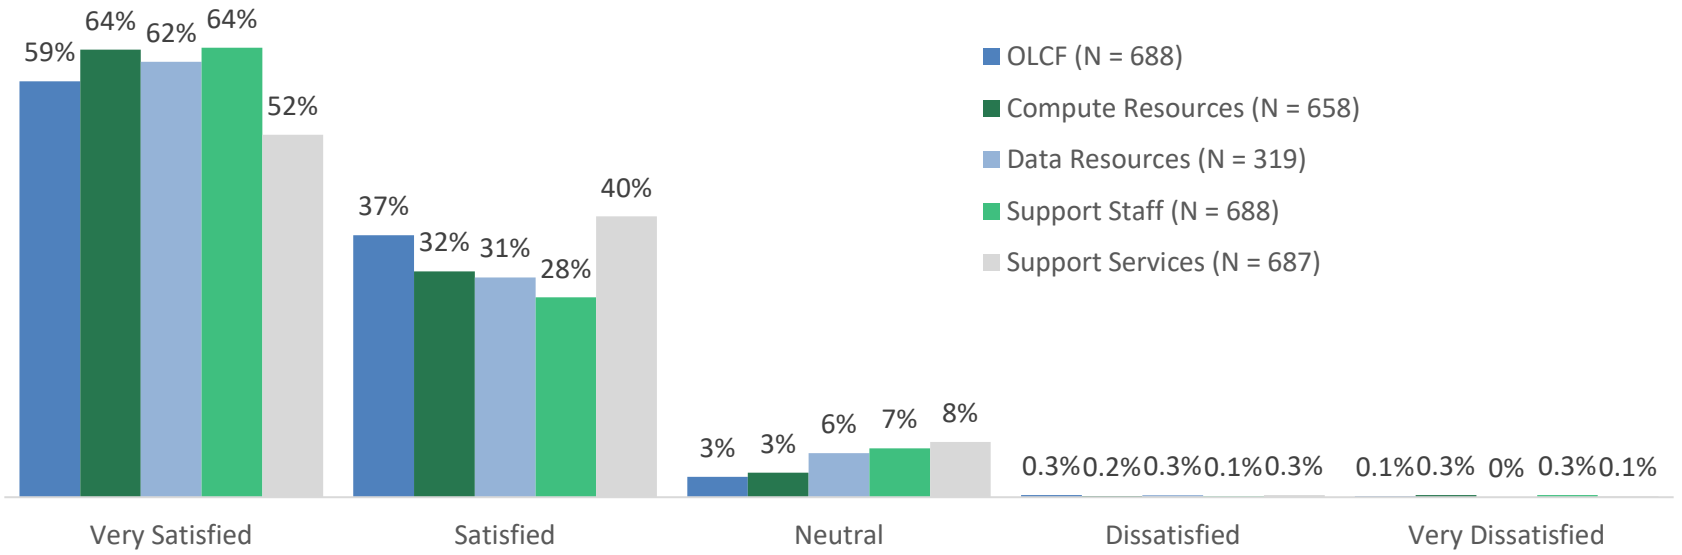

Figure 5. "Overall" satisfaction with OLCF and its major resources/services

Note: Percentages may not add up precisely to 100% due to rounding in each category.

Table 4. *Overall Satisfaction with OLCF and Its Major Resources/Services by PI Status and Totals*

	PI Status					Non-PI Status					Total				
	N	***M***	***Var.***	***SD***	***%Sat***	***N***	***M***	***Var.***	***SD***	***%Sat***	***N***	***M***	***Var.***	***SD***	***%Sat***
OLCF	**93**	**4.7**	**0.38**	**0.62**	**98%**	**595**	**4.5**	**0.34**	**0.58**	**96%**	**688**	**4.6**	**0.34**	**0.58**	**97%**
Compute Resources	**88**	**4.6**	**0.42**	**0.65**	**97%**	**570**	**4.6**	**0.35**	**0.59**	**96%**	**658**	**4.6**	**0.36**	**0.60**	**96%**
Rhea	27	4.6	0.49	0.70	89%	116	4.5	0.34	0.58	96%	143	4.5	0.36	0.60	94%
Summit	86	4.6	0.28	0.53	98%	542	4.5	0.34	0.58	96%	628	4.5	0.32	0.57	96%
Data Resources	**52**	**4.7**	**0.26**	**0.51**	**98%**	**267**	**4.5**	**0.41**	**0.64**	**93%**	**319**	**4.6**	**0.40**	**0.63**	**93%**
Data Transfer Nodes	32	4.4	0.90	0.95	91%	159	4.4	0.64	0.80	91%	191	4.4	0.67	0.82	91%
HPSS	25	4.7	0.21	0.46	100%	94	4.6	0.29	0.54	98%	119	4.6	0.27	0.52	98%
Alpine GPFS Scratch Filesystem	38	4.7	0.21	0.46	100%	189	4.5	0.41	0.64	94%	227	4.5	0.38	0.62	95%
Support Staff	**93**	**4.6**	**0.50**	**0.71**	**92%**	**595**	**4.6**	**0.42**	**0.65**	**93%**	**688**	**4.6**	**0.44**	**0.66**	**93%**
Support Services	**92**	**4.5**	**0.56**	**0.75**	**90%**	**595**	**4.4**	**0.42**	**0.65**	**92%**	**687**	**4.4**	**0.44**	**0.66**	**92%**
Support received*	86	4.7	0.35	0.59	94%	510	4.6	0.37	0.61	95%	596	4.6	0.37	0.61	94%
OLCF website	51	4.4	0.36	0.60	94%	277	4.4	0.37	0.61	94%	328	4.4	0.36	0.60	94%
Communications	89	4.7	0.27	0.52	98%	570	4.5	0.35	0.59	95%	659	4.6	0.34	0.58	96%
Support and training documentation	79	4.5	0.38	0.62	94%	460	4.5	0.42	0.65	93%	539	4.5	0.41	0.64	93%
Problem resolution	79	4.5	0.38	0.62	94%	450	4.5	0.41	0.64	94%	529	4.5	0.41	0.64	94%
Min	**25**	**4.4**	**0.21**	**0.46**	**89%**	**94**	**4.4**	**0.29**	**0.54**	**91%**	**119**	**4.4**	**0.27**	**0.52**	**91%**
Max	**93**	**4.7**	**0.90**	**0.95**	**100%**	**595**	**4.6**	**0.64**	**0.80**	**98%**	**688**	**4.6**	**0.67**	**0.82**	**98%**

*"Support received" lists out and includes the following support areas: User Assistance, Accounts, INCITE Scientific Computing Liaisons, and Advanced Data/Workflow Liaisons.

Table 5. *Overall Satisfaction with OLCF and Its Major Resources/Services by Project Allocation*

	INCITE					DD					ALCC					ECP				
	N	***M***	***Var.***	***SD***	***%Sat***	***N***	***M***	***Var.***	***SD***	***%Sat***	***N***	***M***	***Var.***	***SD***	***%Sat***	***N***	***M***	***Var.***	***SD***	***%Sat***
OLCF	**267**	**4.6**	**0.31**	**0.56**	**97%**	**200**	**4.6**	**0.28**	**0.53**	**98%**	**126**	**4.6**	**0.34**	**0.58**	**97%**	**284**	**4.5**	**0.40**	**0.63**	**96%**
Compute Resources	**258**	**4.6**	**0.35**	**0.59**	**96%**	**190**	**4.6**	**0.31**	**0.56**	**96%**	**124**	**4.6**	**0.38**	**0.62**	**97%**	**274**	**4.5**	**0.44**	**0.66**	**95%**
Rhea	83	4.6	0.37	0.61	94%	61	4.5	0.35	0.59	95%	38	4.4	0.41	0.64	92%	26	4.4	0.25	0.50	100%
Summit	243	4.6	0.30	0.55	97%	182	4.5	0.31	0.56	97%	119	4.5	0.30	0.55	97%	268	4.4	0.35	0.59	96%
Data Resources	**145**	**4.6**	**0.36**	**0.60**	**94%**	**118**	**4.6**	**0.35**	**0.59**	**95%**	**74**	**4.5**	**0.42**	**0.65**	**95%**	**105**	**4.5**	**0.42**	**0.65**	**91%**
Data Transfer Nodes	88	4.4	0.59	0.77	91%	77	4.3	0.85	0.92	87%	50	4.4	0.69	0.83	92%	46	4.2	0.90	0.95	85%
HPSS	69	4.6	0.27	0.52	99%	39	4.7	0.28	0.53	97%	34	4.6	0.31	0.56	97%	28	4.5	0.34	0.58	96%
Alpine GPFS Scratch Filesystem	103	4.5	0.38	0.62	95%	87	4.5	0.40	0.63	95%	48	4.5	0.34	0.58	96%	84	4.5	0.49	0.70	93%
Support Staff	**267**	**4.6**	**0.40**	**0.63**	**94%**	**200**	**4.6**	**0.34**	**0.58**	**95%**	**126**	**4.6**	**0.38**	**0.62**	**93%**	**284**	**4.5**	**0.46**	**0.68**	**92%**
Support Services	**266**	**4.4**	**0.42**	**0.65**	**92%**	**200**	**4.5**	**0.40**	**0.63**	**93%**	**126**	**4.5**	**0.41**	**0.64**	**92%**	**284**	**4.3**	**0.46**	**0.68**	**90%**
Support received*	229	4.7	0.36	0.60	95%	184	4.6	0.35	0.59	95%	119	4.6	0.44	0.66	91%	240	4.5	0.40	0.63	93%
OLCF website	135	4.4	0.44	0.66	92%	106	4.3	0.38	0.62	92%	71	4.4	0.37	0.61	93%	124	4.4	0.34	0.58	95%
Communications	255	4.6	0.32	0.57	97%	196	4.6	0.29	0.54	97%	124	4.5	0.40	0.63	93%	267	4.5	0.34	0.58	96%
Support and training documentation	211	4.5	0.40	0.63	94%	164	4.5	0.38	0.62	93%	103	4.5	0.42	0.65	91%	221	4.5	0.44	0.66	93%
Problem resolution	206	4.6	0.35	0.59	96%	164	4.5	0.42	0.65	91%	105	4.5	0.46	0.68	91%	217	4.5	0.41	0.64	95%
Min	**69**	**4.4**	**0.27**	**0.52**	**91%**	**39**	**4.3**	**0.28**	**0.53**	**87%**	**34**	**4.4**	**0.30**	**0.55**	**91%**	**26**	**4.2**	**0.25**	**0.5**	**85%**
Max	**267**	**4.7**	**0.59**	**0.77**	**99%**	**200**	**4.7**	**0.85**	**0.92**	**98%**	**126**	**4.6**	**0.69**	**0.83**	**97%**	**284**	**4.5**	**0.90**	**0.95**	**100%**

*"Support received" lists out and includes the following support areas: User Assistance, Accounts, INCITE Scientific Computing Liaisons, and Advanced Data/Workflow Liaisons.

Only 8 respondents reported **dissatisfaction with the OLCF overall or with its major resources/services, and** all provided explanations. Individuals cited problems with *support and documentation* (*n* = 5), *Summit* (*n* = 4), *purge policy* (*n* = 1), *security and access* (*n* = 1), and *user survey* (*n* = 1). Illustrative examples include:

Support and documentation	*"A lot of documentation is out of date or missing information about specific tools. It would also be nice if you documented which modules are supported, provide documentation links, and keep it up to date."*
Summit	*"The original users guide for Summit was much better. This one is confusing and feels like information is missing."* *"Over the past year, I have been working on development of production code on Summit. I have run into numerous issues which appear unique to Summit. First in my debugging code, the cuda-gdb is completely unusable. It will crash or idle at various locations within the run which are unpredictable, unreproducible, and unrelated to the errors trying to be debugged . . . A more startling issue is that we discovered instances of where MPI incorrectly transmits data between ranks when using the UVM on multiple nodes. It is particularly worrisome, because MPI and our code does not really have access to the knowledge of whether it is being run on one node or several, but these produce different results. It's possible that all of the calculations using MPI and UVM in this way on Summit are incorrect due to these improper data transfers. Currently my coworker who is more familiar with the finer details of this code has a ticket open trying to understand this issue."*
Purge policy	*"The 90 day purge policy is very inconvenient."*
Security and access	*"The MFA requirement coupled with draconian and unnecessarily intrusive SSH security policies make Summit cumbersome to use for development. Learn from other centers like NERSC who provide the necessary security while prioritizing the user experience. Give me a way to enter my MFA token at most ONCE per work day and use ssh connection sharing or otherwise share that authentication across opening multiple ssh connections . . . The process for yearly renewal of existing user accounts is unnecessarily tedious and a waste of time for everyone involved. You should instead give us a form with all our existing information populated and ask us to update anything that has changed."*

Survey	*"OLCF staff is very pushy about this survey. The survey should be optional, not mandatory, and spamming people by calling it an "obligation" and cc'ing others."*

Finally, respondents described what they perceived to be **"the best qualities of OLCF."** Many illustrative examples praised multiple elements of OLCF:

> "*Access to Summit is the best computing resource we have, without question - it has made it possible for us to consider problems that we would have otherwise not even attempted. We have had little interaction with OLCF staff, simply because we encountered no problems, all of our codes are running very smoothly."*
>
> *"In my experience the best qualities of OLCF are:*
> *1) Leadership class computational resources that enable high risk and impactful science.*
> *2) Collaboration and support from our INCITE Liaison has been extremely helpful.*
> *3) Timeliness and quality of problem resolution is excellent."*
>
> *"The stability of the system, relatively easy-to-use job submission system, availability of tools and libraries."*
>
> *"Provision of cutting edge uncompromising compute facilities with excellent availability, latest compilers and stable run environment. A streamlined governance and reporting infrastructure and a very helpful user support documentation with a wealth of information. Friendly, responsive and helpful staff."*
>
> *"OLCF provides access to computing resources of a magnitude that we would otherwise not have access to. Utilizing Summit allows us to run problems that we cannot run on any other system. The OLCF staff make using the systems straightforward and the support we receive has always been excellent."*
>
> *"The physical resources (Summit etc.) at OLCF are of course among the very best in the world, but in my view a great, if not the greatest asset of OLCF, is their high-quality staff-both those involved with operations (running the computers resources and administration) and the scientific staff (and in particular those with whom I interact). Other supercomputing facilities with top-class computers I currently use or have used in the past (in my country of residence and abroad) may have very good operations staff, but generally lack the scientific staff. The second kind of specialist makes all the difference from the point of view of the domain specialist (in my case) when trying to make the most out of the resources."*

Thematic analysis of user responses identified *computing power/hardware/HPC resources/performance* (39%) and *user support/staff* (31%) as the most valued qualities of the OLCF

Best Qualities of the OLCF for all responses by category; *N* = 630). Appreciation for the power and performance of the facilities has been expressed in user surveys across several years, as has the high frequency of positive references to OLCF staff and user support. These responses were re-examined, excluding individual responses that mentioned **only** computing performance as the best quality (removing 114 responses). The relative frequency of comments reported by this group (*N* = 516), excluding references to computing power/performance is shown in the last column of Table 6. When the responses are examined in this way, *OLCF staff and user support* is prominent as the perceived best OLCF quality, with a gap between this response and the next more frequent quality. There is significant spread across other categories and variety in responses.

For example:

OLCF staff and user support	*"I really like the OLCF support team. I have encountered a couple issues with the Summit system. The support team was very helpful and proactive in assisting me."* *"The support and responsiveness of the OLCF staff in ensuring that we have the resources necessary to get our jobs running."* *"Technical rigor from the staff."*
Summit	*"Summit is a great resource."* *"For my purposes, the best quality of the OLCF is the incredibly high performance capabilities of each Summit compute node. This makes the testing and debugging process very efficient."*
System documentation & website information	*"Excellent documentation that is very easy to navigate."* *"We're pleased with the user guides, which are extremely helpful for navigating the complexities of the system (especially the complexity of using jsrun and resource sets)."*
Management/infrastructure/ maintenance of facility	*"The system is well maintained for large number of users."* *"Always trying to provide the cutting-edge computational capabilities."*
Data transfer/filesystem/I/O	*"Transferring data to and from OLCF over the DTNs using Globus are lightning fast."* *"I/O rate."*

Table 6. *Best Qualities of OLCF (ordered by % of all respondents, high to low)*

	All Responses (*N* = 630)	**Responses Excluding Computing Performance (*N* = 516)**
Computing power/hardware/HPC resources	39%	--
OLCF staff and user support	31%	38%
Summit	11%	14%
System documentation & website information	10%	12%
Availability/uptime	9%	11%
Queue time, turnaround, and scheduling policy	9%	11%
Environment and tools/stack (software, libraries, visualization, & analysis)	8%	10%
Stability/reliability	6%	8%
GPU resources	6%	7%
Management/infrastructure/ maintenance of facility	5%	6%
Ease of use	4%	5%
Training/tutorials	4%	5%
Supports scientific research/demanding problems/"bleeding edge" science	3%	4%
Newer user/unsure of ability to answer	3%	3%
Communication	2%	3%
Data storage/disk space	2%	2%
Data transfer/filesystem/I/O	2%	2%
User-/customer-focused	2%	2%
Accessibility	2%	2%
Overall satisfaction	2%	2%
Miscellaneous/Other	4%	5%

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

Compute and Data Resources

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

- Sufficient notice of scheduled downtimes
- Sufficient disk space
- Bandwidth offered by the OLCF
- I/O performance

Table 7 reports satisfaction for these features by PI status and overall, and Table 8 reports ratings by project allocation. The highest satisfaction ratings (all respondents) were for *notice given prior to scheduled maintenance* (93% satisfied). The lowest overall mean rating was for *I/O performance* (91% satisfied).

Fourteen respondents indicated **dissatisfaction with one or more aspects of the OLCF compute and data resources,** and 11 provided explanations for their dissatisfaction. The majority of complaints related to *hanging or sluggishness* (*n* = 4) and *performance/capabilities* (*n* = 3). Illustrative examples include:

Hanging or sluggishness	*"Frequent system unavailability due to the stale file system and sluggish file server."* *"Recently in the home directory of the headnode there are some random significant delays of accessing files such as ls and writing by vim."*
Performance/capabilities	*"We've run into scaling issues with sharing a 12.5GB/s link across 3 GPUs in a socket."* *"GPFS is too slow to use it for Deep Learning training in the (statistically rigorous) way I'd like to--so I have to make 'hack' solutions."*

Table 7. *Satisfaction Ratings for Features of the OLCF HPC Compute and Data Resources by PI Status and Overall Totals*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Sufficient notice of scheduled downtimes	88	4.6	0.53	98%	545	4.5	0.65	93%	633	4.5	0.63	93%
Sufficient project disk space	89	4.6	0.57	96%	542	4.5	0.69	91%	631	4.5	0.68	92%
Bandwidth offered by the OLCF	88	4.6	0.55	97%	528	4.5	0.68	92%	616	4.5	0.67	92%
I/O performance	86	4.6	0.60	94%	528	4.4	0.71	91%	614	4.5	0.70	91%

Table 8. *Satisfaction Ratings for Features of the OLCF HPC Compute and Data Resources by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Sufficient notice of scheduled downtimes	250	4.6	0.61	94%	190	4.6	0.61	94%	119	4.5	0.69	89%	256	4.5	0.66	93%
Sufficient project disk space	251	4.6	0.66	92%	189	4.6	0.61	94%	120	4.5	0.66	93%	256	4.4	0.73	89%
Bandwidth offered by the OLCF	241	4.6	0.62	93%	187	4.5	0.67	91%	122	4.5	0.66	93%	250	4.5	0.70	92%
I/O performance	246	4.5	0.68	92%	188	4.5	0.73	90%	119	4.4	0.67	92%	244	4.4	0.75	90%

In addition, respondents were asked to indicate their opinion regarding the **performance of computing and data resources compared to the previous year.** Overall, only 22.9% reported *improvements*, just 0.4% perceived *decreases in performance*, and 76.7% reported *no change* (Figure 6). Some differences in these perceptions were observed across years of using the OLCF. Less experienced users (those with only 1-2 years' experience) were slightly less likely to report seeing an improvement over the last year compared to those with greater than 2 years of experience.

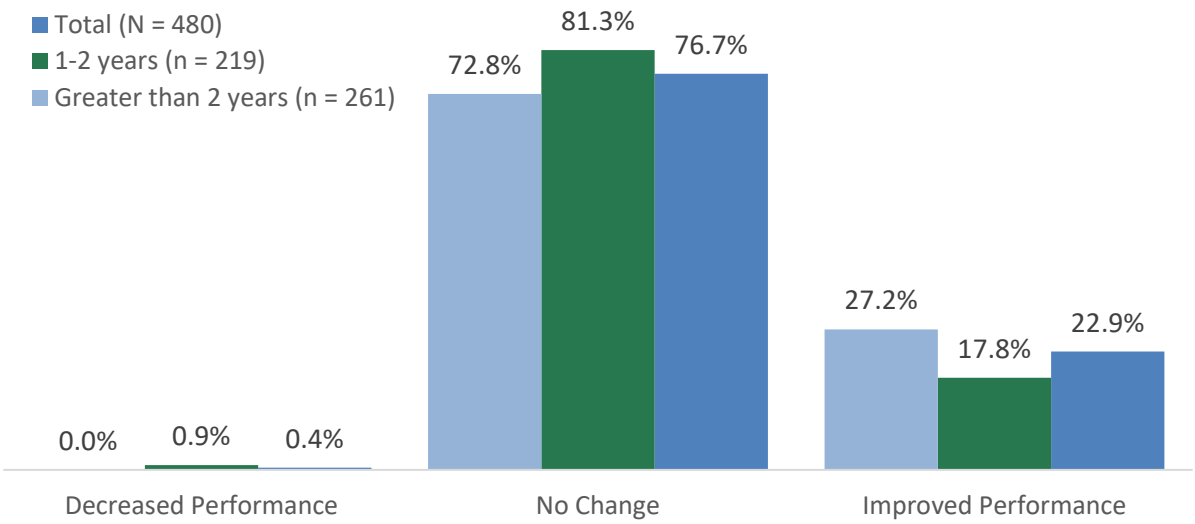

Figure 6. Perceived changes from FY 2019 in performance of computing/data resources, by years of experience using OLCF

Note: Percentages may not add up precisely to 100% due to rounding in each category.

Only two respondents provided comments **describing or explaining decreased performance**:

> *"Job throughput on Summit was low for some period earlier this year. This was likely due to prioritized COVID-19 compute jobs."*

> *"There was an issue with mpi_allreduce when running on a very high node count in Spring of 2020. This issue took a lot of our time to work around. Summit support team [names removed] were very helpful. I believe though initially our complaints were not taken seriously (by vendor?)."*

Compute Resources

Summit

Summit was used by 93% of respondents during the 2020 calendar year (*N* = 641). Summit users were asked to provide satisfaction ratings for multiple aspects of the system, and descriptive statistics of these ratings are shown in Table 9, which also reports satisfaction by PI status. 96% of all respondents were either *satisfied* or *very satisfied* with the system overall. Table 10 summarizes these satisfaction statistics by project allocation. The *system availability* was the highest rated specific aspect of Summit, and the lowest rated specific aspect was *availability of libraries.*

There were 29 Summit users who reported **dissatisfaction with Summit**, and 28 of them provided explanations for their dissatisfaction. Over half of these users (*n* = 16) were unhappy with *tools: software, compilers, libraries, and compatibilities* that had an impact on the work they could conduct on the system. For example:

> *"I had issues with profiling tools such as Tau and ScoreP. There were some instances when latest versions were not installed or had issues. This is a minor issue, but it would be good to have updated, working profiling tools available on Summit."*

> *"No post-processing software. I had to get access to Rhea to be able to post-process using ParaView."*

The next largest group of dissatisfied users (*n* = 10) were unhappy with *scheduling, wait time, and job time.* For example:

> *"Compared to last year, the job waiting time is significantly longer, maybe due to the increase of Summit users."*

> *"It takes sometimes an entire day just for a job to start running which seems a little extreme. I also wish there was a longer wall clock."*

> *"Our challenge run, submitted 1 week before the allocation expired, leaving time unused, then the team used the remainder of our time the next week and we have not been able to get more time."*

Other themes included *support and maintenance/updates* (*n* = 8)*, documentation and reference* (*n* = 2)*,* and *miscellaneous* (*n* = 3)*.*

Table 9. *Satisfaction Ratings of Summit by PI Status and Overall*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Scheduling turnaround	83	4.3	0.71	88%	522	4.3	0.71	90%	605	4.3	0.71	90%
System availability	85	4.6	0.60	94%	536	4.5	0.59	95%	621	4.5	0.59	95%
Availability of tools	81	4.5	0.63	93%	524	4.4	0.73	90%	605	4.4	0.71	91%
Availability of libraries	83	4.5	0.69	92%	524	4.4	0.73	89%	607	4.4	0.72	89%
Programming environment	81	4.5	0.69	91%	530	4.4	0.71	90%	611	4.4	0.71	90%
Overall satisfaction with Summit	86	4.6	0.53	98%	542	4.5	0.58	96%	628	4.5	0.57	96%

Table 10. *Satisfaction Ratings of Summit by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Scheduling turnaround	232	4.4	0.65	92%	177	4.4	0.67	92%	114	4.3	0.69	90%	261	4.2	0.77	86%
System availability	237	4.5	0.60	95%	181	4.6	0.58	96%	119	4.6	0.56	97%	268	4.4	0.62	94%
Availability of tools	230	4.5	0.64	94%	178	4.4	0.69	90%	111	4.3	0.72	89%	263	4.3	0.76	89%
Availability of libraries	232	4.5	0.63	93%	179	4.3	0.74	87%	116	4.3	0.72	87%	262	4.3	0.77	88%
Programming environment	235	4.5	0.66	92%	177	4.4	0.72	88%	116	4.3	0.70	89%	266	4.3	0.75	88%
Overall satisfaction with Summit	243	4.6	0.55	97%	182	4.5	0.56	97%	119	4.5	0.55	97%	268	4.4	0.59	96%

Rhea

Rhea was used by 21% of respondents during the 2020 calendar year (*N* = 146). Rhea users were asked to provide satisfaction ratings for multiple aspects of the system, and descriptive statistics of these ratings are shown in Table 11, which also reports satisfaction statistics by PI status. 94% of respondents were either *satisfied* or *very satisfied* with the system overall. Table 12 summarizes these satisfaction statistics by project allocation. The highest rated specific aspect of Rhea was the *programming environment,* and the lowest rated features were *scheduling turnaround* and the *availability of tools.*

Only five Rhea users expressed **reasons for dissatisfaction**, related to modules, capabilities, and running jobs on Rhea nodes:

> *"Because if I do not manually request updated sysmodules for NCO, which I develop, they will not be created. The whole thing should be automated like conda-forge."*
>
> *"I feel that Rhea is too small of a system to handle the postprocessing load of results generated by simulations on Summit. An alternative documented workflow that allows batch processing on Summit would be a good workaround."*
>
> *"Limit on jobs running even when nodes are idle is annoying and wasteful."*
>
> *"Nodes are occupied by a few users, so I had to wait for hours to get my short and small jobs start to run. I understand that it is important to keep the machine usage high, but it will be nice if there is a small reservation for small (1 to 4 nodes) and short jobs."*
>
> *"There were often wait times to get node access. Furthermore, there is no easy way to run Jupyter notebooks on Rhea."*

Table 11. *Satisfaction Ratings of Rhea by PI Status and Overall Totals*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Scheduling turnaround	27	4.4	0.85	93%	110	4.4	0.71	89%	137	4.4	0.73	90%
System availability	27	4.6	0.57	96%	116	4.5	0.65	93%	143	4.5	0.64	94%
Availability of tools	26	4.4	0.75	85%	113	4.5	0.68	91%	139	4.4	0.69	90%
Availability of libraries	26	4.5	0.65	92%	111	4.5	0.63	93%	137	4.5	0.63	93%
Programming environment	26	4.6	0.50	100%	111	4.5	0.63	95%	137	4.5	0.61	96%
Overall satisfaction with Rhea	27	4.6	0.70	89%	116	4.5	0.58	96%	143	4.5	0.60	94%

Table 12. *Satisfaction Ratings of Rhea by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Scheduling turnaround	80	4.4	0.78	89%	59	4.4	0.72	90%	38	4.3	0.87	84%	25	4.4	0.64	92%
System availability	83	4.5	0.65	94%	61	4.6	0.62	93%	38	4.4	0.72	87%	26	4.5	0.51	100%
Availability of tools	80	4.5	0.73	89%	61	4.4	0.67	90%	37	4.3	0.63	92%	26	4.3	0.62	92%
Availability of libraries	79	4.5	0.68	90%	60	4.5	0.60	95%	36	4.4	0.59	94%	26	4.3	0.60	92%
Programming environment	79	4.6	0.61	96%	59	4.6	0.56	97%	36	4.4	0.60	94%	26	4.5	0.51	100%
Overall satisfaction with Rhea	83	4.6	0.61	94%	61	4.5	0.59	95%	38	4.4	0.64	92%	26	4.4	0.50	100%

Data Resources

Data Transfer Nodes

DTNs were used by 28% of respondents during the 2020 calendar year (*N* = 193), and 91% were either *satisfied* or *very satisfied* with the DTNs (Table 13 and Table 14). Six respondents indicated dissatisfaction with DTNs, but only three of these users provided comments explaining their **reasons for dissatisfaction** with the DTNs:

> *"Ease of use."*
>
> *"Slow."*
>
> *"This is not an issue with OLCF support staff, but rather a globus + (most likely) [institution name removed] issue, I bring it up as a use case that came up and I couldn't get a resolution when working with both support staff sides. I was trying to transfer a good sized data set (~10 TB) from[institution], and the transfer rate was extremely low. But neither side (OLCF or [institution]) could pin down the issue (it wasn't an issue transferring the other way, from OLCF to [institution])."*

HPSS

HPSS was used by 18% of respondents during the 2020 calendar year (*N* = 123). HPSS users were asked to provide satisfaction ratings for multiple aspects of the system, and descriptive statistics of these ratings are shown in Table 15, which also reports satisfaction statistics by PI status. 98% of respondents were overall either *satisfied* or *very satisfied* with the system. The highest rated specific aspects of HPSS were *reliability (data integrity)* and the *ability to store/retrieve files.* The lowest rated specific aspect was the *frequency of outages.* Table 16 summarizes these satisfaction statistics by project allocation.

Of the six respondents that indicated dissatisfaction, five reported **reasons for dissatisfaction with HPSS**:

> *"hsi/htar is clunky! No tab completion, unintuitive command line flags, etc."*
>
> *"I could not get globus to work. I mostly use scp to hpss."*
>
> *"It would be nice if the frequency of the outages were less and the duration smaller."*
>
> *"It's hard to use, without many options."*
>
> *"Write times to HPSS seem to have been slower for my projects this year than in the past, and I am not entirely sure why. It seems to be somewhat sporadic, so the issue is normally something I can ignore, but every so often I end up waiting for a transfer to HPSS and it will take more than a day or so."*

Table 13. *Satisfaction Ratings of Data Transfer Nodes by PI Status and Overall*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Overall satisfaction with Data Transfer Nodes	32	4.4	0.95	91%	159	4.4	0.80	91%	191	4.4	0.82	91%

Table 14. *Satisfaction Ratings of Data Transfer Nodes by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Overall satisfaction with Data Transfer Nodes	88	4.4	0.77	91%	77	4.3	0.92	87%	50	4.4	0.83	92%	46	4.2	0.95	85%

Table 15. *Satisfaction Ratings of HPSS by PI Status and Overall*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
hsi/htar interface	24	4.5	0.59	96%	73	4.4	0.78	90%	97	4.5	0.74	92%
Globus interface	22	4.9	0.29	100%	75	4.6	0.69	93%	97	4.7	0.64	95%
Ability to store/retrieve files	26	4.7	0.47	100%	94	4.6	0.58	96%	120	4.6	0.55	97%
Reliability (data integrity)	26	4.7	0.53	96%	92	4.6	0.53	98%	118	4.7	0.53	97%
Time to store/retrieve files	26	4.5	0.81	88%	93	4.5	0.60	95%	119	4.5	0.65	93%
Frequency of outages	23	4.6	0.59	96%	88	4.3	0.77	86%	111	4.4	0.74	88%
Overall satisfaction with HPSS	25	4.7	0.46	100%	94	4.6	0.54	98%	119	4.6	0.52	98%

Table 16. *Satisfaction Ratings of HPSS by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
hsi/htar interface	61	4.5	0.77	90%	33	4.4	0.75	91%	27	4.3	0.73	93%	20	4.3	0.75	95%
Globus interface	56	4.8	0.48	96%	33	4.8	0.36	100%	30	4.6	0.72	93%	20	4.4	0.82	90%
Ability to store/retrieve files	70	4.6	0.54	97%	40	4.7	0.47	100%	34	4.5	0.61	94%	28	4.4	0.57	96%
Reliability (data integrity)	70	4.7	0.52	97%	38	4.7	0.52	97%	33	4.6	0.56	97%	28	4.6	0.50	100%
Time to store/retrieve files	70	4.5	0.68	93%	39	4.5	0.68	95%	34	4.5	0.66	91%	28	4.3	0.61	93%
Frequency of outages	67	4.5	0.68	90%	36	4.5	0.74	92%	33	4.3	0.81	85%	23	4.4	0.58	96%
Overall satisfaction with HPSS	69	4.6	0.52	99%	39	4.7	0.53	97%	34	4.6	0.56	97%	28	4.5	0.58	96%

Alpine GPFS Scratch Filesystem

Alpine GPFS Scratch Filesystem was used by 34% of respondents during the 2020 calendar year (*N* = 231). Alpine GPFS users were asked to provide satisfaction ratings for multiple aspects of the system, and descriptive statistics of these ratings are shown in Table 17, which also reports satisfaction statistics by PI status. 95% of respondents were overall either *satisfied* or *very satisfied* with the system. The *reliability (data integrity)* and *size* were the highest rated specific aspects of Alpine GPFS, and the lowest rated features were the *frequency of outages* and *file and directory operations*. Table 18 summarizes these satisfaction statistics by project allocation.

There were only five users who **indicated dissatisfaction** with at least one aspect of the Alpine GPFS Scratch Filesystem, and four of them provided reasons for their dissatisfaction, primarily related to performance issues:

> *"Frequent stale file system."*
>
> *"Purging timelines are too short and hpss is too complicated to use in workflows (especially since Globus certificates expire)."*
>
> *"Sometimes it freezes, but this behavior has reduced over time, but still freezes sometimes."*
>
> *"There are occasions when simple actions, like listing a directory or open a file, have huge delays in the order of tens of seconds."*

Table 17. *Satisfaction Ratings of Alpine GPFS Scratch Filesystem by PI Status and Overall*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Size	37	4.8	0.40	100%	188	4.5	0.61	94%	225	4.6	0.59	95%
I/O bandwidth	35	4.7	0.52	97%	185	4.5	0.65	92%	220	4.5	0.64	93%
File and directory operations	37	4.5	0.61	95%	186	4.5	0.67	91%	223	4.5	0.66	91%
Reliability (data integrity)	36	4.6	0.64	97%	185	4.5	0.57	96%	221	4.5	0.58	96%
Frequency of outages	36	4.6	0.50	100%	181	4.4	0.70	90%	217	4.4	0.68	91%
Overall satisfaction with Alpine GPFS Scratch filesystem	38	4.7	0.46	100%	189	4.5	0.64	94%	227	4.5	0.62	95%

Table 18. *Satisfaction Ratings of Alpine GPFS Scratch Filesystem by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Size	102	4.6	0.62	93%	86	4.6	0.58	95%	47	4.5	0.59	96%	83	4.6	0.57	96%
I/O bandwidth	100	4.6	0.59	95%	83	4.4	0.70	90%	46	4.4	0.69	89%	80	4.5	0.67	93%
File and directory operations	101	4.5	0.61	94%	85	4.4	0.69	91%	47	4.3	0.73	85%	82	4.5	0.67	90%
Reliability (data integrity)	100	4.6	0.57	98%	86	4.6	0.61	97%	46	4.4	0.58	96%	82	4.5	0.57	96%
Frequency of outages	98	4.4	0.63	93%	81	4.4	0.78	86%	47	4.3	0.74	87%	79	4.4	0.69	91%
Overall satisfaction with Alpine GPFS Scratch filesystem	103	4.5	0.62	95%	87	4.5	0.63	95%	48	4.5	0.58	96%	84	4.5	0.70	93%

Support Services

The Support Services element of the survey was substantially re-written in 2019, and was revised again in 2020 to differentiate between overall satisfaction with support *staff* and details of satisfaction with support *services*. The revised survey addressed the following areas of OLCF service: overall satisfaction with support received, the OLCF website, OLCF communication, OLCF's problem resolution, and the use of OLCF data analysis and visualization tools.

Overall Satisfaction with Support

Users were asked to provide their overall satisfaction with support received from the wide variety of OLCF services available. Most respondents (94%) were either *satisfied* or *very satisfied* with support received from user assistance, accounts, INCITE Scientific Computing Liaisons, and Advanced Data/Workflow Liaisons (Table 19 and Table 20). Only one respondent reported **reasons for dissatisfaction**:

> *"I've requested help on running LLVM 11 OpenMP over the last month. I haven't received the help after 3 emails yet. (Understandably, the support staff may be busy and may need time to get back to me)."*

Table 19. *Satisfaction Ratings of Support Received by PI Status and Overall*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Support received (user assistance, accounts, INCITE Liaisons, Advanced Data/Workflow Liaisons)	86	4.7	0.59	94%	510	4.6	0.61	95%	596	4.6	0.61	94%

Table 20. *Satisfaction Ratings of Support Received by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Support received (user assistance, accounts, INCITE Liaisons, Advanced Data/Workflow Liaisons)	229	4.7	0.60	95%	184	4.6	0.59	95%	119	4.6	0.66	91%	240	4.5	0.63	93%

OLCF Website

Forty-nine percent of survey respondents indicated that they had visited the OLCF website during 2020 (*N* = 334). Before indicating their satisfaction with various aspects of the website, these users were asked how frequently they visit the OLCF website (http://olcf.ornl.gov); only 332 provided responses to this item, as displayed in Figure 7.

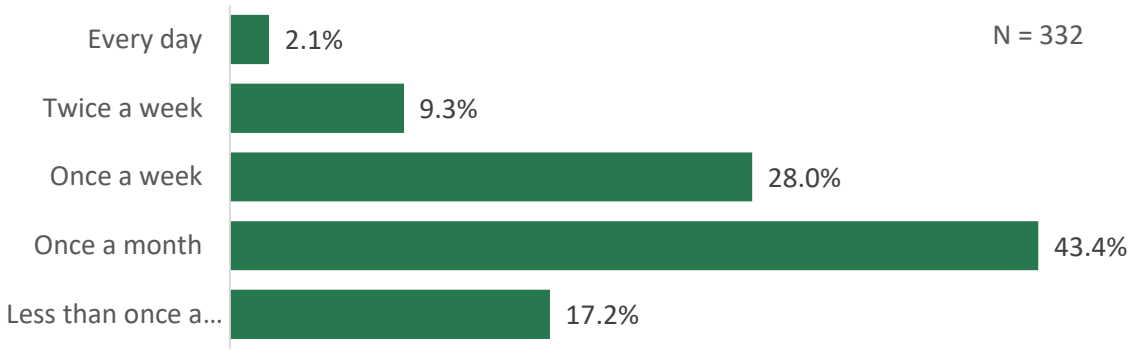

Figure 7. Frequency with which OLCF users visit the OLCF website

Note: Percentages may not add up precisely to 100% due to rounding in each category.

Users rated several aspects of the website (Table 21 and Table 22). 94% of respondents were either *satisfied* or *very satisfied* overall with the website. The highest rated specific aspect of the website was *usefulness of content,* while the lowest rated aspect was *search capabilities.*

There were four users who reported explanations for their **dissatisfaction with one or more aspects of the website** out of the six who indicated dissatisfaction:

> *"All of the allocation tables in My OLCF are non-specific or missing."*
>
> *"It can be a little clunky to find things. You are able to find them but it may take longer than necessary, particularly when you are looking for things like the name of a queue etc."*
>
> *"It is difficult to reach the summit manual from google or from the top page of OLCF. In the manual, I could not find the job class information (e.g., <46 nodes / max 2 hours, <92 nodes / max 6 hours), so I had to search around the previous tutorial slides."*
>
> *"More description is needed for some software available like conda."*

Communication with Users

As seen in Table 23, most respondents (96%) were either *satisfied* or *very satisfied* with how OLCF keeps them informed of changes, events, downtimes, and current issues (Table 23 and Table 24). None of the respondents indicated **dissatisfaction with OLCF communication.**

Table 21. *Satisfaction Ratings of the OLCF Website by PI Status and Overall Totals*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Ease of navigation	51	4.3	0.62	92%	275	4.3	0.67	88%	326	4.3	0.66	88%
Search capabilities	45	4.2	0.77	84%	257	4.2	0.72	84%	302	4.2	0.73	84%
Usefulness of content	51	4.4	0.61	94%	276	4.5	0.60	96%	327	4.4	0.60	95%
Overall satisfaction with the OLCF website	51	4.4	0.60	94%	277	4.4	0.61	94%	328	4.4	0.60	94%

Table 22. *Satisfaction Ratings of the OLCF Website by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Ease of navigation	135	4.3	0.70	87%	106	4.3	0.68	88%	70	4.3	0.66	89%	123	4.2	0.63	89%
Search capabilities	128	4.2	0.78	81%	95	4.1	0.83	76%	64	4.2	0.77	84%	117	4.1	0.70	85%
Usefulness of content	134	4.5	0.63	94%	106	4.4	0.62	95%	71	4.4	0.64	94%	124	4.4	0.60	96%
Overall satisfaction with the OLCF website	135	4.4	0.66	92%	106	4.3	0.62	92%	71	4.4	0.61	93%	124	4.4	0.58	95%

Table 23. *Satisfaction Ratings of Communication by PI Status and Overall Totals*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Overall communications	89	4.7	0.52	98%	570	4.5	0.59	95%	659	4.6	0.58	96%

Table 24. *Satisfaction Ratings of Communications by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Overall communications	255	4.6	0.57	97%	196	4.6	0.54	97%	124	4.5	0.63	93%	267	4.5	0.58	96%

Problem Resolution

Figure 8 shows how frequently respondents **submitted queries** to OLCF (via phone or email) in 2020. Over half submitted between one and five requests, while approximately **one-third had not submitted any queries at all.**

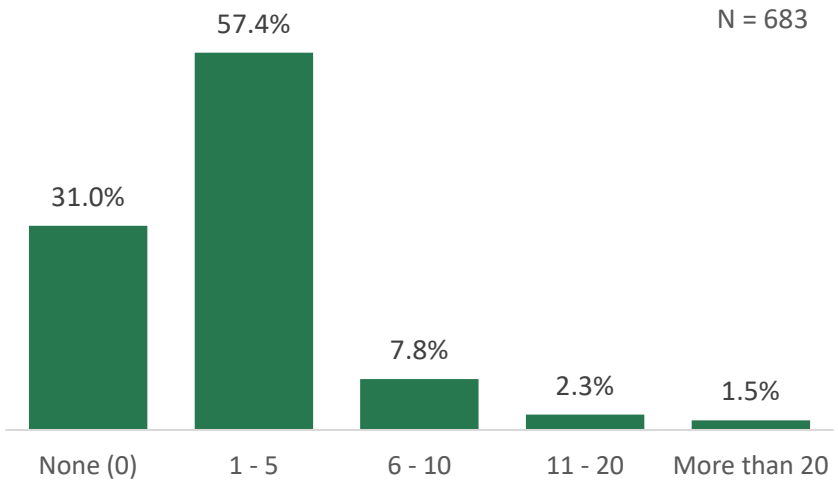

Figure 8. Distribution of number of queries submitted to OLCF in 2020

Note: Percentages may not add up precisely to 100% due to rounding in each category.

Users were asked to provide satisfaction ratings for their overall satisfaction with OLCF's problem resolution and three specific aspects (Table 25 and Table 26). 94% of respondents were overall either *satisfied* or *very satisfied* with problem resolution. The *quality of OLCF response to reported issues* was the highest rated specific aspect, while the *usefulness of support and training documentation* and *timeliness of OLCF responses to reported issues* were rated only slightly lower. Nine respondents reported **reasons for dissatisfaction with problem resolution.** The most common reasons for dissatisfaction were related to *documentation* (*n* = 3) and to *ticketing and follow-up* (*n* =3). Illustrative responses include:

> *"Documentation on tools and compilers is poor."*
>
> *"Email responses, even acknowledgement and ticketing, same day would be better."*
>
> *"The support ticket system should be OPEN to allow users to see and update their current ticket status."*

Table 25. *Satisfaction Ratings of OLCF's Problem Resolution by PI Status and Overall*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Quality of OLCF response to reported issues	78	4.5	0.57	96%	443	4.6	0.63	93%	521	4.6	0.62	94%
Timeliness of OLCF responses to reported issues	78	4.5	0.66	94%	443	4.6	0.64	93%	521	4.6	0.64	93%
Usefulness of support and training documentation	79	4.5	0.62	94%	460	4.5	0.65	93%	539	4.5	0.64	93%
Overall satisfaction with problem resolution	79	4.5	0.62	94%	450	4.5	0.64	94%	529	4.5	0.64	94%

Table 26. *Satisfaction Ratings of OLCF's Problem Resolution by Project Allocation*

	INCITE				DD				ALCC				ECP			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Quality of OLCF response to reported issues	202	4.6	0.59	95%	163	4.5	0.63	93%	102	4.5	0.69	91%	214	4.6	0.62	94%
Timeliness of OLCF responses to reported issues	202	4.6	0.65	95%	163	4.5	0.66	92%	102	4.5	0.68	91%	214	4.6	0.62	94%
Usefulness of support and training documentation	211	4.5	0.63	94%	164	4.5	0.62	93%	103	4.5	0.65	91%	221	4.5	0.66	93%
Overall satisfaction with problem resolution	206	4.6	0.59	96%	164	4.5	0.65	91%	105	4.5	0.68	91%	217	4.5	0.64	95%

Data Analysis and Visualization

Respondents were not asked whether they used data analysis and visualization services and were not asked for overall satisfaction ratings of this aspect of OLCF.

Users were asked to indicate where they analyze data. Figure 9 shows that the largest proportion of users analyzed all or most of their data "elsewhere" and the smallest proportion analyzed it all at OLCF.

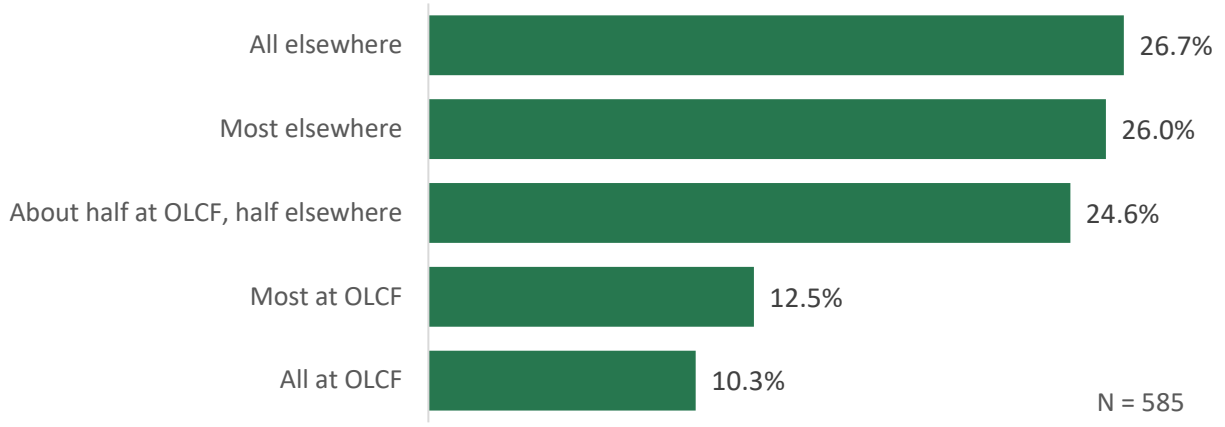

Figure 9. Locations for analysis of data by OLCF users

Note: Percentages may not add up precisely to 100% due to rounding in each category.

To put these results in context, users were also asked about the source of their data, displayed in Figure 10. The largest proportion of users are working with data that is primarily sourced from OLCF jobs, but an almost equal proportion work with data mostly from outside OLCF.

Figure 10. Source of user data

Note: Percentages may not add up precisely to 100% due to rounding in each category.

Finally, users were asked whether they were interested in scheduling one-on-one conversations with OLCF analysis and visualization specialists in order to consult on needs and approaches. 52 of those who provided a response to this question (8%, *N* = 681) indicated they were interested, and their contact information was referred to the OLCF to arrange consultations.

User Suggestions for Improvement

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

OLCF Experience

When asked "What additional services, resources, and/or other improvements are needed to enhance your experience at the OLCF?" 141 respondents supplied comments; 4% indicated satisfaction, i.e., that *no additional services and/or resources* are needed to enhance their experience at the OLCF (Table 27). Among those expressing a need or preference, *tools, software, libraries, installations, and updates* was mentioned most frequently, followed by *Summit* and *jobs, queue, and scheduling policy*. See Appendix E: User Suggestions for Improvement for all responses by category. Select comments include:

Tools, software, libraries, installations, and updates	*"Keeping libraries (such as CUDA, PyTorch, etc.) more up-to-date would be appreciated."* *"More libraries/options, making it easier to add different MPI implementations, etc."*
Summit	*"More courses or talks on profiling tools for Summit (multi-node)."* *"The login nodes on Summit hang a lot this year."*
Jobs, queue, and scheduling policy	*"A faster scheduling of interactive jobs."* *"Sometimes turnaround for debugging runs is very long."* *"Support for short/small jobs that are necessary in the overall workflow."* *"The job submission system is a little bit complicated."*
Development, debugging, and test resources/queue	*"Debugging queue and small job testing server will be great."* *"It can take some time to get an interactive session for those debug/test circumstances where a single node would be useful to solve difficult problems. More availability for nodes for this type of work would be useful."*
Documentation	*"Better user documentation, i.e. how resource pools are allocated and tools for launching jobs."* *"I understand Rhea is to be decommissioned soon, but the documentation could have been more extensive, precisely in terms of the relevant slurm script keywords."*

Table 27. *Users' Suggestions for Additional Services and/or Resources Needed to Enhance Their Experience at the OLCF*

Category	*N* = 141	Percentage
Tools, software, libraries, installations, and updates	36	26%
Summit	17	12%
Jobs, queue and scheduling policy	16	11%
Development, debugging and test resources/queue	15	11%
Documentation	15	11%
Performance, hardware, or failures/hanging/reliability	11	8%
Accounts, user portal, login, security & authentication	10	7%
Tech support/assistance and ticketing	10	7%
Allocations, usage, and project monitoring	9	6%
Training and tutorials	9	6%
Filesystem, I/O and data transfer	8	6%
Visualization, preprocessing and analysis	8	6%
Data storage, backup, retention, and purge	7	5%
GPU resources	6	4%
Satisfaction	6	4%
Wall/run time	5	4%
Architecture	4	3%
OLCF community	2	1%
Miscellaneous/other	12	9%

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

Compute or Data Resources

When asked, "Please describe how the OLCF can improve your experience using any of the HPC resources (i.e., Summit, Rhea, DTNs, HPSS, Alpine GPFS) and/or tell us if any additional resources are needed," 98 respondents provided comments. Of these respondents, 18%, expressed their *satisfaction* with OLCF compute and data resources. The highest proportion of users provided suggestions related to *Summit* (21%), *file systems, data storage, and data transfer* (19%), and *documentation, training, tutorials and/or resources* (11%). See Appendix E: User Suggestions for Improvement for all responses by category. Refer to Table 28 for all themes identified.

Select comments include:

Summit	*"On rare occasion, a job will terminate sporadically on Summit due to inability to access GPU resources. I'd love to see some NVIDIA A100 GPUs become available on Summit. Sometimes, queue times are a little long."* *"I've had builds fail on the Summit home directory because of file system failure. In some cases I've found I need to build in a temporary directory (which is presumably mounted on a local disk) to get it to work. This is not a performance issue, this is outright failure. Errors are non-deterministic. One symptom of the problem is that when a build fails, I cannot delete the build directory where the failure occurred because it is locked. (Usually, after some period of time, like 24 hours, then I can delete the directory.) Instead I have to move it aside to try again."*
File systems, data storage, and data transfer	*"As with the Lustre file system on Titan, GPFS on Summit seems to routinely have heavy lag that I've not experienced at other leadership-class facilities. Still, it's an annoyance rather than a blocker. Additionally, with users being responsible for more software installs these days, the project storage is not sufficient. We end up installing all our software on GPFS and keeping all our datasets on GPFS, and built a pipeline to restore stuff after purges. It's workable but not ideal."* *"Space is limited especially for large-scale I/O experiments."* *"Very happy with HPSS proper. My project needs are essentially incremental backups of experimental data over multiyear durations. Are there any book keeping tools well integrated with HPSS and Globus to assist in this?"*
More documentation, training, tutorials, and/or resources	*"A detailed and thorough guide on mpi implementations via python and R for summit and rhea available via website would be highly convenient. If there is such a guide, I was unable to find it."* *"For someone getting started, the documentation is scattered. "Getting started" guides should be more comprehensive."* *"It would be nice if there is a webpage to monitor the status of HPC resources in OLCF. ALCF is managing such a webpage (status.alcf.anl.gov) It might be great if OLCF has such a webpage. I think the documentation for Rhea is not rich comparing to the Summit. For example, it will be nice to have more documents for heavy Python users, e.g. MPI4PY."*

Table 28. *Users' Suggestions for Improvements to HPC Compute and Data Resources*

Category	*N* = 98	Percentage
Summit	21	21%
File systems, data storage, and data transfer	19	19%
Positive comments/specific expressions of satisfaction	18	18%
More documentation, training, tutorials and/or resources	11	11%
Issues with software/tools/modules	10	10%
Need for specific software/tools/modules	9	9%
Performance/reliability/stability/uptime of systems	9	9%
Data retention/purge policy and communication about purges	8	8%
Interactive data access, visualization, and processing	8	8%
Wall time/runtime and queues	8	8%
Monitoring project allocations and usage/user portal/MyOLCF	7	7%
Rhea/Titan/Eos functionality replacements	6	6%
Support and ticketing	6	6%
Andes	5	5%
Allocations and quotas	3	3%
SSH/accessibility and authentication	2	2%
Miscellaneous/Other	5	5%

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

Data Analysis, Visualization, and/or Workflow

When asked, "What additional data analysis, visualization, and/or workflow services would you like the OLCF to provide?" 100 users responded. Because this item asks respondents to provide concrete recommendations, respondents who provided only a comment indicating satisfaction with current services, and did not provide any additional suggestions, were not included in this count. Among those 100 users who expressed a need/preference, the largest proportions were interested in a variety of analysis and visualization tools/software. Large proportions also specifically mentioned *Jupyter, ParaView, VisIt,* and *Python*-related capabilities; tools and software would have dominated responses if these types of responses were combined into a single category (Figure 11). *Tutorials, documentation, and training* were the second most common suggestions.

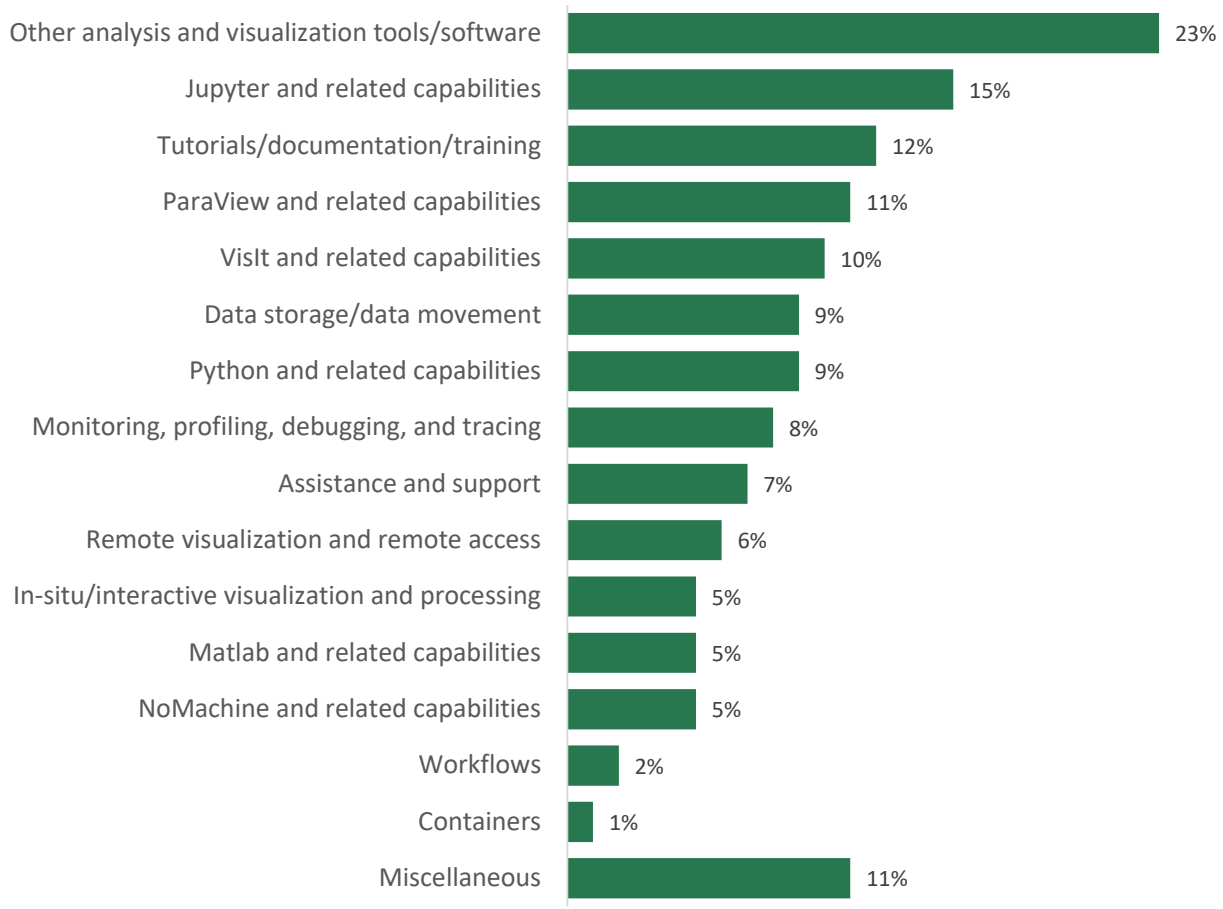

Figure 11. Users' suggestions for additional data analysis, visualization, and/or workflow services

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

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

Other analysis and visualization tools/software	*"Parallelized visualization tools and training for those. Most of my issues on Rhea are memory related when working with large files."* *"Blender implementation on Rhea (or Andes?)? for command line rendering. Data will be exported through other visualization tools already installed (i.e. VisIt), or generated in Blender for analysis."* *"I'd like to use tools like HPCToolkit very easily with applications that use OpenMP offload."*
Jupyter and related capabilities	*"Jupyter notebook frontends have been very useful to us at other computing centers, exploring data with an easy interface."* *"Web interface for Jupyter notebooks."* *"Analysis using Jupyter. I wasn't able open Jupyter on Summit."*
Visualization and analysis tutorials/documentation/training	*"After the pandemic, I would be very interested in hands-on workflow trainings. I am very interested in having workflows run automated analyses after simulations finish, but I haven't had a chance to look into it. I've been to a few talks, but they've been more conceptual than concrete. I also am still in the process of writing robust analysis scripts that would be effective in a workflow."* *"I would be interested in learning about possibilities using VR for visualization of my data."* *"Tutorials on compiling and benchmarking of codes."*

Other OLCF Issues

When asked to comment on any **additional areas of importance not covered elsewhere** in the survey, 39 individuals replied. The largest proportion expressed *satisfaction/positive comments or highlights* (31%), followed by comments related to *staff and user services* (13%) and *Summit* (13%). Other comments were distributed as seen in Table 29 (refer to Appendix E: User Suggestions for Improvement for text of these comments).

Table 29. *Respondent Comments on Other Issues Not Addressed within the Survey*

Category	***N* = 39**	**Percentage**
Satisfaction/positive comments or highlights	12	31%
Staff and user services	5	13%
Summit	5	13%
Allocations, queue policy, and job prioritization	4	10%
Software/libraries/compilers/tools	4	10%
Accounts, access and credentials	3	8%
Development and debugging needs	2	5%
Purge policy	2	5%
Support materials and training	2	5%
Visualization, remote viz, and interactive sessions	2	5%
Documentation/website	1	3%
Performance and support issues	1	3%
Suggestions for additional survey queries/topics	1	3%
Miscellaneous	7	18%

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

Summary of Survey Observations

In most respects, users were highly satisfied with the OLCF resources/services. Table 30 summarizes satisfaction (*satisfied* or *very satisfied*) ratings. The color scale indicates the relative magnitude of cell values: high-med-low = green-yellow-red. Examination of the table suggests that **satisfaction was highest** (across respondent types) for Summit, HPSS, Alpine GPFS Scratch Filesystem, and communications. While the **lowest ratings** were found for data transfer nodes, support staff, support services, and support/training documentation, these ratings still reflect a generally high satisfaction among users. When "All" respondents are considered as a group, all items were rated as either *satisfied* or *very satisfied* by 91% or more of users.

Table 30. *Summary of Overall Satisfaction with Aspects of OLCF, by PI Status and Project Allocation*

		PI Status		Project Allocation			
	All	PI	Non-PI	INCITE	DD	ALCC	ECP
Max N responding:	688	93	595	267	200	126	284
OLCF	**97%**	**98%**	**96%**	**97%**	**98%**	**97%**	**96%**
Compute Resources	**96%**	**97%**	**96%**	**96%**	**96%**	**97%**	**95%**
Rhea	94%	89%	96%	94%	95%	92%	100%
Summit	96%	98%	96%	97%	97%	97%	96%
Data Resources	**93%**	**98%**	**93%**	**94%**	**95%**	**95%**	**91%**
Data Transfer Nodes	91%	91%	91%	91%	87%	92%	85%
HPSS	98%	100%	98%	99%	97%	97%	96%
Alpine GPFS Scratch Filesystem	95%	100%	94%	95%	95%	96%	93%
Support Staff	**93%**	**92%**	**93%**	**94%**	**95%**	**93%**	**92%**
Support Services	**92%**	**90%**	**92%**	**92%**	**93%**	**92%**	**90%**
Support received*	94%	94%	95%	95%	95%	91%	93%
OLCF website	94%	94%	94%	92%	92%	93%	95%
Communications	96%	98%	95%	97%	97%	93%	96%
Support/training documentation	93%	94%	93%	94%	93%	91%	93%
Problem resolution	94%	94%	94%	96%	91%	91%	95%
Min	**91%**	**89%**	**91%**	**91%**	**87%**	**91%**	**85%**
Max	**98%**	**100%**	**98%**	**99%**	**98%**	**97%**	**100%**

*"Support received" lists out and includes the following support areas: User Assistance, Accounts, INCITE Scientific Computing Liaisons, Advanced Data/Workflow liaisons.

Longitudinal Comparisons of User Responses

This section reviews the results from the 2006 through 2020 OLCF User Surveys and reports information about long-term response trends related to respondent years of experience with OLCF, project allocations, and overall satisfaction with OLCF.

OLCF Users

Figure 12 shows that **length of time using OLCF** (i.e., experience as an OLCF user) reported by most survey respondents has changed substantially between 2006 and 2020. Prior to 2009, about half of respondents reported using OLCF less than one year, and this category comprised the largest proportion of users. However, between 2009 and 2011, the largest proportion of respondents indicated having greater than two years of experience at OLCF. In 2012, user experience shifted back to the largest proportion of respondents reporting using OLCF less than one year. From 2013 to 2020, users who had been with OLCF for more than two years once again made up the greatest proportion of respondents.

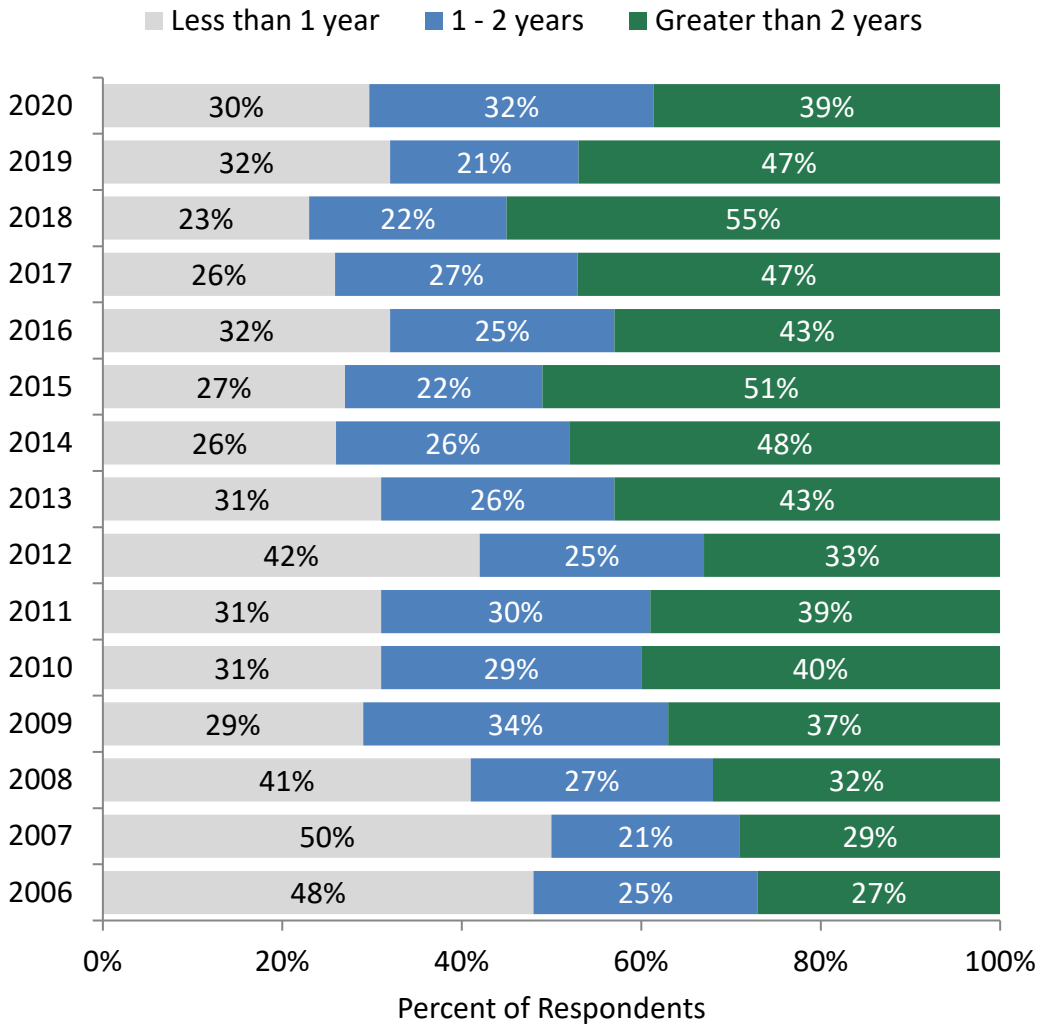

Figure 12. Respondent years of experience with OLCF, 2006-2020

Note: Percentages may not add up precisely to 100% due to rounding in each category.

With respect to **project classifications** (Figure 13), *survey respondent* data is available from 2007 to the present year, and OLCF data for *the entire pool of OLCF users* is available from 2014 to present. The figure shows these side-by-side and indicates that the distribution of respondents has tracked the overall user pool. In 2018, the ECP project allocation was added to the dataset. Note that this longitudinal tracking excludes "other" project classifications, such as NOAA projects, General projects, Staff projects, and 2019's Early Science (ES) projects.

Generally,

- Until 2017, INCITE projects have shown a downward trend in share of both the respondent and the user pool. In 2017 through 2019, INCITE projects have grown only modestly in the user pool and substantially in the respondent pool, before reversing course in 2020.
- Director's Discretion projects remained relatively constant between 2007 and 2011 for respondents, and generally trended upward between 2012 and 2018. DD decreased substantially in the respondent and user pools in 2019, and remained relatively stable in 2020. Notably, in both 2019 and 2020, DD was noticeably underrepresented among survey respondents.
- ASCR Leadership Computing Challenge (ALCC) supported projects began in 2010 and supported only 2% of respondents, but grew significantly by 2014. Among both users and respondents, there has been a dip since 2018.
- Exascale Computing Project (ECP) supported projects began in 2018 and have grown since then among both users and respondents.

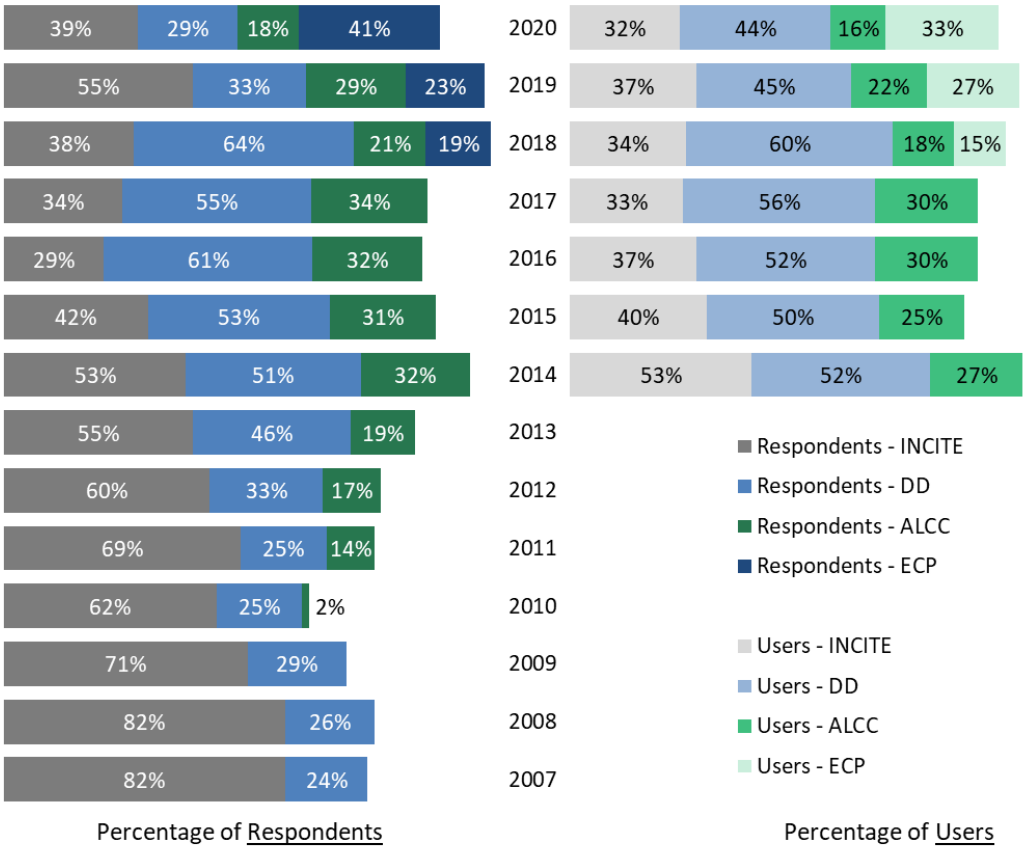

Figure 13. Survey respondent project allocations, 2007-2020, and OLCF user project allocations, 2014-2020

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

Satisfaction with OLCF Overall

With regard to **overall satisfaction with OLCF**, the percent of *very satisfied* respondents showed a nearly uninterrupted trend upward from 2007 to 2018. The proportion of *very satisfied* respondents more than doubled from the 2007 value to 69% in 2017 and 70% in 2018 (Figure 14). The exceptions to this trend were moderate decreases in 2011 and 2012. In 2019 and again in 2020, the proportion of *very satisfied* respondents dropped to 59%, with a noticeable shift toward respondents selecting the *satisfied* option. The overall proportion of respondents indicating satisfaction (*satisfied* and *very satisfied* responses) has grown as well, from 91% in 2012 to 94-97% in each year from 2013 to 2020. The proportion of respondents indicating satisfaction was 97% in 2020, tied with 2014 for the year with the most satisfied user base.

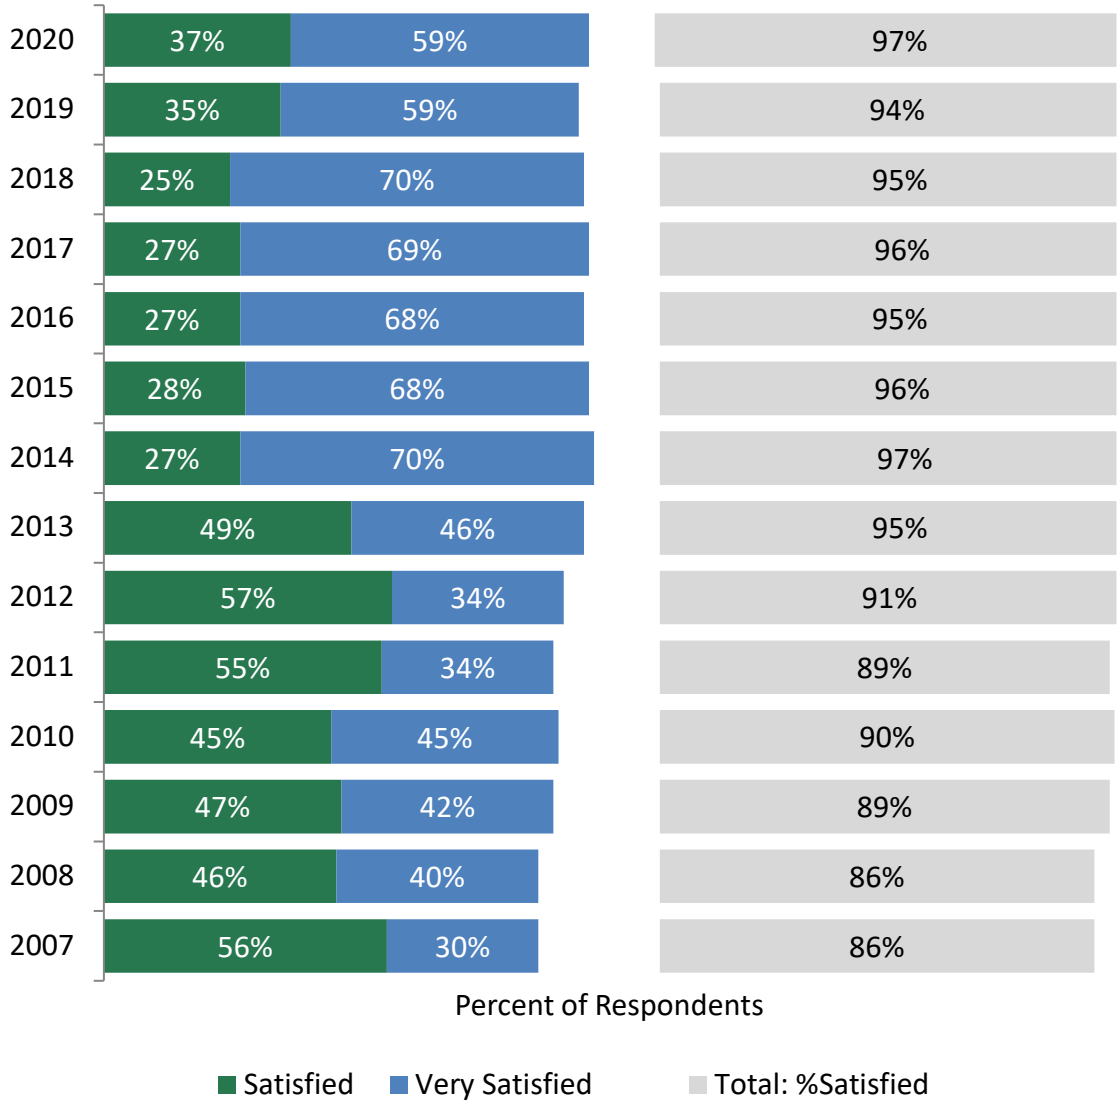

Figure 14. Proportion of respondents reporting being *satisfied* and *very satisfied* overall with OLCF and the total of %Sat respondents, 2007-2020

Note: Indicated percentages may not add up precisely to %Satisfied due to rounding in each category. In 2020, 37.4% and 59.3% of respondents, respectively, were *satisfied* or *very satisfied*; this rounds to 97% satisfaction.