2019 Oak Ridge Leadership Computing Facility User Survey

Findings and Recommendations

PREPARED BY:

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

Kelly P. Townsend, MSBA
Evaluation Specialist
Scientific Assessment & Workforce Development

Kate Miller-Bains, Ph.D.
Evaluation Specialist
Scientific Assessment & Workforce Development

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

PREPARED FOR:

Oak Ridge Leadership Computing Facility
Oak Ridge National Laboratory
February 2020

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

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

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

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

ORAU provides innovative scientific and technical solutions to advance national priorities in science, education, security and health. Through specialized teams of experts, unique laboratory capabilities and access to a consortium of more than 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

Table of Contents iii
List of Tables v
List of Figures vi
Executive Summary vii
Introduction 1
Data Collection and Analysis 1
- Survey Revisions 1
- Data Collection 2
- Data Analysis 3

Findings 7
- Respondents 7
- Resource Utilization 8
- Overall Satisfaction 11
- Compute and Data Resources 19
- Compute Resources 22
  - Summit 22
  - Titan 24
  - Eos 26
  - Rhea 26
- Data Resources 29
  - Data Transfer Nodes 29
  - HPSS 29
  - Lustre/Spider Scratch Filesystem 32
  - Alpine GPFS Scratch Filesystem 32
- Support Services 35
  - Overall Satisfaction with Support 35
  - OLCF Website 37
  - Communication with Users 37
  - Problem Resolution 40
  - Data Analysis and Visualization 42
- User Suggestions for Improvement 43
  - OLCF Experience 43
  - Compute or Data Resources 44
  - Data Analysis, Visualization, and/or Workflow 47
  - Other OLCF Issues 48
- Summary of Survey Observations 49

List of Tables

Table 1. *Major Categories and Subcategories Used to Organize Open-Ended Responses* 5
Table 2. *Project Allocations by OLCF Users and Survey Respondents* 7
Table 3. *HPC and Support Resources Used by PI status, Project Allocation and Overall Totals* 10
Table 4. *Overall Satisfaction with OLCF and Its Major Resources/Services by PI Status and Totals* 12
Table 5. *Overall Satisfaction with OLCF and Its Major Resources/Services by Project Allocation* 13
Table 6. *Overall Satisfaction with OLCF and Its Major Resources/Services by Project Allocation (Continued)* 14
Table 7. *Best Qualities of OLCF (ordered by % of all respondents, high to low)* 18
Table 8. *Satisfaction Ratings for Features of the OLCF HPC Compute and Data Resources by PI Status and Overall Totals* 20
Table 9. *Satisfaction Ratings for Features of the OLCF HPC Compute and Data Resources by Project Allocation* 20
Table 10. *Satisfaction Ratings of Summit by PI Status and Overall* 23
Table 11. *Satisfaction Ratings of Summit by Project Allocation* 23
Table 12. *Satisfaction Ratings of Titan by PI Status and Overall* 25
Table 13. *Satisfaction Ratings of Titan by Project Allocation* 25
Table 14. *Satisfaction Ratings of Eos by PI Status and Overall* 27
Table 15. *Satisfaction Ratings of Eos by Project Allocation* 27
Table 16. *Satisfaction Ratings of Rhea by PI Status and Overall Totals* 28
Table 17. *Satisfaction Ratings of Rhea by Project Allocation* 28
Table 18. *Satisfaction Ratings of Data Transfer Nodes by PI Status and Overall* 30
Table 19. *Satisfaction Ratings of Data Transfer Nodes by Project Allocation* 30
Table 20. *Satisfaction Ratings of HPSS by PI Status and Overall* 31
Table 21. *Satisfaction Ratings of HPSS by Project Allocation* 31
Table 22. *Satisfaction Ratings of Lustre/Spider Scratch Filesystem by PI Status and Overall* 33
Table 23. *Satisfaction Ratings of Lustre/Spider Scratch Filesystem by Project Allocation* 33
Table 24. *Satisfaction Ratings of Alpine GPFS Scratch Filesystem by PI Status and Overall* 34
Table 25. *Satisfaction Ratings of Alpine GPFS Scratch Filesystem by Project Allocation* 34
Table 26. *Satisfaction Ratings of Support Received by PI Status and Overall* 36
Table 27. *Satisfaction Ratings of Support Received by Project Allocation* 36
Table 28. *Satisfaction Ratings of the OLCF Website by PI Status and Overall Totals* 38
Table 29. *Satisfaction Ratings of the OLCF Website by Project Allocation* 38
Table 30. *Satisfaction Ratings of Communication by PI Status and Overall Totals* 39
Table 31. *Satisfaction Ratings of Communications by Project Allocation* 39
Table 32. *Satisfaction Ratings of OLCF's Problem Resolution by PI Status and Overall* 41
Table 33. *Satisfaction Ratings of OLCF's Problem Resolution by Project Allocation* 41
Table 34. *Users' Suggestions for Additional Services and/or Resources Needed to Enhance Their Experience at the OLCF* 44
Table 35. *Users' Suggestions for Improvements to HPC Compute and Data Resources* 46

Figure 1. Respondent occupational affiliation (*N* = 578) 7
Figure 2. Project allocations for OLCF Users (*N* = 1251) and for Respondents (*N* = 578) 8
Figure 3. PI Status for OLCF Users (*N* = 1251) and for Respondents (*N* = 578) 8
Figure 4. Experience using the OLCF (*N* = 578) 9
Figure 5. "Overall" Satisfaction with OLCF and its major resources/services (maximum *N* = 578) 11
Figure 6. Perceived changes from FY 2018 computing/data resources performance by years using OLCF (*N* = 390) 21
Figure 7. Frequency with which OLCF users visit the OLCF website (*N* = 305) 37
Figure 8. Distribution of number of queries submitted to OLCF in 2019 (*N* = 574) 40
Figure 9. Locations for analysis of data by OLCF users (*N* = 499) 42
Figure 10. Source of user data (*N* = 574) 42
Figure 11. Users' suggestions for additional data analysis, visualization, and/or workflow services (*N* = 82) 47

Executive Summary

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

Respondents

At the end of the nine-week survey period, 578 users completed the survey out of 1,251 possible respondents, giving an overall response rate of 46.2%. Respondents' projects were supported by INCITE (55%), Director's Discretion (33%), ALCC (29%), ECP (23%), ES (9%) and General (5%) sources.

Findings Highlights

Overall Evaluation

The proportions of all respondents *satisfied* or *very satisfied* with OLCF resources/services, ranged from 90% to 98% for "overall" evaluation items. Specifically, ratings for major categories of resources/services were a) OLCF (94%), b) Compute Resources (96%), c) Data Resources (95%), and d) Support Services (91%). Thematic analysis of open-ended comments identified *computing power/hardware/HPC resources* (46% of respondents) and *user support/staff* (37% of respondents) as the most valued OLCF qualities.

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

	All	PI	Non-PI	INCITE	DD	ALCC	ECP	ES*	General*
Max N responding:	*578*	*102*	*476*	*320*	*189*	*170*	*135*	*51*	*27*
OLCF	**94%**	**94%**	**94%**	**93%**	**95%**	**94%**	**96%**	**94%**	**100%**
Compute Resources	**96%**	**96%**	**96%**	**94%**	**97%**	**97%**	**98%**	**98%**	**100%**
Titan	91%	95%	91%	90%	91%	92%	85%	88%	91%
Eos	90%	93%	89%	83%	86%	92%	100%	--	--
Rhea	96%	94%	97%	95%	98%	95%	94%	92%	100%
Summit	95%	94%	95%	96%	94%	94%	98%	96%	100%
Data Resources	**95%**	**95%**	**95%**	**94%**	**96%**	**95%**	**95%**	**100%**	**100%**
Data Transfer Nodes	94%	83%	97%	95%	94%	97%	84%	85%	100%
HPSS	98%	97%	98%	99%	98%	97%	100%	100%	100%
Lustre/Spider Scratch Filesystem	91%	88%	92%	92%	86%	85%	84%	77%	89%
Alpine GPFS Scratch Filesystem	97%	97%	97%	97%	96%	96%	96%	95%	100%
Support Services	**91%**	**94%**	**91%**	**91%**	**94%**	**94%**	**93%**	**92%**	**96%**
Support received**	93%	93%	94%	94%	94%	93%	96%	84%	96%
OLCF website	93%	93%	93%	92%	96%	92%	96%	82%	100%
Communications	95%	98%	94%	94%	97%	94%	95%	90%	96%
Support/training documentation	91%	92%	91%	90%	92%	93%	93%	93%	100%
Problem resolution	93%	92%	94%	94%	94%	93%	93%	95%	100%

*27 respondents are allocated to "General" projects and 51 respondents to "ES" projects; as a result, some questions received too few responses for SD or %Sat to be meaningful. **Support received lists out: user assistance, accounts, INCITE Liaisons, Advanced Data/Workflow Liaisons.

OLCF Systems, Data Resources, and Compute Resources

Summit, which became available to users in 2019, was utilized by 76% of users. Due to the availability of Summit, Titan and Eos were used at a decreased rate compared to results from 2014-2018. Rhea, however, was used still used at a similar rate compared to past years. Half of users (51.8%) noted no changes in overall OLCF computing performance over the last year, while 47.7% cited improved performance; only two users (0.5%, *N* = 390) noted a decrease in performance compared to FY 2018. Overall satisfaction across the compute resources and data resources ranged from 90% (Eos) to 98% (HPSS) of users either *satisfied* or *very satisfied*. 94% of users were *satisfied* or *very satisfied* with both notice for scheduled maintenance 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 (93%) 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

95% of respondents were overall *satisfied* or *very satisfied* with how OLCF keeps them informed of changes, events, downtimes, and current issues, up from 91% in FY 2018.

Problem Resolution

Half (53%) of respondents submitted between one and five queries to OLCF (via phone or email) in 2019. 93% 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* (93% satisfaction) followed by *timeliness of OLCF responses to reported issues* (92% satisfaction) and *usefulness of support and training documentation* (91% satisfaction).

Website

Half (53%) of respondents indicated that they had visited the OLCF website during 2019. 43% of respondents indicated they visit the OLCF website (http://olcf.ornl.gov) once a week or more frequently. More than 9 in 10 respondents indicated they were satisfied with the OLCF Website (93%). The highest rated aspect of the OLCF website was the *usefulness of content* (94% satisfaction). *Search capabilities* were the lowest rated aspect of website usability (81% satisfaction).

Data Analysis and Visualization

13% of 574 respondents who answered this question indicated they do not need or use data analysis from OLCF and were excluded from the following results. 22% of respondents indicated they analyze *most* or *all* of their data at OLCF while 56% analyze most or all of their data elsewhere. 21% 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 are working with data that is primarily (*most* or *all*) sourced from OLCF jobs (66%). Thirty-six (9%, *N* = 414) 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.

Recommendations

OLCF Resources/Services

Recommendations offered here are based on examination of the relative satisfaction ratings, respondent reasons for dissatisfaction, and user recommendations for OLCF improvement. Note that since the satisfaction ratings across resources/services were relatively consistent and typically 90% or higher, recommendations for change are best found in the **expressed reasons** for user dissatisfaction in conjunction with their **suggestions for improvement**.

As expected, many responses to open-ended questions noted an interest in *additional* systems, tools, libraries, and other software resources. While many such resources were cited, Jupyter and Python-related software and capabilities were particularly prominent. Many other responses focused on the new Summit system and needs for support and tools for Summit. Development, debugging, and testing tools have been requested across multiple years of the survey, but users in 2019 noted particular needs for these capabilities on Summit.

Across multiple elements of the survey, many users asked for *more assistance* or specific kinds of support (ranging from specialists in particular areas of HPC, to startup guides/tutorials, to documentation about OLCF systems, capabilities, and common tasks taking advantage of OLCF tools and software). In addition, some users were interested in public documentation of support, such as a ticketing system visible to all OLCF users, so that others could benefit from the lessons learned and the solutions tested in particular situations. Finally, issues of access, including multi-factor authentication, use of SSH, and difficulty with cross-system workflows due to both MFA and SSH were raised across the survey responses. In terms of performance, lags and hang-ups on the file system were cited.

Examination of overall satisfaction suggests that data resources requiring the most attention may be the support services, including documentation, the website, and support/problem resolution. Among the compute systems, Titan and Eos received the lowest rating; however, no users indicated they were *dissatisfied* with Eos and therefore no users provided comments, and Titan comments cited hardware failures.

These findings reflect a range of open-ended comments that called attention to the ever-changing nature of both HPC and the OLCF, and the challenges users face in becoming acquainted with new systems and tools.

Support Services	
Support/training documentation	Users are interested in many types of support, training, and documentation. Some comments indicated a need for clearer documentation including tutorials, information about "common mistakes" or tips and tricks from other users, and general documentation about specific systems or software. The need is particularly urgent for new systems like Summit and during transitions when users may encounter difficulty running existing workflows. Summit was particularly cited as an area of OLCF requiring more thorough documentation.
OLCF website	Concerns with the website centered on difficulty searching for or navigating *to* useful information, indicating that some information is hard to find. Other comments noted an interest in updated or more informative content (such as the documentation and other material noted above).
Problem resolution	While very few users provided explanations for their dissatisfaction with problem resolution, their concerns centered around support queries that went unanswered and, for some, a sense that they have had to solve problems on their own. Users are also interested in a ticketing system or documentation of resolved issues that can be of use to teams beyond those who directly received the support/guidance.
Overall support received	Several users cited specific areas of HPC, and technical needs for their own projects, for which they desire specialized support and the ability to consult with an expert or receive ongoing help with their workflows.
Compute Resources	
Summit	Users noted limitations to Summit when it first came online, as well as difficulty running existing software, code, or workflows to accommodate project needs. For some users, Summit has presented persistent challenges or has been a barrier to project work.

Additionally, OLCF should consider the following areas of emphasis:

Hardware Computing Resources	
File System	Explore the reasons for freezing, non-responsiveness, or other issues related to file system lag.
Access	
Authentication, Two-Factor, and SSH	Explore options for meeting users' interest in a less-cumbersome authentication process, in particular for workflows that are distributed across multiple systems.
Running Jobs	
Scheduling Policy	Explore ways that users with relatively small jobs, or jobs whose resource needs are quantified differently than the majority, can more readily compete with larger job users for OLCF resources.
Debug, Test, and Profiling	In addition to overall exploration of a scheduling policy that would accommodate smaller jobs, consider queues and other resources to assist users with debugging and testing. One user noted that the ability

	to test such jobs is particularly important while preparing to submit an OLCF proposal. Users are also interested in tools for profiling, benchmarking, and tracing performance.
Data Management	
Data Retention/Purge Procedures	Re-evaluate communication of the purge policy, the purge schedule, and documentation of how to plan for purges. Many users are very understanding of the need for the purge policy and made suggestions, like an email warning system, for improving user experiences even if the policy cannot be adjusted. If tips or lessons learned exist to help users maintain their code or custom libraries within the constraints of the purge policy, this should be documented and disseminated.
Software	
Libraries & Updates	Make more software and libraries available, maintain updates, and document them. The text of users' full comments indicate an array of software, tools, and libraries. Support and resources for containers, visualization tools (including remote visualization), and data access and analysis were frequently requested. Increase communication about updates. Document the software and libraries that are available across various systems, as well as instructions or user guides for updates or custom installs.

OLCF Evaluation

The following suggestions are offered with respect to the assessment of OLCF performance:

- Maintain the survey at its current length, which encouraged participation, streamlined analysis, and did not attract negative comments.
- Utilize the findings of the 2019 survey to make some minor adjustments to the 2020 survey.
- During annual survey refinement, highlight significant OLCF changes from the previous FY and planned/potential changes or rollouts in the upcoming FY, and ensure those areas are adequately probed by existing items in the survey.

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 longitudinal comparisons of user responses from 2006 through 2019. 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 preparation for the FY 2019 data collection year. The previous surveys, such as the most recent FY 2018 survey, were lengthy, and each year several respondents provided survey comments to this effect. For context, the FY 2018 survey included 181 available survey items. Because each respondent is only shown certain questions, the maximum number of these items answered by any one respondent was 150, but the average respondent answered 90 survey items. The average completion time, as measured and reported by the survey software interface, was 18.5 minutes.

The survey revisions were undertaken with the objectives of reducing the burden on participants and focusing the survey on the key insights that would provide useful information with OLCF. Ideally, the reduction in burden and redesign of the survey would also increase response rates and increase data quality.

Several approaches were taken to achieve the revised survey. First, the ORAU team calculated inter-item correlations to identify pairs of items that were redundant. Those with a correlation factor greater than 0.8 were reviewed to determine whether the items covered redundant content. When possible, such items were combined. For instance, items about *scheduled* vs. *unscheduled* outages/downtime were highly correlated, because users are unlikely to differentiate their satisfaction or dissatisfaction with outages based on whether they were scheduled or not. This pair of two items was combined into a single item about outages.

Second, the survey software was reconfigured to incorporate more “show/hides,” questions that are only displayed to relevant respondents based on their early responses. This technique was already in place to some extent, but the survey was reviewed and rearranged to optimize show/hides and minimize the number of items shown to each respondent. For instance, a question was added to the start of the survey to ask users which services or OLCF capabilities they had used in the past year. These responses were used by the survey software to show only relevant questions on the subsequent pages. As another example, comments boxes for dissatisfaction were only displayed if the respondent had indicated dissatisfaction with one or more elements within a given section of the survey.

Ultimately, all survey questions were reviewed with OLCF staff. Because of changes to OLCF and changes in the interests of the OLCF staff, additional questions were eliminated before the final survey was rolled out to OLCF users.

As a result of the revision process, the total number of possible items in the survey was reduced by approximately 56%, from 181 to 80. On average, respondents to the 2019 OLCF survey answered 38 items on the survey (compared to 90 items in FY 2018), and the maximum number of items shown to a user was 71. The average response time was 9.8 minutes, approximately half of the 18.5-minute response time in FY 2018. Additionally, because respondents were prevented from seeing and answering irrelevant items, the data cleaning and analysis process took less time.

Data Collection

The survey sampling frame was constituted by first collecting the names of individuals who had logged into an OLCF system between 1/1/2019 through 9/30/2019. OLCF staff and vendors, as well as individuals with invalid email addresses, were then removed from the list.

OLCF staff invited all OLCF users from this list to participate in the survey, which was hosted online beginning on October 2, 2019 and remained open for completion through December 2, 2019 (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.

Overall, this process resulted in a sampling frame with 1,251 OLCF users. A total of 578 users completed or partially completed the survey, resulting in a response rate of 46.2%. Figure 22, within Appendix B: Survey Administration Timeline, highlights the value of each reminder email in increasing the response rate.

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

OLCF (Overall)

OLCF Computing Resources

- Summit
- Titan
- Eos
- Rhea

OLCF Data Resources

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

OLCF 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 the scaled reply) with a resource or service (note: not all dissatisfied individuals supplied open-ended comments).

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

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);

ES The Summit Early Science (ES) program began in 2019 with the selection of a small number of projects that were given access to Summit from January 1, 2019 to June 30, 2019. The goals of the Early Science program were threefold. First, this was an opportunity to realize early scientific achievements on Summit. Second, the Early Science projects demonstrated the scalability and performance of applications ported to the Summit architecture. Third, the OLCF benefited from the "hardening" of both the hardware and software environment using production-ready codes at scale in important mission-relevant scientific computationally challenging projects; and

General Other projects fall into the General program. Staff projects are not reported.

Note that, in recent years, an "Other" category has combined General, Vendor, and other programs. In the FY 2019 user data, only the General program was represented, and Other has therefore been replaced throughout this report with this more specific category.

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

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

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 only three items are presented in a table.

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

[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.

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

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

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

Access
Accessibility and authentication
SSH
Inter-system access/transfer
Hardware Computing Resources
Performance
CPU resources
Architecture 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
File systems
Software
Software/tools/modules
Libraries and updates
Compilers
Visualization
Development tools
Debugging tools
Testing tools
Monitoring/profiling
User Support
Documentation
User guides

Tutorials
Training
Tech support and ticketing
Website
Communication
Example Additional Categories
Satisfaction
Miscellaneous
Survey suggestions
Administrative issues
Support for scientific research

Findings

Respondents

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

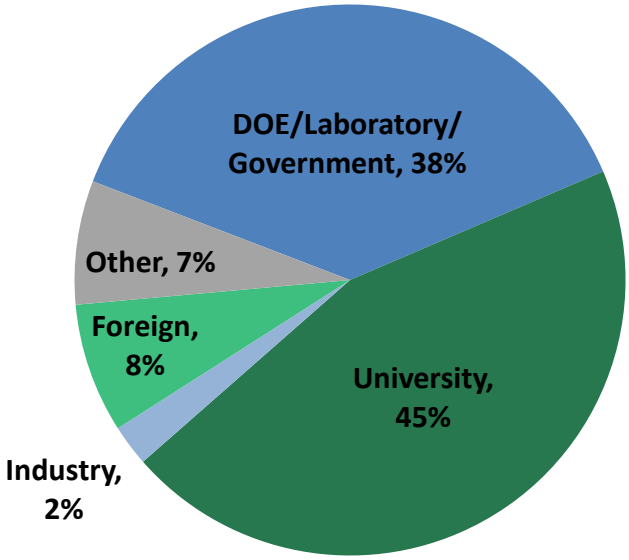

Figure 1. Respondent occupational affiliation (*N* = 578)

The distribution of OLCF users across **project allocations** is shown in Figure 2 and in detail in Table 2. The pool of survey respondents is generally representative of OLCF's ALCC, ECP, ES, and General users; INCITE users are somewhat overrepresented among respondents, and DD users are somewhat underrepresented. 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 62% of respondents reported a single project allocation (i.e., assignment to only INCITE, only DD, only ALCC, only ECP, or only ES).

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

	OLCF Users (*N* = 1251)		**Survey Respondents (*N* = 578)**	
	Percentage	***n***	**Percentage**	***n***
INCITE	37%	461	55%	320
DD	45%	560	33%	189
ALCC	22%	269	29%	170
ECP	27%	335	23%	135
ES	6%	81	9%	51
General	4%	50	5%	27

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

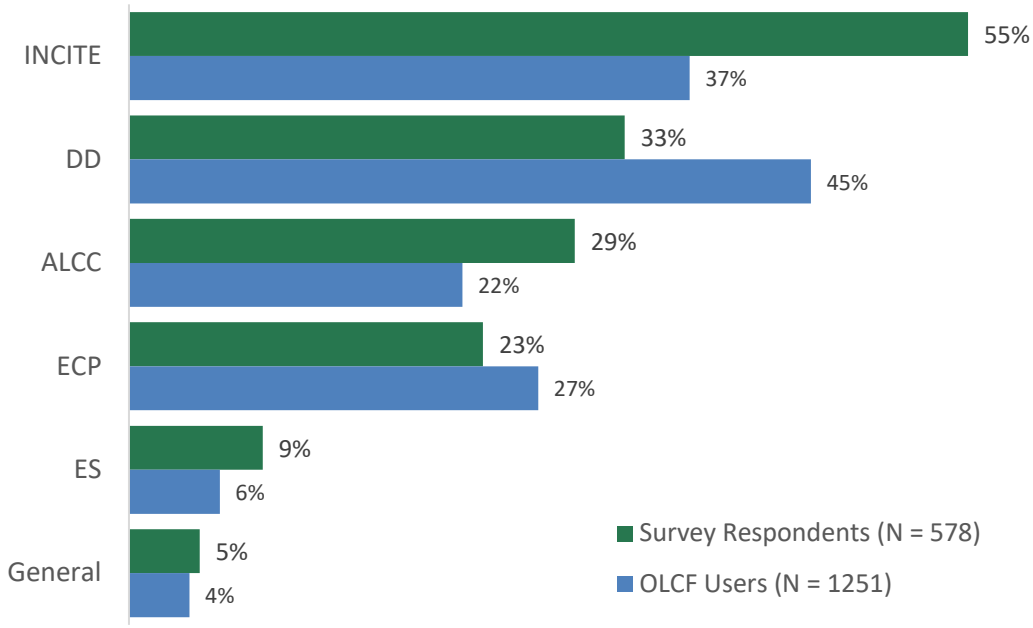

Figure 2. Project allocations for OLCF Users (*N* = 1251) and for Respondents (*N* = 578)

The proportions of OLCF users and of 2019 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.

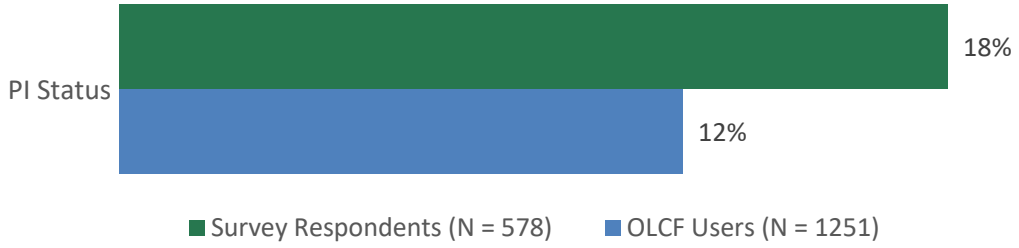

Figure 3. PI Status for OLCF Users (*N* = 1251) and for Respondents (*N* = 578)

Resource Utilization

Overall experience using the OLCF was approximately evenly split between users with more than 2 years of experience and those with 1-2 years or less than 1 year of OLCF experience. The largest proportion of respondents (nearly one-half) had used the OLCF for more than 2 years (Figure 4).

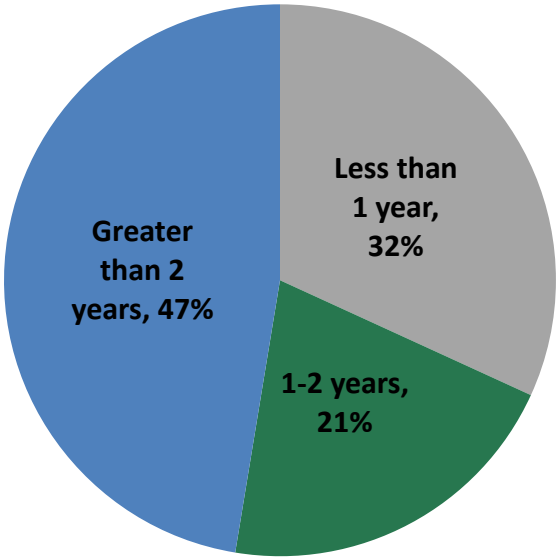

Figure 4. Experience using the OLCF (*N* = 578)

Respondents were asked to indicate **which OLCF HPC resources they utilized** during the 2019 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; Eos 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		ES		General		Total	
	n	**% Users**	***n***	**% Users**	***n***	**% Users**	***n***	**% Users**	***n***	**% Users**	***n***	**% Users**	***n***	**% Users**	***n***	**% Users**
Summit	79	77%	252	79%	138	73%	128	75%	126	93%	46	90%	26	96%	442	76%
Titan	57	56%	150	47%	96	51%	96	56%	54	40%	26	51%	11	41%	260	45%
Rhea	35	34%	87	27%	60	32%	38	22%	16	12%	13	25%	4	15%	137	24%
Eos	14	14%	25	8%	22	12%	26	15%	2	1%	0	0%	0	0%	59	10%
Data Transfer Nodes	35	34%	81	25%	63	33%	33	19%	25	19%	13	25%	4	15%	145	25%
HPSS	35	34%	82	26%	51	27%	32	19%	24	18%	10	20%	2	7%	124	21%
Lustre/Spider	43	42%	103	32%	74	39%	50	29%	40	30%	13	25%	9	33%	180	31%
Alpine GPFS	34	33%	119	37%	74	39%	59	35%	56	41%	20	39%	9	33%	192	33%
OLCF Website	61	60%	171	53%	118	62%	92	54%	83	61%	33	65%	15	56%	308	53%
I have not used any of the listed resources	3	3%	13	4%	3	2%	6	4%	4	3%	1	2%	0	0%	21	4%

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

Overall Satisfaction

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

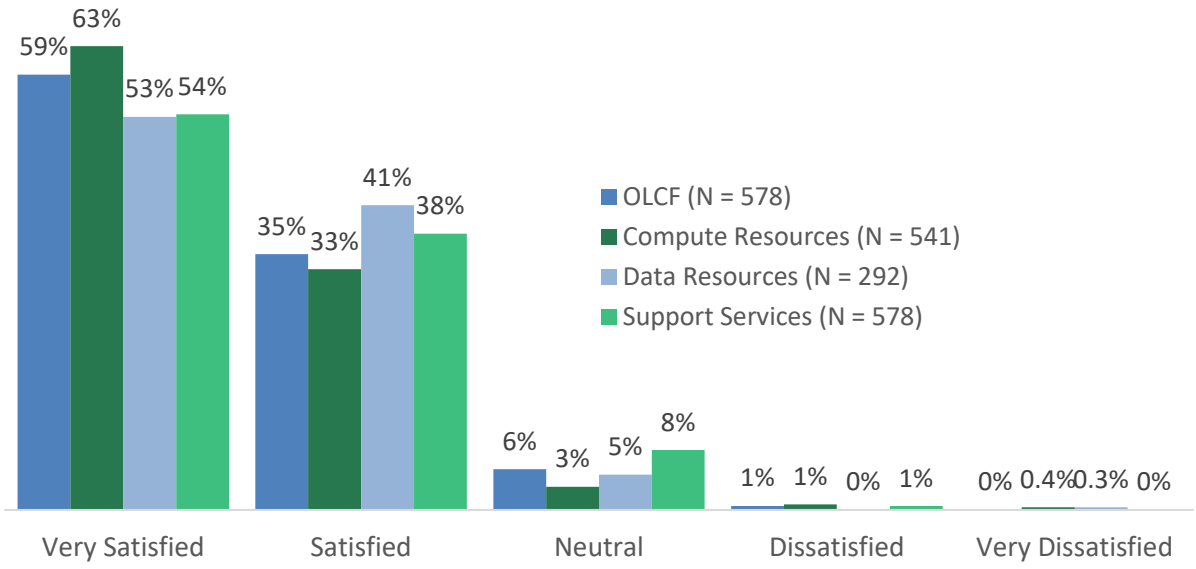

Figure 5. "Overall" Satisfaction with OLCF and its major resources/services (maximum *N* = 578)

Table 4 summarizes descriptive statistics for these overall satisfaction ratings for all respondents and broken down by PI status, while Table 5 and Table 6 report satisfaction statistics across project allocations. The tables also include ratings of **specific compute resources** (i.e., Titan, Eos, Rhea, and Summit), data resources (i.e., Data Transfer Nodes, HPSS, Lustre/Spider, and Alpine GPFS), 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 13 items and the full set of respondents, the tables show that:

- *%Sat* ranged from 90% to 98%,
- Means ranged from 4.3 to 4.6, and
- SDs ranged from 0.57 to 0.73.

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

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
OLCF	**102**	**4.5**	**0.61**	**94%**	**476**	**4.5**	**0.63**	**94%**	**578**	**4.5**	**0.63**	**94%**
Compute Resources	**94**	**4.6**	**0.64**	**96%**	**447**	**4.6**	**0.63**	**96%**	**541**	**4.6**	**0.63**	**96%**
Titan	56	4.5	0.66	95%	200	4.4	0.67	91%	256	4.4	0.67	91%
Eos	14	4.7	0.61	93%	44	4.5	0.70	89%	58	4.5	0.68	90%
Rhea	34	4.6	0.61	94%	102	4.5	0.56	97%	136	4.5	0.57	96%
Summit	79	4.6	0.60	94%	355	4.5	0.62	95%	434	4.5	0.62	95%
Data Resources	**61**	**4.5**	**0.72**	**95%**	**231**	**4.5**	**0.60**	**95%**	**292**	**4.5**	**0.62**	**95%**
Data Transfer Nodes	35	4.2	0.97	83%	110	4.5	0.63	97%	145	4.4	0.73	94%
HPSS	35	4.5	0.78	97%	85	4.6	0.54	98%	120	4.5	0.62	98%
Lustre/Spider Scratch Filesystem	43	4.4	0.69	88%	133	4.3	0.64	92%	176	4.3	0.65	91%
Alpine GPFS Scratch Filesystem	34	4.5	0.56	97%	154	4.4	0.56	97%	188	4.4	0.56	97%
Support Services	**102**	**4.5**	**0.67**	**94%**	**476**	**4.4**	**0.66**	**91%**	**578**	**4.4**	**0.66**	**91%**
Support received**	97	4.6	0.71	93%	405	4.6	0.64	94%	502	4.6	0.66	93%
OLCF website	61	4.4	0.66	93%	245	4.4	0.65	93%	306	4.4	0.65	93%
Communications	101	4.7	0.51	98%	453	4.5	0.61	94%	554	4.6	0.59	95%
Support and training documentation	88	4.4	0.71	92%	359	4.5	0.66	91%	447	4.4	0.67	91%
Problem resolution	91	4.5	0.74	92%	353	4.5	0.63	94%	444	4.5	0.65	93%
Min	14	4.2	0.51	83%	44	4.3	0.54	89%	58	4.3	0.57	90%
Max	102	4.7	0.97	98%	476	4.6	0.70	98%	578	4.6	0.73	98%

**Support received lists out: user assistance, accounts, INCITE Liaisons, Advanced Data/Workflow Liaisons.

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

	INCITE				DD				ALCC			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
OLCF	**320**	**4.5**	**0.66**	**93%**	**189**	**4.6**	**0.59**	**95%**	**170**	**4.5**	**0.63**	**94%**
Compute Resources	**300**	**4.5**	**0.69**	**94%**	**177**	**4.6**	**0.58**	**97%**	**161**	**4.6**	**0.55**	**97%**
Titan	146	4.4	0.71	90%	96	4.3	0.68	91%	95	4.4	0.64	92%
Eos	24	4.4	0.78	83%	22	4.5	0.74	86%	25	4.6	0.65	92%
Rhea	87	4.4	0.59	95%	59	4.6	0.53	98%	38	4.6	0.59	95%
Summit	246	4.5	0.61	96%	137	4.5	0.63	94%	124	4.4	0.66	94%
Data Resources	**173**	**4.5**	**0.60**	**94%**	**112**	**4.4**	**0.66**	**96%**	**86**	**4.4**	**0.58**	**95%**
Data Transfer Nodes	81	4.5	0.59	95%	63	4.4	0.78	94%	33	4.3	0.77	97%
HPSS	80	4.6	0.52	99%	50	4.5	0.71	98%	32	4.5	0.57	97%
Lustre/Spider Scratch Filesystem	101	4.3	0.59	92%	72	4.3	0.71	86%	48	4.2	0.73	85%
Alpine GPFS Scratch Filesystem	116	4.5	0.55	97%	72	4.4	0.57	96%	56	4.3	0.54	96%
Support Services	**320**	**4.4**	**0.67**	**91%**	**189**	**4.5**	**0.65**	**94%**	**170**	**4.5**	**0.62**	**94%**
Support received**	272	4.6	0.65	94%	175	4.6	0.67	94%	150	4.6	0.62	93%
OLCF website	170	4.3	0.67	92%	118	4.4	0.60	96%	92	4.3	0.62	92%
Communications	303	4.6	0.62	94%	188	4.6	0.56	97%	165	4.5	0.61	94%
Support and training documentation	239	4.4	0.68	90%	157	4.5	0.68	92%	138	4.4	0.65	93%
Problem resolution	238	4.6	0.60	94%	156	4.5	0.69	94%	136	4.5	0.68	93%
Min	24	4.3	0.52	83%	22	4.3	0.53	86%	25	4.2	0.55	85%
Max	320	4.6	0.78	99%	189	4.6	0.78	98%	170	4.6	0.77	97%

**Support received lists out: user assistance, accounts, INCITE Liaisons, Advanced Data/Workflow Liaisons.

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

	ECP				ES*				General*			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
OLCF	**135**	**4.5**	**0.62**	**96%**	**51**	**4.5**	**0.61**	**94%**	**27**	**4.6**	**0.49**	**100%**
Compute Resources	**128**	**4.6**	**0.64**	**98%**	**48**	**4.6**	**0.54**	**98%**	**27**	**4.7**	**0.48**	**100%**
Titan	53	4.3	0.72	85%	26	4.3	0.67	88%	11	4.4	0.67	91%
Eos	2	4.5	0.71	100%	0	--	--	--	0	--	--	--
Rhea	16	4.5	0.63	94%	13	4.5	0.66	92%	4	4.5	0.58	100%
Summit	124	4.5	0.55	98%	46	4.5	0.59	96%	26	4.5	0.51	100%
Data Resources	**75**	**4.5**	**0.60**	**95%**	**31**	**4.4**	**0.50**	**100%**	**13**	**4.6**	**0.51**	**100%**
Data Transfer Nodes	25	4.2	0.82	84%	13	4.4	0.77	85%	4	4.0	0.00	100%
HPSS	22	4.7	0.48	100%	10	4.7	0.48	100%	2	5.0	0.00	100%
Lustre/Spider Scratch Filesystem	38	4.3	0.73	84%	13	4.1	0.76	77%	9	4.1	0.60	89%
Alpine GPFS Scratch Filesystem	55	4.4	0.57	96%	20	4.5	0.60	95%	9	4.3	0.50	100%
Support Services	**135**	**4.4**	**0.63**	**93%**	**51**	**4.5**	**0.70**	**92%**	**27**	**4.5**	**0.58**	**96%**
Support received**	118	4.6	0.60	96%	45	4.4	0.81	84%	23	4.7	0.56	96%
OLCF website	82	4.4	0.56	96%	33	4.3	0.76	82%	15	4.5	0.52	100%
Communications	129	4.6	0.60	95%	50	4.5	0.68	90%	26	4.7	0.55	96%
Support and training documentation	105	4.5	0.65	93%	43	4.5	0.63	93%	20	4.6	0.51	100%
Problem resolution	105	4.6	0.65	93%	42	4.5	0.59	95%	19	4.6	0.50	100%
Min	2	4.2	0.48	84%	10*	4.1	0.48	77%	2*	4.0	0.00	89%
Max	135	4.7	0.82	100%	51	4.7	0.81	100%	27	5.0	0.67	100%

*27 respondents are allocated to "General" projects and 51 respondents to "ES" projects; as a result, some questions received too few responses for SD or *%Sat* to be meaningful. Items receiving 0 responses from a subgroup are excluded from min and max calculations. **Support received lists out: user assistance, accounts, INCITE Liaisons, Advanced Data/Workflow Liaisons.

Only 10 respondents reported reasons for **dissatisfaction with the OLCF overall or with its major resources/services.** Individuals cited problems with *Summit* (*n* = 4), *performance* (*n* = 3), *support, training, and documentation (n* = 1), *usefulness/user-friendliness* (*n* = 1), and *miscellaneous* (*n* = 1). Illustrative examples include:

Summit	*"Working on OLCF's Summit was a mixed-feelings experience. Having access to Summit's hardware is thrilling and exciting and the user documentation is satisfactory. In the end however, our largest obstacle and ultimate stumbling block was Summit's software stack: to run our software (cp2k, https://github.com/cp2k/cp2k, http://cp2k.org) on Summit, we need an optimized & multithreaded BLAS compatible with gcc, which is also known to run without errors on Summit. This requirement is a common dependency of so many HPC programs - you would expect a leading supercomputing facility to provide such a library that one could simply ""module load"".* *We were able to run our software correctly with netlib's BLAS & lapack, but with unsatisfactory performance, unsurprisingly.* *Using ESSL, on the contrary, wasn't an option in our case since cp2k (in particular, its C++ files) cannot be compiled with IBM's XL compiler, and XL and gcc's OpenMP runtimes are incompatible.*
Performance	*"Our experience was that carrying out sustained 'capability' computing runs on Titan was very difficult due to frequent hardware failures that were not easily managed within existing tools. These issues, combined with limitations on memory and communication speed significantly increased computational and wall-clock time and cost for our work. I am happy to report that my colleagues indicate significantly better experience with Summit."*
Support, Training, and Documentation	*"I find jsrun cumbersome and poorly explained, and in particular its interaction with cudaSetDevice and cuda device numbering. This may be IBM's fault. But ORNL could explain more with code examples, and how this differs from other systems where mpi-rank N should set the device. How this interacts with which socket a rank & device is on etc... should be documented so code can rely on it."*

	“User access policies are obnoxious. Either enable ssh session sharing or find a way to let me authenticate via MFA once per day instead of once per login."
Usefulness and User-Friendliness	*The OLCF support ticket system lacks end-user visibility, which is a MAJOR problem for tracking of ongoing problems, especially those affecting entire research groups. This is a terrible way to do user support. The support ticket system should be directly visible to users submitting tickets, their project members and anyone they explicitly CC."*

Finally, respondents described what they perceived to be **“the best qualities of OLCF.”** Thematic analysis of user responses identified *computing power/hardware/HPC resources/performance* (46%) and *user support/staff* (37%) as the most valued qualities of the OLCF (Table 7; see Appendix C: Best Qualities of the OLCF for all responses by category; *N* = 513). Many illustrative examples praised multiple elements of OLCF:

"Comprehensive computational user facility offering resources across scales that serve jobs across the spectrum. The software ecosystem is robust and the information dissemination is spot on."

"Excellent software stack and stability. I've seen several large-scale HPC installations (and even more small-scale installations). However, your installation clearly outperforms all of them!"

"Extremely scientist-friendly facility, which reflects on the quality of the research and publication output. The staff works hands-on even with the young and junior-most researchers in the INCITE teams to get them up and running. Requesting for priority boosts is always handled with a lot of care, keeping the researchers publication and proposal interests in mind. Discretionary acceptance of additional allocation requests often lead to new discoveries - a direct evidence of the foresight in the OLCF staff!"

"In my experience, the following are OLCF's best qualities: 1) OLCF's has the most powerful supercomputer in the world which is crucial for the large scale electronic-structure studies we are conducting, 2) OLCF's support team's quick response and quality of issue resolution is top notch, and finally 3) comprehensive user documentation which includes critical things like known issues. OLCF's GPU Hackathon experience was also extremely useful for GPU porting of our code."

"OLCF has a history of providing leadership class systems with state of the art hardware and software technologies, while also making the systems relatively easy to use. OLCF generally provides better job scheduling turnaround than other centers, and historically does a fantastic job with scheduling small debugging and software development oriented jobs of the sort that I'm

typically involved with. I often use OLCF as a positive example for comparison when giving advice to other computing centers about how they might improve their services."

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 individuals that mentioned **only** computing performance as the best quality (removing 95 responses). The relative frequency of comments reported by this group (*N* = 418), excluding references to computing power/performance is shown in the last column of Table 7. *User support/staff* is prominent as the perceived best OLCF quality when the responses are examined in this way, but there is significant spread across other categories and variety in responses.

For example:

User Support/Staff	*"Help -- the folks at OLCF are the most helpful of staff I have worked with. Turn around with questions, setting up libraries and requirements and providing reservations, all these are the best aspects of OLCF."* *"I have had good experiences with user support and found the OLCF to respond very quickly."* *"In past experience, the quick response and technical knowledge of the support staff has been the greatest asset."*
System Documentation & Website Information	*"90% of my questions are answered on the easily accessed user guides."* *"Coming in as an inexperienced user, I found the documentation/online support to be very useful."*
Summit	*"Early access to Summit was managed well. The team supporting our early use was great."* *"I have observed that Summit's reliability has improved significantly from the previous year, making it an excellent system for conducting computational science research."*

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

	All Responses (*N* = 513)	**Responses Excluding Computing Performance (*N* = 418)**
Computing power/hardware/HPC resources	46%	n/a
User support/staff	37%	46%
System documentation & website information	11%	13%
Summit	11%	13%
Stability/reliability	7%	8%
Tools (software, libraries, visualization, & analysis)	7%	8%
Queue time/turnaround time	6%	7%
Availability/uptime	5%	7%
GPU resources	4%	6%
Ease of use	4%	6%
Resource management/infrastructure/ maintenance	4%	5%
Supports scientific research/demanding problems	4%	5%
Overall satisfaction	4%	5%
Communication	4%	4%
Training/tutorials	3%	3%
Customer-focused	2%	3%
Accessibility	2%	3%
Data storage/disk space	2%	2%
Scheduling	1%	1%
Data transfer	1%	1%
Shared filesystem/cross-system use cases	1%	1%
Miscellaneous/Other	4%	6%

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

Compute and Data Resources

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

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

Table 8 reports satisfaction for these features by PI status and overall, and Table 9 reports ratings by project allocation. The highest satisfaction ratings (all respondents) were for *notice given prior to scheduled maintenance* (94% satisfied) and bandwidth offered by the OLCF (94% satisfied). The lowest overall mean rating was for *I/O performance* (89% satisfied).

Of the 10 respondents that reported reasons for **dissatisfaction with one or more aspects of the OLCF compute and data resources,** the majority of complaints had to do with *performance/reliability* (*n* = 4), *I/O or data transfer* (*n* = 3), and *bandwidth or lag* (*n* = 3).

Performance/Reliability	*"Constantly freezing of the system..."* *"The FS is slow and unpredictable."*
I/O or Data Transfer	*"I/O is always a big bottleneck for us. I'm not sure if this is an OLCF problem or a code/implementation problem, but I/O limitations definitely limit us."*
Bandwidth or Lag	*"Bandwidth offsite could be substantially improved."*

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

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Sufficient notice of scheduled downtimes	94	4.6	0.73	94%	428	4.5	0.67	94%	522	4.5	0.68	94%
Sufficient project disk space	92	4.5	0.73	93%	431	4.4	0.72	90%	523	4.5	0.72	91%
Bandwidth offered by the OLCF	91	4.4	0.82	95%	426	4.5	0.66	94%	517	4.5	0.69	94%
I/O performance	86	4.4	0.83	90%	413	4.4	0.74	89%	499	4.4	0.75	89%

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

	INCITE				DD				ALCC			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Sufficient notice of scheduled downtimes	287	4.5	0.72	94%	178	4.6	0.58	97%	154	4.5	0.66	92%
Sufficient project disk space	289	4.4	0.76	90%	176	4.5	0.65	94%	157	4.4	0.71	89%
Bandwidth offered by the OLCF	288	4.4	0.74	91%	176	4.5	0.62	97%	153	4.5	0.64	95%
I/O performance	280	4.4	0.77	89%	166	4.4	0.73	90%	147	4.3	0.73	89%
	ECP				**ES**				**General**			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Sufficient notice of scheduled downtimes	125	4.6	0.60	95%	49	4.5	0.84	88%	24	4.7	0.46	100%
Sufficient project disk space	123	4.5	0.67	92%	49	4.3	0.88	88%	26	4.6	0.58	96%
Bandwidth offered by the OLCF	119	4.6	0.62	97%	47	4.4	0.77	94%	24	4.7	0.46	100%
I/O performance	114	4.3	0.80	87%	45	4.3	0.95	82%	22	4.4	0.50	100%

In addition, respondents were asked to indicate their opinion regarding the **performance of computing and data resources compared to the previous year.** Overall, 47.7% reported *improvements*, just 0.5% perceived *decreases in performance*, and 51.8% reported *no change* (Figure 6). Some differences in these perceptions were observed across years of using the OLCF. Less experienced users (those with only 1-2 years' experience) were slightly less likely to report seeing an improvement over the last year.

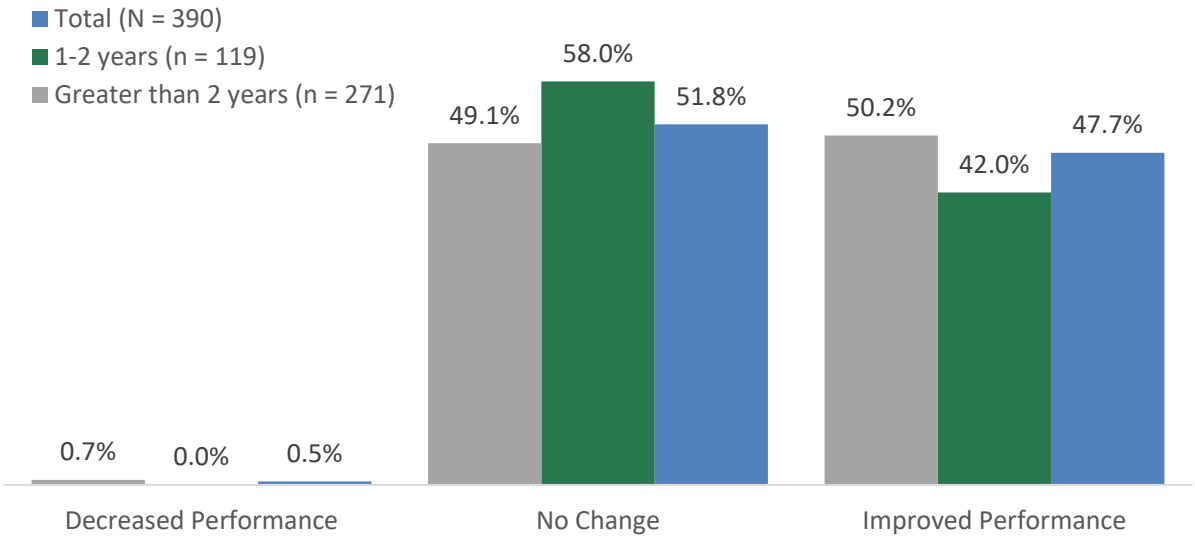

Figure 6. Perceived changes from FY 2018 computing/data resources performance by years using OLCF (*N* = 390)

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

> *"Summit does not seem to be as reliable regarding MPI communications as Titan was. We are working with our CAAR liaison to try to understand this better."*

> *"Unknown hardware failures have prevented us from running on the full system, and performance at scale has been limited by individual GPU or network issues."*

Compute Resources

Summit

Summit was used by 76% of respondents during the 2019 calendar year (*N* = 442). Summit users were asked to provide satisfaction ratings for multiple aspects of the system, and descriptive statistics of these ratings are shown in Table 12, which also reports satisfaction by PI status. 95% of all respondents were either *satisfied* or *very satisfied* with the system. Table 13 summarizes these satisfaction statistics by project allocation.

The *system availability and scheduling turnaround* were the highest rated specific aspects of Summit, and lowest rated aspects were *availability of tools* and *programming environment.*

There were 16 Summit users who reported at least one reason for **dissatisfaction with Summit**. Over two-thirds of these users (*n* = 11) were unhappy with *software, libraries, and compatibilities* that had an impact on the work they could conduct on the system. For example:

> *"Long-standing simulation tools remain missing on Summit and porting of tools and software from developers seems unlikely given the expected short-lifetime of the machine compared to the upcoming Frontier."*

> *"The IBM software stack needs serious improvements. It is a pity to have probably some of the best hardware around but a poor software stack."*

The next largest group of dissatisfied users (*n* = 3) were unhappy with the *development and debugging capabilities/time,* particularly because Summit was a new system. For example:

> *"Queues are quite busy; makes testing hard. Low wall-clock limit (2 hrs) for 1-node jobs makes debugging and optimization hard."*

Table 10. *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	75	4.5	0.62	93%	340	4.5	0.66	92%	415	4.5	0.65	92%
System availability	78	4.6	0.65	91%	349	4.5	0.64	93%	427	4.5	0.64	93%
Availability of tools	73	4.4	0.79	86%	348	4.3	0.73	86%	421	4.3	0.74	86%
Availability of libraries	74	4.4	0.83	85%	348	4.3	0.70	88%	422	4.3	0.73	88%
Programming environment	75	4.3	0.86	83%	346	4.3	0.79	88%	421	4.3	0.80	87%
Overall satisfaction with Summit	79	4.6	0.60	94%	355	4.5	0.62	95%	434	4.5	0.62	95%

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

	INCITE				DD				ALCC			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Scheduling turnaround	234	4.5	0.66	91%	131	4.5	0.61	95%	119	4.4	0.63	94%
System availability	241	4.6	0.63	93%	135	4.5	0.62	95%	123	4.5	0.62	93%
Availability of tools	239	4.4	0.76	87%	132	4.2	0.80	83%	119	4.2	0.70	86%
Availability of libraries	239	4.4	0.70	91%	132	4.2	0.84	81%	120	4.3	0.63	92%
Programming environment	237	4.3	0.80	89%	134	4.2	0.91	81%	122	4.3	0.81	86%
Overall satisfaction with Summit	246	4.5	0.61	96%	137	4.5	0.63	94%	124	4.4	0.66	94%
	ECP				**ES**				**General**			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Scheduling turnaround	122	4.5	0.63	94%	45	4.4	0.78	82%	26	4.4	0.75	92%
System availability	124	4.5	0.59	97%	45	4.5	0.69	89%	26	4.4	0.80	88%
Availability of tools	122	4.3	0.69	88%	43	4.1	0.82	77%	26	4.0	0.77	69%
Availability of libraries	123	4.4	0.62	93%	42	4.2	0.76	81%	25	4.1	0.78	84%
Programming environment	123	4.3	0.72	91%	43	4.1	0.83	77%	26	4.0	0.87	73%
Overall satisfaction with Summit	124	4.5	0.55	98%	46	4.5	0.59	96%	26	4.5	0.51	100%

Titan

Titan was used by 45% of respondents during the 2019 calendar year (*N* = 260). Titan users were asked to provide overall satisfaction ratings for the system, and descriptive statistics of these ratings are shown in Table 12, which also reports satisfaction by PI status. 91% of all respondents were either *satisfied* or *very satisfied* with the system. Table 13 summarizes these satisfaction statistics by project allocation.

Only two users reported reasons for **dissatisfaction with Titan**:

> *"As described above, our experience was that carrying out sustained "capability" computing runs on Titan was very difficult due to frequent hardware failures that were not easily managed within existing tools. These issues, combined with limitations on memory and communication speed significantly increased computational and wall-clock time and cost for our work. I am happy to report that my colleagues indicate significantly better experience with Summit."*

> *"Lots of node failures."*

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

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Overall satisfaction with Titan	56	4.5	0.66	95%	200	4.4	0.67	91%	256	4.4	0.67	91%

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

	INCITE				DD				ALCC			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Overall satisfaction with Titan	146	4.4	0.71	90%	96	4.3	0.68	91%	95	4.4	0.64	92%
	ECP				**ES**				**General**			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Overall satisfaction with Titan	53	4.3	0.72	85%	26	4.3	0.67	88%	11	4.4	0.67	91%

Eos

Eos was used by 10% of respondents during the 2019 calendar year (*N* = 59). Eos users were asked to provide overall satisfaction ratings for the system, and descriptive statistics of these ratings are shown in Table 14, which also reports satisfaction statistics by PI status. Most (90%) respondents were either *satisfied* or *very satisfied* with the system. Table 15 summarizes these satisfaction statistics by project allocation.

No Eos users indicated dissatisfaction or provided **reasons for dissatisfaction**.

Rhea

Rhea was used by 24% of respondents during the 2019 calendar year (*N* = 137). Rhea users were asked to provide satisfaction ratings for multiple aspects of the system, and descriptive statistics of these ratings are shown in Table 16, which also reports satisfaction statistics by PI status. 96% of respondents were either *satisfied* or *very satisfied* with the system. Table 17 summarizes these satisfaction statistics by project allocation.

The highest rated aspect of Rhea was the *system availability,* and the lowest rated features was the *availability of libraries.*

Only three Rhea users expressed **reasons for dissatisfaction**, all related to the environment:

> *"Having the cray environment and wrappers makes compiling our data analysis code so much easier. And the memory limit that is set automatically caused a lot of run time issues that required odd workarounds."*
>
> *"I had some problems using higher version of GCC, and ended up using the 4.x version."*
>
> *"There is no clang support on Rhea."*

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

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Overall satisfaction with Eos	14	4.7	0.61	93%	44	4.5	0.70	89%	58	4.5	0.68	90%

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

	INCITE				DD				ALCC			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Overall satisfaction with Eos	24	4.4	0.78	83%	22	4.5	0.74	86%	25	4.6	0.65	92%
	ECP				**ES**				**General**			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Overall satisfaction with Eos	2	4.5	0.71	100%	0	--	--	--	0	--	--	--

Table 16. *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	34	4.5	0.66	91%	96	4.6	0.63	93%	130	4.5	0.64	92%
System availability	35	4.7	0.59	94%	100	4.7	0.59	94%	135	4.7	0.59	94%
Availability of tools	32	4.3	0.77	81%	99	4.4	0.67	90%	131	4.4	0.70	88%
Availability of libraries	33	4.4	0.75	85%	99	4.4	0.72	88%	132	4.4	0.72	87%
Programming environment	34	4.4	0.82	85%	97	4.4	0.75	89%	131	4.4	0.76	88%
Overall satisfaction with Rhea	34	4.6	0.61	94%	102	4.5	0.56	97%	136	4.5	0.57	96%

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

	INCITE				DD				ALCC			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Scheduling turnaround	83	4.5	0.67	90%	56	4.6	0.55	96%	37	4.6	0.65	92%
System availability	86	4.6	0.62	93%	59	4.8	0.49	97%	37	4.7	0.52	97%
Availability of tools	84	4.4	0.71	87%	56	4.4	0.70	88%	36	4.4	0.69	89%
Availability of libraries	85	4.3	0.76	85%	56	4.4	0.73	89%	36	4.4	0.65	92%
Programming environment	82	4.3	0.78	85%	58	4.4	0.84	88%	37	4.5	0.65	92%
Overall satisfaction with Rhea	87	4.4	0.59	95%	59	4.6	0.53	98%	38	4.6	0.59	95%
	ECP				**ES**				**General**			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Scheduling turnaround	16	4.7	0.48	100%	12	4.6	0.51	100%	4	4.8	0.50	100%
System availability	16	4.9	0.34	100%	13	4.8	0.44	100%	4	5.0	0.00	100%
Availability of tools	16	4.2	0.75	81%	12	4.5	0.67	92%	4	4.0	0.82	75%
Availability of libraries	16	4.3	0.60	94%	11	4.5	0.69	91%	4	3.8	0.50	75%
Programming environment	16	4.1	1.02	81%	11	4.2	0.98	82%	4	4.0	0.82	75%
Overall satisfaction with Rhea	16	4.5	0.63	94%	13	4.5	0.66	92%	4	4.5	0.58	100%

Data Resources

Data Transfer Nodes

DTNs were used by 25% of respondents during the 2019 calendar year (*N* = 145), and 94% were either *satisfied* or *very satisfied* with the DTNs (Table 18 and Table 19). Only two users provided comments explaining their **reasons for dissatisfaction** with the DTNs:

> *"It does not have public IP address which makes it very hard to use user specific data transfer tools."*

> *"My dissatisfaction is related to getting data off via Globus. It's convenient to have a DTN for HPSS, but the bandwidth and connectivity has not been great."*

HPSS

HPSS was used by 21% of respondents during the 2019 calendar year (*N* = 124). HPSS users were asked to provide satisfaction ratings for multiple aspects of the system, and descriptive statistics of these ratings are shown in Table 20, which also reports satisfaction statistics by PI status. 98% of respondents were overall either *satisfied* or *very satisfied* with the system. The highest rated items were *reliability (data integrity)* and the *ability to store/retrieve files.* The lowest rated aspect was the *frequency of outages.* Table 21 summarizes these satisfaction statistics by project allocation.

Of the two respondents that reported **reasons for dissatisfaction** with HPSS, both mentioned outages:

> *"Transferring data from HPSS offsite via Globus has been a terrible experience -- frequent re-authorizations and outages have prevented me from moving ~100 TB data sets offsite. I'm actually worried they will be lost because we can't get them off HPSS."*

> *"Very frequent downtimes, always accompanied by email blasts!"*

Table 18. *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	35	4.2	0.97	83%	110	4.5	0.63	97%	145	4.4	0.73	94%

Table 19. *Satisfaction Ratings of Data Transfer Nodes by Project Allocation*

	INCITE				DD				ALCC			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Overall satisfaction with Data Transfer Nodes	81	4.5	0.59	95%	63	4.4	0.78	94%	33	4.3	0.77	97%
	ECP				**ES**				**General**			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Overall satisfaction with Data Transfer Nodes	25	4.2	0.82	84%	13	4.4	0.77	85%	4	4.0	0.00	100%

Table 20. *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	30	4.4	0.72	93%	73	4.3	0.63	92%	103	4.3	0.65	92%
Globus interface	23	4.5	0.95	91%	65	4.5	0.69	89%	88	4.5	0.76	90%
Ability to store/retrieve files	35	4.5	0.66	97%	84	4.5	0.59	95%	119	4.5	0.61	96%
Reliability (data integrity)	35	4.5	0.66	97%	80	4.6	0.59	95%	115	4.6	0.61	96%
Time to store/retrieve files	35	4.3	0.79	94%	84	4.4	0.65	90%	119	4.3	0.69	92%
Frequency of outages	33	4.3	0.85	91%	82	4.2	0.77	82%	115	4.3	0.79	84%
Overall satisfaction with HPSS	35	4.5	0.78	97%	85	4.6	0.54	98%	120	4.5	0.62	98%

Table 21. *Satisfaction Ratings of HPSS by Project Allocation*

	INCITE				DD				ALCC			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
hsi/htar interface	75	4.4	0.61	93%	38	4.4	0.68	95%	23	4.4	0.59	96%
Globus interface	55	4.5	0.66	91%	37	4.5	0.87	89%	22	4.7	0.57	95%
Ability to store/retrieve files	80	4.6	0.55	98%	49	4.6	0.64	96%	31	4.5	0.62	94%
Reliability (data integrity)	78	4.6	0.55	97%	47	4.6	0.68	94%	30	4.6	0.63	93%
Time to store/retrieve files	79	4.4	0.61	94%	50	4.3	0.77	92%	32	4.4	0.71	88%
Frequency of outages	76	4.4	0.67	89%	49	4.2	0.92	80%	31	4.1	0.88	74%
Overall satisfaction with HPSS	80	4.6	0.52	99%	50	4.5	0.71	98%	32	4.5	0.57	97%
	ECP				**ES**				**General**			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
hsi/htar interface	18	4.5	0.51	100%	9	4.4	0.53	100%	2	5.0	0.00	100%
Globus interface	17	4.5	0.72	88%	8	4.5	0.76	88%	2	5.0	0.00	100%
Ability to store/retrieve files	22	4.6	0.59	95%	10	4.5	0.71	90%	2	5.0	0.00	100%
Reliability (data integrity)	22	4.6	0.59	95%	9	4.7	0.71	89%	2	5.0	0.00	100%
Time to store/retrieve files	22	4.6	0.50	100%	9	4.6	0.53	100%	2	4.5	0.71	100%
Frequency of outages	21	4.4	0.60	95%	9	4.4	0.73	89%	2	4.5	0.71	100%
Overall satisfaction with HPSS	22	4.7	0.48	100%	10	4.7	0.48	100%	2	5.0	0.00	100%

Lustre/Spider Scratch Filesystem

Lustre/Spider was used by 31% of respondents during the 2019 calendar year (*N* = 180). Lustre/Spider users were asked to provide overall satisfaction ratings for the system, and descriptive statistics of these ratings are shown in Table 22, which also reports satisfaction statistics by PI status. 91% of respondents were overall either *satisfied* or *very satisfied* with the system. Table 23 summarizes these satisfaction statistics by project allocation.

Only one user **indicated dissatisfaction** with at least one aspect of the Lustre/Spider Scratch Filesystem:

> *"Purge frequency was high."*

Alpine GPFS Scratch Filesystem

Alpine GPFS Scratch Filesystem was used by 33% of respondents during the 2019 calendar year (*N* = 192). 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 22, which also reports satisfaction statistics by PI status. 97% of respondents were overall either *satisfied* or *very satisfied* with the system. The *reliability (data integrity)* and *size* were the highest rated Alpine GPFS features, and the lowest rated feature was the *frequency of outages.* Table 23 summarizes these satisfaction statistics by project allocation.

There were only three users who **indicated dissatisfaction** with at least one aspect of the Alpine GPFS Scratch Filesystem, and all comments related to outages or periods of instability:

> *"Outages have been more frequent than desirable, and communication of these outages has not been as rapid as desired (more than once I learned of an outage from experience--e.g. not being able to see a directory, only to learn of the outage by searching for it on OLCF's website). However, Alpine is still relatively new, and this already seems to be happening less frequently."*

> *"Presumably when jobs are doing large I/O, gpfs becomes unusable for periods of 5-15 minutes at a time. Not as bad as Lustre on Titan, but comparable. This is a major quality of life issue for users, particularly for external collaborators who are confused when their session is unresponsive for large periods of time."*

> *"There were a lot of unscheduled outages."*

Table 22. *Satisfaction Ratings of Lustre/Spider Scratch Filesystem by PI Status and Overall*

	PI Status				Non-PI Status				Total			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Overall satisfaction with Lustre/Spider filesystem	43	4.4	0.69	88%	133	4.3	0.64	92%	176	4.3	0.65	91%

Table 23. *Satisfaction Ratings of Lustre/Spider Scratch Filesystem by Project Allocation*

	INCITE				DD				ALCC			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Overall satisfaction with Lustre/Spider filesystem	101	4.3	0.59	92%	72	4.3	0.71	86%	48	4.2	0.73	85%
	ECP				**ES**				**General**			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Overall satisfaction with Lustre/Spider filesystem	38	4.3	0.73	84%	13	4.1	0.76	77%	9	4.1	0.60	89%

Table 24. *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	34	4.4	0.66	91%	154	4.5	0.59	95%	188	4.5	0.6	95%
I/O bandwidth	34	4.5	0.62	94%	152	4.4	0.64	92%	186	4.4	0.63	92%
File and directory operations	34	4.4	0.69	88%	151	4.4	0.62	95%	185	4.4	0.63	94%
Reliability (data integrity)	34	4.4	0.66	91%	148	4.5	0.57	97%	182	4.5	0.58	96%
Frequency of outages	33	4.3	0.68	88%	145	4.2	0.71	88%	178	4.2	0.71	88%
Overall satisfaction with Alpine GPFS Scratch filesystem	34	4.5	0.56	97%	154	4.4	0.56	97%	188	4.4	0.56	97%

Table 25. *Satisfaction Ratings of Alpine GPFS Scratch Filesystem by Project Allocation*

	INCITE				DD				ALCC			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Size	116	4.5	0.58	96%	73	4.4	0.62	93%	57	4.4	0.59	95%
I/O bandwidth	116	4.5	0.58	96%	71	4.3	0.69	87%	55	4.3	0.64	89%
File and directory operations	116	4.4	0.61	94%	70	4.3	0.67	91%	56	4.2	0.65	91%
Reliability (data integrity)	114	4.5	0.57	96%	68	4.5	0.58	96%	55	4.5	0.54	98%
Frequency of outages	109	4.2	0.73	89%	69	4.2	0.7	86%	54	4.1	0.66	85%
Overall satisfaction with Alpine GPFS Scratch filesystem	116	4.5	0.55	97%	72	4.4	0.57	96%	56	4.3	0.54	96%
	ECP				**ES**				**General**			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Size	55	4.4	0.62	93%	20	4.7	0.59	95%	8	4.5	0.53	100%
I/O bandwidth	54	4.4	0.6	94%	20	4.7	0.59	95%	9	4.3	0.5	100%
File and directory operations	53	4.4	0.66	94%	20	4.6	0.6	95%	9	4.6	0.53	100%
Reliability (data integrity)	53	4.5	0.54	98%	20	4.5	0.6	95%	9	4.3	0.5	100%
Frequency of outages	49	4.3	0.63	92%	20	4.3	0.73	85%	8	4	0.53	88%
Overall satisfaction with Alpine GPFS Scratch filesystem	55	4.4	0.57	96%	20	4.5	0.6	95%	9	4.3	0.5	100%

Support Services

The Support Services element of the survey was substantially re-written in 2019, excluding many questions that had been asked of users in previous years. 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 (93%) were either *satisfied* or *very satisfied* with support received from user assistance, accounts, INCITE Scientific Computing Liaisons, and Advanced Data/Workflow Liaisons (Table 26 and Table 27). Five respondents reported **reasons for dissatisfaction**:

> *"Long delays in communication response and lack of support for technically demanding GPU-enabled implementations on IBM processing architecture of Summit."*
>
> *"The user support ticket system should be directly accessible to ticket submitters, their projects and anyone they CC."*
>
> *"We additionally had trouble getting serial processes to work efficiently with the wraprun commands. The user support were very helpful in providing us with some workarounds to prevent us from being significantly delayed in the actual data acquisition portion of our project."*
>
> *"We rarely get support unless it directly benefits the liaison or OLCF. When we have identified issues, the first response is always that the problem is with our code (even when it isn't). This simply exacerbates the issue, since we then have to spend effort demonstrating that the problem really is with the system and not with our code base."*

Table 26. *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)	97	4.6	0.71	93%	405	4.6	0.64	94%	502	4.6	0.66	93%

Table 27. *Satisfaction Ratings of Support Received by Project Allocation*

	INCITE				DD				ALCC			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Support received (user assistance, accounts, INCITE Liaisons, Advanced Data/Workflow Liaisons)	272	4.6	0.65	94%	175	4.6	0.67	94%	150	4.6	0.62	93%
	ECP				**ES**				**General**			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Support received (user assistance, accounts, INCITE Liaisons, Advanced Data/Workflow Liaisons)	118	4.6	0.60	96%	45	4.4	0.81	84%	23	4.7	0.56	96%

OLCF Website

Fifty-three percent of survey respondents indicated that they had visited the OLCF website during 2019 (*N* = 308). 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), as displayed in Figure 7.

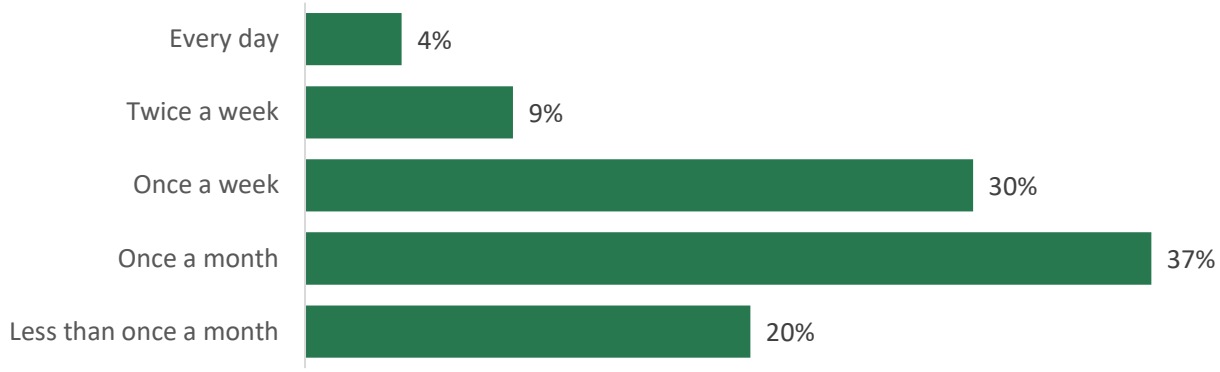

Figure 7. Frequency with which OLCF users visit the OLCF website (*N* = 305)

Users rated several aspects of the website (Table 28 and Table 29). 93% 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;** most commonly, they related to difficulty finding information:

> *"I usually have a hard time finding the information I need or feel like it is not in a logical location. This issue has been improving with recent updates to the site though and is less of a problem now."*
>
> *"The Titan and Summit documentation is hard to navigate."*

Communication with Users

As seen in Table 30, most respondents (95%) were either *satisfied* or *very satisfied* with how OLCF keeps them informed of changes, events, downtimes, and current issues (Table 30 and Table 31). Only one respondent provided an explanation for their **dissatisfaction with OLCF communication:**

> *"Last I looked I did not see a contact specifically for Summit. I have some questions, but am hesitant to use the general support@ address."*

Table 28. *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	61	4.4	0.66	93%	244	4.2	0.71	88%	305	4.3	0.70	89%
Search capabilities	59	4.2	0.74	85%	227	4.1	0.73	80%	286	4.1	0.73	81%
Usefulness of content	61	4.5	0.65	95%	245	4.4	0.61	94%	306	4.4	0.62	94%
Overall satisfaction with the OLCF website	61	4.4	0.66	93%	245	4.4	0.65	93%	306	4.4	0.65	93%

Table 29. *Satisfaction Ratings of the OLCF Website by Project Allocation*

	INCITE				DD				ALCC			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Ease of navigation	169	4.2	0.75	86%	118	4.3	0.60	95%	92	4.2	0.67	86%
Search capabilities	158	4.1	0.74	80%	114	4.1	0.74	80%	85	4.1	0.71	79%
Usefulness of content	170	4.4	0.63	94%	118	4.5	0.60	97%	92	4.4	0.60	95%
Overall satisfaction with the OLCF website	170	4.3	0.67	92%	118	4.4	0.60	96%	92	4.3	0.62	92%
	ECP				**ES**				**General**			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Ease of navigation	81	4.3	0.65	91%	33	4.2	0.83	79%	15	4.4	0.63	93%
Search capabilities	74	4.2	0.72	84%	32	4.1	0.76	75%	14	4.3	0.47	100%
Usefulness of content	82	4.5	0.57	96%	33	4.3	0.69	88%	15	4.5	0.52	100%
Overall satisfaction with the OLCF website	82	4.4	0.56	96%	33	4.3	0.76	82%	15	4.5	0.52	100%

Table 30. *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	101	4.7	0.51	98%	453	4.5	0.61	94%	554	4.6	0.59	95%

Table 31. *Satisfaction Ratings of Communications by Project Allocation*

	INCITE				DD				ALCC			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Overall communications	303	4.6	0.62	94%	188	4.6	0.56	97%	165	4.5	0.61	94%
	ECP				**ES**				**General**			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Overall communications	129	4.6	0.60	95%	50	4.5	0.68	90%	26	4.7	0.55	96%

Problem Resolution

Figure 8 shows how frequently respondents **submitted queries** to OLCF (via phone or email) in 2019. Half submitted between one and five requests, while **one-third had not submitted any queries at all.**

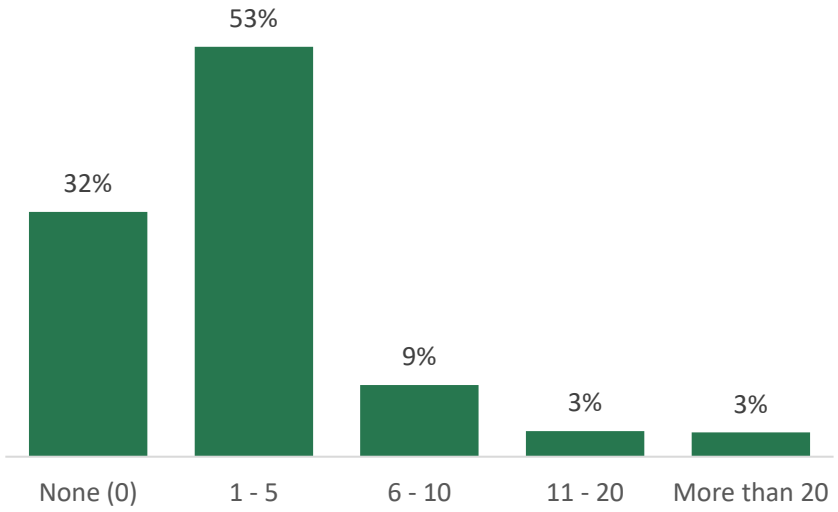

Figure 8. Distribution of number of queries submitted to OLCF in 2019 (*N* = 574)

Users were asked to provide satisfaction ratings for their overall satisfaction with OLCF's problem resolution and three specific aspects (Table 32 and Table 33). 93% 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 was the* lowest rated feature. Six respondents reported **reasons for dissatisfaction with problem resolution**:

> *"I was raising issues of node and communication failures, and I haven't got a good explanation or solution."*
>
> *"Support ticket system should be open to users."*
>
> *"The only ticket I've submitted was nearly eight weeks ago and I haven't gotten a follow-up."*
>
> *"Too much latency."*

Table 32. *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	90	4.6	0.73	92%	344	4.6	0.64	93%	434	4.6	0.66	93%
Timeliness of OLCF responses to reported issues	90	4.5	0.71	94%	343	4.5	0.65	92%	433	4.5	0.66	92%
Usefulness of support and training documentation	88	4.4	0.71	92%	359	4.5	0.66	91%	447	4.4	0.67	91%
Overall satisfaction with problem resolution	91	4.5	0.74	92%	353	4.5	0.63	94%	444	4.5	0.65	93%

Table 33. *Satisfaction Ratings of OLCF's Problem Resolution by Project Allocation*

	INCITE				DD				ALCC			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Quality of OLCF response to reported issues	234	4.6	0.63	93%	155	4.6	0.66	94%	129	4.5	0.70	91%
Timeliness of OLCF responses to reported issues	233	4.5	0.64	92%	155	4.5	0.67	94%	129	4.5	0.67	91%
Usefulness of support and training documentation	239	4.4	0.68	90%	157	4.5	0.68	92%	138	4.4	0.65	93%
Overall satisfaction with problem resolution	238	4.6	0.60	94%	156	4.5	0.69	94%	136	4.5	0.68	93%
	ECP				**ES**				**General**			
	N	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***	***N***	***M***	***SD***	***%Sat***
Quality of OLCF response to reported issues	104	4.5	0.70	92%	39	4.5	0.68	90%	19	4.6	0.60	95%
Timeliness of OLCF responses to reported issues	104	4.5	0.62	93%	39	4.6	0.64	92%	19	4.7	0.48	100%
Usefulness of support and training documentation	105	4.5	0.65	93%	43	4.5	0.63	93%	20	4.6	0.51	100%
Overall satisfaction with problem resolution	105	4.6	0.65	93%	42	4.5	0.59	95%	19	4.6	0.50	100%

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. Of the 574 respondents who answered this question, 13% (*n* = 75) indicated that they do not need or use data analysis. These responses were then excluded from further analysis, and the distribution of other responses in 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.

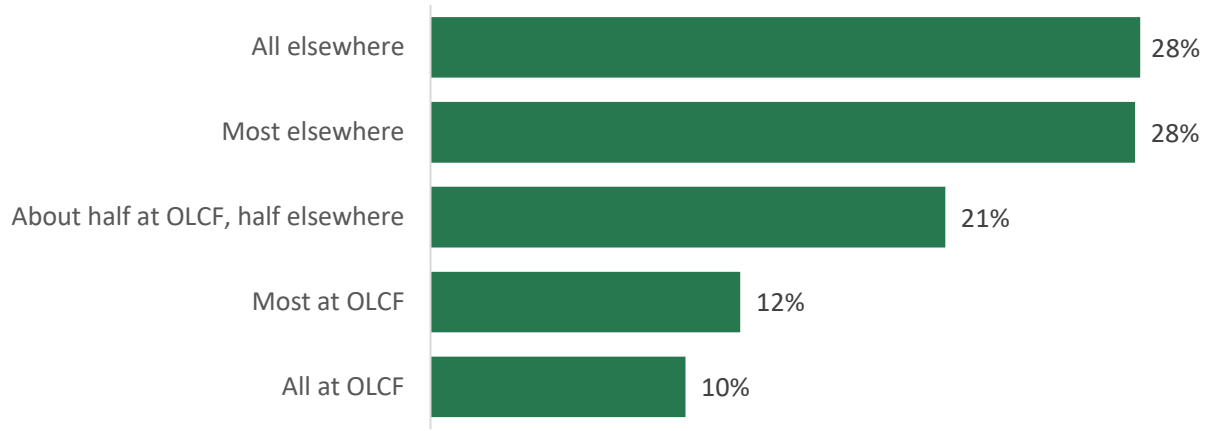

Figure 9. Locations for analysis of data by OLCF users (*N* = 499)

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.

Figure 10. Source of user data (*N* = 574)

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. Thirty-six (9%, *N* = 414) respondents 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?" 113 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 34). Among those expressing a need or preference, *Summit* was mentioned most frequently, followed by *tools/software/installations, documentation* and *development/debugging and test resources/queue*. See Appendix E: User Suggestions for Improvement for all responses by category. Select comments include:

Summit	*"Better stability with the updates to Summit, or frankly just less of them. I understand that this is partly from IBM and OLCF may not have as much control over this."* *"When new systems arrive (example Summit), they should be made available only when they are stable. We experienced several issues with Summit not being ready for production. If a system is used for production at an early stage, then usage of allocated time should be heavily discounted."*
Tools/Software/Installations	*"Please make more tools to help make things easier like filling burst buffers on job submission, and breaking apart job allocations to run lots of smaller parallel jobs."* *"I have not tried out the Pytorch implementation, but deep learning frameworks that can easily incorporate new methods/algorithms are very important for me. I encourage keeping at the cutting edge of this field with implementation on OLCF resources."*
Documentation	*"A section in the user guide collecting "common mistakes", similar to the known issues, but not necessarily bugs, would be very helpful."* *"More documents are needed for Summit."*
Development/Debugging and Test Resources/Queue	*"A development queue that allows for longer duration small test problems. We often need that to prepare for a major run."* *"When our allocation is exhausted and we can continue to run only a single job, it would be VERY valuable to have additional access to a debug queue so many people could prepare code for later proposals without stopping what is typically a very large and valuable production job."*

Table 34. *Users' Suggestions for Additional Services and/or Resources Needed to Enhance Their Experience at the OLCF*

Category	*N* = 113	Percentage
Summit	19	17%
Tools/software/installations	16	14%
Documentation	16	14%
Development/debugging and test resources/queue	11	10%
Login, access, SSH & authentication	10	9%
Libraries and updates	9	8%
Training and tutorials	9	8%
Queue time and scheduling policy	9	8%
General tech support/assistance	8	7%
Smaller Jobs	8	7%
GPU resources	7	6%
Specialized support/developers/experts	7	6%
Performance/performance upgrade	6	5%
Allocations	6	5%
Satisfaction	5	4%
Compilers	5	4%
Filesystem, I/O and data transfer	5	4%
Wall/run time	5	4%
Containers	4	4%
Stability/reliability	4	4%
Visualization and analysis	3	3%
Data storage	3	3%
Website	2	2%
Accounts and account status information	2	2%
Data retention/purge policy and procedures	1	1%
Architecture	1	1%
Workflow	1	1%
Miscellaneous/Other	12	11%

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, Titan, Eos, Rhea, DTNs, HPSS, Lustre/Spider, Alpine GPFS) and/or tell us if any additional resources are needed," 111 respondents provided comments. The largest proportion, 28%, indicated their *satisfaction* with OLCF compute and data resources. This was followed by *need for specific software/tools/modules* (23%), *more documentation* (8%), and *file systems and data transfer* (8%). See Appendix E: User Suggestions for Improvement for all responses by category. Refer to Table 35 for all themes identified.

Select comments include:

Need for Specific Software/Tools/Modules	*"Helpful to have a "bleeding edge" set of compilers available as a module. e.g. Latest gcc, clang, intel, XL, maybe beta versions etc."* *"My only complaint is that there were modules available in Titan that were not ported over to Summit."* *"On Summit, several of the provided Python versions do not come with default optimized numpy package, which is a must-have for scientific applications."*
More Documentation	*"As with Titan, a tutorials page listing different tutorials for running, debugging, and optimizing applications on Summit would be welcome. For example, OpenACC examples showing how to utilize multiple GPUs using simple codes such as SAXPY would be appreciated."* *"It might be useful to develop documentation for the use of common machine learning tools like Tensorflow, and sci-kit learn."*
File Systems and Data Transfer	*"I have still had issues with Lustre, both in when interacting via simple file system commands, and with I/O, although to be truthful, I haven't spent much time with the I/O issue this year, so take this as somewhat anecdotal. Could be metadata-related in both cases, which would be consistent with past experience. On the other hand, I had very good experience this (fiscal) year using the data transfer and hpss, with no problems, so...."* *"Shells on the login nodes working within directories on the alpine filesystem stop responding for a few to several minutes on a semi-regular basis."* *"The OLCF compute facilities are truly exemplary. It has been an absolute pleasure to compute on the OLCF systems. One persistent issue seems that the GPFS filesystem will occasionally (rarely, but with some frequency) hang from simple filesystem requests (e.g., ls). The filesystem will then return to normal operation."*

Table 35. *Users' Suggestions for Improvements to HPC Compute and Data Resources*

Category	*N* = 111	Percentage
Expressions of satisfaction	31	28%
Need for specific software/tools/modules	25	23%
More documentation	9	8%
File systems and data transfer	9	8%
Tutorials/Training/Instructional resources	8	7%
Wall time/runtime and queues	7	6%
Support and ticketing	7	6%
Administrative issues	7	6%
Issues with software	6	5%
Provide unpurged storage/scratch space	6	5%
Improve debugging/development/test	6	5%
Reliability/stability/uptime of systems	5	5%
CPU/Titan/EOS functionality replacements	5	5%
Updates to data retention/purge policy and procedures	4	4%
Support for small jobs	4	4%
SSH/accessibility and authentication	3	3%
Performance	3	3%
GPU computing/resources	3	3%
Miscellaneous/Other	11	10%

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?" 82 users responded. An additional 76 users provided comments indicating they were satisfied with current services. Among those who expressed a need/preference, the largest proportions were interested in *a variety of analysis and visualization tools/software,* and specifically with *Jupyter* and *Python*-related capabilities (Figure 11).

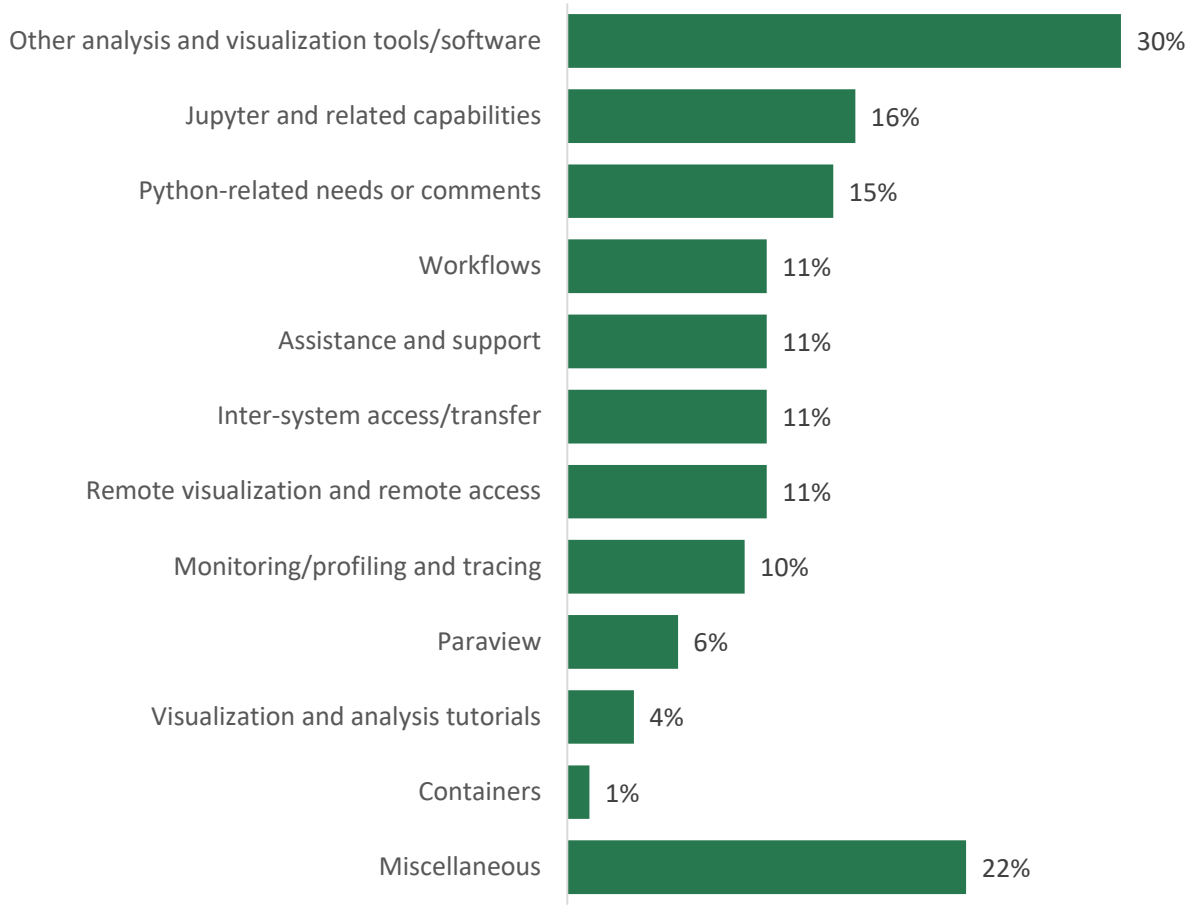

Figure 11. Users' suggestions for additional data analysis, visualization, and/or workflow services (*N* = 82).

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	*"I believe python and perl scripting are already supported? My current workflow uses those, as well as various plotting tools (gnuplot, matplotlib, pgfplot). It'd be nice to have the latter three (and they might already be supported), but not required."*