ORNL/SPR-2017/486

US Department of Energy, Office of Science

High Performance Computing Facility Operational Assessment 2016 Oak Ridge Leadership Computing Facility

Ryan Adamson
Ashley D. Barker
Arthur S. Bland
James J. Hack
Jason Hill
Stephen T. McNally
Gordon Rhyne
Mallikarjun Shankar
T. P. Straatsma
Kevin G. Thach
Suzy Tichenor
Sudharshan S. Vazhkudai
Jack C. Wells

Approved for public release.
Distribution is unlimited.

March 2017

OAK RIDGE NATIONAL LABORATORY
MANAGED BY UT-BATTELLE FOR THE US DEPARTMENT OF ENERGY

ORNL/SPR-2017/486

Oak Ridge Leadership Computing Facility

HIGH PERFORMANCE COMPUTING FACILITY OPERATIONAL ASSESSMENT 2016
OAK RIDGE LEADERSHIP COMPUTING FACILITY

Ryan Adamson
Ashley D. Barker
Arthur S. Bland
James J. Hack
Jason Hill
Stephen T. McNally
Gordon Rhyne
Mallikarjun Shankar
T. P. Straatsma
Kevin G. Thach
Suzy Tichenor
Sudharshan S. Vazhkudai
Jack C. Wells

Date Published: March 2017

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, TN 37831-6283
managed by
UT-BATTELLE, LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

Page

LIST OF FIGURES

Figure **Page**

LIST OF TABLES

Table **Page**

ACRONYMS

ACCEL	Accelerating Competitiveness through Computational ExceLlence
ADIOS	Adaptable I/O System
ADW	Advanced Data and Workflows Group
ALCC	ASCR Leadership Computing Challenge
ALCF	Argonne Leadership Computing Facilities
ANL	Argonne National Laboratory
API	application programming interface
ARM	Atmospheric Radiation Measurement
ASCR	Advanced Scientific Computing Research
CAAR	Center for Accelerated Application Readiness
CADES	Compute and Data Environment for Science
CANDLE	CANcer Distributed Learning Environment
CFD	computational fluid dynamics
CSEEN	Computational Scientists for Energy, the Environment, and National Security
CSGF	Computational Science Graduate Fellowship
CT	computer tomography
CY	calendar year
DART	days away, restricted, or transferred
DD	Director's Discretionary
DEC	Divide-Expand-Consolidate
DOE	Department of Energy
DOI	digital object identifier
DL	deep learning
DTN	data transfer node
ECP	Exascale Computing Project
EVEREST	Exploratory Visualization Environment for Research in Science and Technology
FP	functional partitioning
HPC	high-performance computing
HPSS	high-performance storage system
IDEAS	Interoperable Design of Extreme-Scale Application Software
INCITE	Innovative and Novel Computational Impact on Theory and Experiment
ISSAOS	International Summer School on Atmospheric and Oceanic Sciences
JIF	journal impact factor
LANL	Los Alamos National Laboratory
LCF	Leadership Computing Facility
LES	large eddy simulation
LBNL	Lawrence Berkeley National Laboratory
LLNL	Lawrence Livermore National Laboratory
LOA	level of assurance
MFA	multifactor authentication
micro-CT	micro-computed tomography
MPI	Message Passing Interface
MTTF	mean time to failure
MTTI	mean time to interrupt
NAM	not a metric
NCCS	National Center for Computational Sciences
NERSC	National Energy Research Scientific Computing Center
NICS	National Institute for Computational Sciences

NUCCOR	Nuclear Coupled Cluster Oak Ridge
OA	overall availability
OAR	Operational Assessment Report
OLCF	Oak Ridge Leadership Computing Facility
OMB	Office of Management and Budget
CNMS	Center for Nanophase Materials Science
ORNL	Oak Ridge National Laboratory
OUG	OLCF User Group
pbdR	Programming with Big Data
PI	principal investigator
PIV	personal identity verification
RATS	Resource and Allocation Tracking System
RMS	root-mean-square
RSS	research safety summaries
RTM	Reverse Time Migration
SciComp	Scientific Computing Group
SA	scheduled availability
SM	streaming multiprocessor
SMP	symmetric multiprocessing
SNL	Sandia National Laboratories
UA	User Assistance
UAO	User Assistance and Outreach
UT	University of Tennessee
WiC	Women in Computing

Executive Summary

HIGH PERFORMANCE COMPUTING FACILITY
2016 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2017

EXECUTIVE SUMMARY

Oak Ridge National Laboratory's (ORNL's) Leadership Computing Facility (OLCF) continues to surpass its operational target goals: supporting users; delivering fast, reliable systems; creating innovative solutions for high-performance computing (HPC) needs; and managing risks, safety, and security associated with operating one of the most powerful computers in the world. The results can be seen in the cutting-edge science delivered by users and the praise from the research community.

Calendar year (CY) 2016 was filled with outstanding operational results and accomplishments: a very high rating from users on overall satisfaction for the third year in a row; the greatest number of core hours delivered to research projects; the largest number of research publications for the life of Titan; and success in delivering on the allocation of 60%, 30%, and 10% of core hours offered for the Innovative and Novel Computational Impact on Theory and Experiment (INCITE), Advanced Scientific Computing Research Leadership Computing Challenge (ALCC), and Director's Discretionary (DD) programs, respectively. These accomplishments, coupled with the extremely high utilization rate, represent the fulfillment of the promise of Titan: maximum use by maximum-size simulations.

The impact of these successes and more is reflected in the accomplishments of OLCF users, with publications this year in notable journals such as *Nature*, *Chemical Reviews*, *Science*, *Nature Physics*, *Nature Climate Change*, *Nano Letters*, *ACS Nano*, *Journal of the American Chemical Society*, *Nature Chemical Biology*, *Nature Geoscience*, *Nature Energy*, *Nature Plants*, *Advanced Functional Materials*, *Nature Communications*, *Proceedings of the National Academy of Sciences*, *Physical Review X*, *Physical Review Letters*, as well as many others. The achievements included in the *2016 OLCF Operational Assessment Report* (OAR) reflect crucial domain-specific discoveries aided by resources at the OLCF; for example, Titan enabled incorporation of micro-computed tomography data in the study of multiphase flows in porous media within the earth's subsurface leading to breakthroughs in oil reservoir modeling and applications (Section 3.2.4).

The Titan system provides the largest extant heterogeneous architecture for computing and computational science. Usage is high, delivering on the promise of a system well-suited for capability simulations for science. This success is a result of the extraordinary work of the staff supporting the nation's leading HPC facility for the Department of Energy. The OLCF staff are pivotal to identifying, developing, and deploying the innovative processes and technologies that have helped the OLCF and other facilities realize success. The facility's leadership in the area of data analysis and workflows continued in 2016 with the addition of the highly successful Compute and Data Environment for Science (CADES) user facility to the National Center for Computational Sciences (NCCS). This organizational move demonstrates the OLCF's commitment to enabling workflows that span security domains and facilities.

Operating a world-class HPC facility that offers the breadth of resources that the OLCF offers is not without its challenges. The Titan system turned four in 2016 and began to show increased but manageable hardware failure signatures. An increased failure rate is certainly not ideal, but it is expected in a large heterogeneous system that is aging. The OLCF worked closely with Cray and NVIDIA to characterize, quantify, and address the increased hardware failures. In December 2016, the OLCF moved from a replace-on-failure service agreement with Cray to a predictive replacement strategy to minimize the impact on actively running jobs. Through rigorous tracking and careful daily management of this process, the OLCF has observed a drop in the number and frequency of failures in less than 1 month from beginning the predictive replacements.

Effective operations of the OLCF play a key role in the scientific missions and accomplishments of its users. Building on the exemplary accomplishments in 2015—shown in the 2015 OAR review committee response (Appendix A)—this OAR delineates the policies, procedures, and innovations implemented by the OLCF to continue delivering a leadership class resource for cutting-edge research. This report covers CY 2016, which denotes the period from January 1, 2016, to December 31, 2016, unless otherwise specified.

COMMUNICATIONS WITH KEY STAKEHOLDERS

Communication with the Program Office

The OLCF regularly communicates with the Advanced Scientific Computing Research Program Office through a series of regularly occurring events. These include weekly Integrated Project Team calls with the local Department of Energy ORNL Site Office and the ASCR Program Office, monthly highlight reports, quarterly reports, the annual OAR, an annual "Budget Deep Dive," an annual independent project review, and the OLCF annual report. Through a team of communications specialists and writers working with our users and management, the OLCF produces a steady flow of reports and highlights for sponsors, current and potential users, and the public.

Communication with the User Community

The OLCF's communications with users take a wide variety of forms and are tailored to the objectives of relating science results to the larger community and helping users to more efficiently and effectively use OLCF systems. The OLCF offers many training and educational opportunities throughout the year for current facility users and the next generation of HPC users (Section 1.4.7).

The impact of OLCF communications is assessed as part of an annual user survey. The mean rating for users' overall satisfaction with OLCF communications was 4.4 for the second year in a row. Ninety-five percent of respondents (346) rated their overall satisfaction with communications from the OLCF as "satisfied" or "very satisfied." In addition, nearly all of the 346 users felt well informed about OLCF changes (99%), events (99%), and current issues (98%). The OLCF uses a variety of methods to communicate with users, including the following:

- weekly email message
- welcome packet
- general email announcements
- automated notifications of system outages
- OLCF website
- conference calls
- OLCF User Council and Executive Board
- one-on-one interactions with liaisons and analysts
- social networking vehicles
- annual face-to-face OLCF User Meeting
- targeted training events (i.e., GPU Hackathons or New User Training)

SUMMARY OF 2016 METRICS

In consultation with the Department of Energy program manager and as proposed in the 2015 OAR, a series of metrics and targets were identified to assess the operational performance of the OLCF in CY 2016. The 2016 metrics, target values, and actual results as of December 31, 2016, are noted throughout this report and are summarized in Section 8. The OLCF exceeded all of the agreed on metric targets.

RESPONSES TO RECOMMENDATIONS FROM THE 2015 OPERATIONAL ASSESSMENT REVIEW

The OLCF received one recommendation from the 2015 OAR Review, which reads:

"OLCF created a clear and useful formula to define %gpu_activity, which the center is able to measure. However, values of %gpu_activity were not reported as such in the document. A table like 2.13 for GPU activity percentage would be interesting, including the overall usage."

The OLCF provided the following written response addressing the recommendation:

"GPU_activity values were included in the 2015 OLCF OAR in Figure 2.4. However, a table similar to table 2.13 from the 2015 OLCF OAR is provided below for convenience. An initial hypothesis was that the GPU_activity data would show usage efficiencies of the GPUs. While reviewing the data, it became clear that the data is not particularly helpful at the macro analysis level. Given that this measure reports GPU SM usage and does not take into account memory or I/O usages, is it short sighted and assumes that SM usage is the one and only measure that is relevant when evaluating GPU usage effectiveness. As such, the OLCF decided to collect and monitor this data, but that it was non-essential to track for the operation of Titan. Rather, the GPU_activity data provides a reasonable high-level view of the overall activity which is useful for internal planning and tracking purposes. This is especially helpful when OLCF staff partner with research teams who are interested in optimizing their applications. It is also worth noting that while the industry remains very focused on effective usage of GPUs, we do not currently measure usage efficiency on other system components such as disk, cpu, or memory. For these reasons, the OLCF has decided to continue to track and measure GPU_activity for internal purposes but will not include it as a standard reportable metric."

Table ES.1. GPU_activity by core hours

	GPU_activity (in millions)				*GPU_enabled (in millions)*			
	INCITE	**ALCC**	**DD**	**Total Hours**	**DD**	**ALCC**	**INCITE**	**Total Hours**
Jan	43.51	24.72	2.50	70.73	6.95	90.51	61.51	158.97
Feb	42.48	10.76	2.17	55.41	7.55	37.77	91.84	137.17
Mar	38.10	18.58	1.01	57.69	3.45	60.32	115.68	179.45
Apr	34.90	8.35	4.76	48.00	12.86	31.38	137.79	182.03
May	35.79	11.50	2.07	49.37	8.65	52.49	126.78	187.93
Jun	24.27	13.55	3.37	41.19	10.00	31.95	122.77	164.72
Jul	24.94	13.06	3.37	41.37	7.97	18.22	126.05	152.25
Aug	60.25	31.54	5.89	97.68	17.50	52.60	115.82	185.92
Sep	39.21	61.23	3.12	103.55	6.14	91.08	136.37	233.59
Oct	51.98	26.12	4.10	82.20	12.36	65.57	168.73	246.66
Nov	47.51	12.12	4.30	63.94	13.17	29.13	150.82	193.12
Dec	54.09	8.80	2.23	65.13	4.27	13.86	184.89	203.02
Total	497.04	240.33	38.89	776.26	110.88	574.88	1,539.06	2,224.83

OPERATIONAL REALIGNMENTS TO BETTER SERVE OLCF STAKEHOLDERS

In 2016, ORNL's Compute and Data Environment for Science (CADES) facility, which provides a compute and science-data environment to accelerate scientific discovery for ORNL staff and collaborators, was realigned by ORNL management to better serve the goals of this ambitious initiative. This realignment moved the CADES team from ORNL's information technology organization to the NCCS. This move represents an important evolution and an innovative step in the deployment of new computational and data capabilities that are highly desirable to existing CADES and OLCF users. Since the NCCS now oversees both user facilities, this move allows better integration and reuse of data services and analytics capabilities. For years, the OLCF has been a flagship user facility for leadership HPC capabilities and has done so within the operational constraints of a moderate security enclave. Similarly, the CADES user facility has been successful with ORNL-based research teams who require advanced workflows and collaborations that an open computing environment provides. ORNL management recognized the complementary nature of these functions and opted to ensure they are consistently managed by implementing the reorganization. While the move occurred late in CY 2016, work to integrate the two activities began immediately after the decision.

The CADES user facility was integrated into the NCCS ADW group. Given the success of CADES, the recent formation of the ADW group, and the overlapping missions of both, the CADES user facility found a natural home within the NCCS organization. Mallikarjun "Arjun" Shankar retained his role as CADES Director and was also given oversight of the ADW group as its group leader.

This move has tremendous potential to facilitate greater innovation in the provision of a wide range of computational and data capabilities. Early integration meetings have shown promise toward establishing and advancing high performing data services across CADES and the OLCF. While the integration will take time to fully execute, users of both facilities will gain much in the way of closer coupled workflows between the DOE facilities and will benefit from the tiered security environments tailored to specific use cases.

User Results

HIGH PERFORMANCE COMPUTING FACILITY
2016 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2017

1. USER RESULTS

CHARGE QUESTION 1: Are the processes for supporting the customers, resolving problems, and Outreach effective?

OLCF RESPONSE: Yes. In 2016, the Oak Ridge Leadership Computing Facility (OLCF) supported 1,186 users and 318 projects. The OLCF continued to leverage an established user support model for effectively supporting users based on continuous improvement, regular assessment, and a strong customer focus. One key element of internal assessment is the annual user survey. As part of the survey, users are asked to rate their overall satisfaction with the OLCF on a scale of 1 to 5, with a rating of 5 indicating "very satisfied." The mean rating for overall satisfaction with the OLCF in 2016 was 4.6. Overall ratings for the OLCF were positive; 95% of users reported being "satisfied" or "very satisfied."

The center measures its performance using a series of quantifiable metrics. The metric targets are structured to ensure that users are provided prompt and effective support and that the organization responds quickly and effectively to improve its support process for any item that does not meet a minimum satisfactory score. The OLCF exceeded all metric targets for user satisfaction in 2016 with 92% of tickets being resolved within 3 business days. The OLCF continued to enhance its technical support, collaboration, training, outreach, and communication and engaged in activities that promoted high-performance computing (HPC) to the next generation of researchers.

1.1 USER RESULTS SUMMARY

The OLCF's user support model comprises customer support interfaces, including user satisfaction surveys, formal problem-resolution mechanisms, user assistance analysts, and scientific liaisons; multiple channels for stakeholder communication, including the OLCF User Council; and training programs, user workshops, and tools to reach and train both current facility users and the next generation of computer and computational scientists. The success of these activities and identification of areas for development are tracked using the annual OLCF user survey.

To promote continual improvement at the OLCF, users are sent surveys soliciting their feedback regarding support services and their experience as users of the facility. The 2016 survey was launched on October 10, 2016, and remained open for participation through November 21, 2016. The survey was sent to 950 users of the Innovative and Novel Computational Impact on Theory and Experiment (INCITE), Advanced Scientific Computing Research (ASCR) Leadership Computing Challenge (ALCC), and Director's Discretionary (DD) projects who logged into an OLCF system between January 1, 2016, and September 30, 2016. OLCF staff members were excluded from participation. A total of 369 users completed the survey, for an overall response rate of 39%. The results of the 2016 survey can be found on the OLCF website.[1]

The effectiveness of the processes for supporting customers, resolving problems, and conducting outreach are defined by the metrics in Table 1.1 and are assessed by the user survey.

[1] https://www.olcf.ornl.gov/media-center/center-reports/2016-outreach-survey/

Table 1.1. 2016 user result metrics summary

Metric description	2015 target	2015 actual	2016 target	2016 actual
Overall OLCF satisfaction score on the user survey	3.5/5.0	4.6/5.0	3.5/5.0	4.6/5.0
Show improvement on results that scored below satisfactory in the previous period	Show improvement in ½ of questions that scored below satisfactory (3.5) in the previous period.	No question scored below satisfactory (3.5/5.0) on the 2015 survey.	Show improvement in ½ of questions that scored below satisfactory (3.5) in the previous period.	No question scored below satisfactory (3.5/5.0) on the 2016 survey.
OLCF survey results related to problem resolution	3.5/5.0	4.6/5.0	3.5/5.0	4.6/5.0
Percentage of user problems addressed within 3 business days	80%	92%	80%	92%
Average of all user support services ratings	3.5/5.0	4.6/5.0	3.5/5.0	4.6/5.0

1.2 USER SUPPORT METRICS

The OLCF exceeded all user support metrics for 2016. The OLCF metric targets and actual results by calendar year (CY) for user support are shown in Table 1.2.

Table 1.2. OLCF user support summary: Metric targets and calendar year results

Survey area	CY 2015		CY 2016	
	Target	Actual	Target	Actual
Overall OLCF satisfaction rating	3.5/5.0	4.6/5.0	3.5/5.0	4.6/5.0
Average of all user support services ratings	3.5/5.0	4.6/5.0	3.5/5.0	4.6/5.0

1.2.1 Overall Satisfaction Rating for the Facility

Users were asked to rate their satisfaction on a 5-point scale, where a score of 5 indicates a rating of "very satisfied," and a score of 1 indicates a rating of "very dissatisfied." The metrics are agreed on by the Department of Energy (DOE) and OLCF program manager, who defined 3.5/5.0 as satisfactory. Overall ratings for the OLCF were positive with 95% of users responding that they are satisfied or very satisfied with the OLCF overall.

Key indicators from the survey, including overall satisfaction, are shown in Table 1.3. They are summarized and presented by program respondents. The data show the satisfaction among all allocation programs is similar for the four key satisfaction indicators.

Table 1.3. Satisfaction rates by program type for key indicators

Indicator	Mean	Program		
		INCITE	ALCC	DD
Overall satisfaction with the OLCF	4.6/5.0	4.6/5.0	4.6/5.0	4.6/5.0
Overall satisfaction with support services	4.5/5.0	4.6/5.0	4.5/5.0	4.5/5.0
Overall satisfaction with compute resources	4.5/5.0	4.4/5.0	4.5/5.0	4.6/5.0
Overall satisfaction with data resources	4.5/5.0	4.5/5.0	4.4/5.0	4.5/5.0

1.2.2 Average Rating across All User Support Questions

The calculated mean of answers to the user support services specific questions on the 2016 survey was 4.6/5.0, indicating that the OLCF exceeded the 2016 user support metric target and that users have a high degree of satisfaction with user support services. Respondents described what they perceived to be “the best qualities of OLCF.” Thematic analysis of user responses identified computing power/performance and user tech support/staff as the most valued qualities of the OLCF. Included below are several open-ended responses to “What are the best qualities of the OLCF?”

“It has some of the best computational resources in the world managed by a very capable team of computer scientists/engineers, who make freely available very useful supporting tools to take advantage of these resources.”

“The high level of technical expertise of the staff.”

“Great technical support.”

“Documentation and user support is excellent.”

“Excellent support services from our INCITE liaison.”

“Excellent facility for industrial research. Strong, well informed support staff.”

“Good user support and training opportunities (e.g., Hackathon).”

1.2.3 Improvement on Past Year Unsatisfactory Ratings

Each year the OLCF works to show improvement on no less than half of any questions that scored below satisfactory (3.5/5.0) in the previous year’s survey. All questions scored above 3.5 on both the 2015 and 2016 surveys. However, based on feedback received from the 2015 User Survey in conjunction with other feedback channels, the OLCF took the following actions in 2016 to enhance the user experience at the OLCF.

- Upgraded the OLCF’s data storage services, which included quadrupling the disk cache of high-performance storage system (HPSS), which increased data intake by a factor of five and made it easier for users to migrate, preserve, and access data in long-term archival storage.
- The OLCF purchased and deployed the new Allinea Forge product, which combines Allinea’s DDT debugger and the MAP parallel profiler.
- The OLCF began a major data transfer node (DTN) hardware refresh, which included the addition of 44 high performance nodes.
- Based on direct user feedback, the OLCF made the project home file system (/ccs/proj) available to users on compute nodes of Titan. This file system provides a purge exempt area to store project files such as executables, shared objects, Python modules, etc.
- The OLCF continued to enhance the optimized job-launching tool called Wraprun. Wraprun is a utility that enables independent execution of multiple Message Passing Interface (MPI) applications under a single aprun call.
- Rhea’s queue policy was modified to permit four eligible jobs per user to match the user experience on the other OLCF computing resources.
- Purchased Google Site Search to incorporate in the OLCF website as part of the website refresh coming in 2017 to improve the search capabilities on the OLCF website.

1.2.4 Assessing the Effectiveness of the OLCF User Survey

The survey was created by Oak Ridge Associated University's Assessment and Evaluation team in collaboration with OLCF staff. Before sending the user survey, OLCF staff met with the Oak Ridge Institute for Science Education evaluation specialist to review the content of the survey questions to ensure they accurately addressed the concerns of the OLCF and that all technical terminology was used appropriately.

Several targeted notifications were sent to those eligible to participate in the survey. Ashley Barker, the User Assistance and Outreach Group Leader, sent the initial survey invitation on October 10, 2016, and subsequent follow-up reminders were sent by Jack Wells, the National Center for Computational Sciences (NCCS) Director of Science, on October 26, 2016, by the OLCF Executive Board on November 1, 2016, and again by Ashley Barker on November 10, 2016. The final reminder was sent on November 16, 2016.

The survey was advertised on the OLCF website and in the weekly communications email sent to all users. Survey responses were tracked daily to assess the effectiveness of the various communication methods. The number of responses increased after every targeted notification, but the results show other efforts, such as including the notice in the weekly communication, also contributed to the survey response rate.

The OLCF has a relatively balanced distribution of new users, users who have been at the center for 1–2 years, and users who have been at the OLCF greater than 2 years (Table 1.4).

Table 1.4. User survey participation

	2015 survey	**2016 survey**
Total number of respondents (Total percentage responding to survey)	308 (35%)	369 (39%)
New users (OLCF user <1 year)	27%	32%
OLCF user 1–2 years	22%	25%
OLCF user >2 years	51%	43%

1.2.4.1 Statistical Analysis of the Results

The survey collected feedback about user needs, preferences, and experience with the OLCF and its support capabilities. Attitudes and opinions on the performance, availability, and possible improvements of OLCF resources and services were also solicited. Oak Ridge Associated University provided the OLCF with a written report that included the results and a summary of the findings. The findings section presents results summarized numerically that report responded levels of satisfaction. This is followed by a verbal summary of the open-ended comments from individuals who indicated they were dissatisfied (via the scaled reply) with a resource or service (note: not all dissatisfied individuals supplied open-ended comments).

The survey assessed satisfaction with OLCF resources and services using a 5-point scale, from "very dissatisfied" (1) to "very satisfied" (5). These responses were closed-ended and summarized by using frequency distributions, proportions, means, and standard deviations. Respondents who were "very dissatisfied" or "dissatisfied" with OLCF resources and services were asked to provide comments explaining their dissatisfaction. To better understand the how responses, needs, and preferences varied by types of OLCF users, closed-ended responses were frequently further separated by principal investigator (PI) status and project allocation.

Table 1.5 displays responses for five of the overall satisfaction categories broken down by allocation program. As Table 1.5 illustrates, the metrics are highly comparable across the three allocation programs, and the variations are statistically insignificant.

Table 1.5. Statistical analysis of survey results

Overall satisfaction with	INCITE			ALCC			DD			ALL		
	Mean	Var	Std Dev	Mean	Var	Std Dev	Mean	Var	Std Dev	Mean	Var	Std Dev
the OLCF	4.6	0.8	0.6	4.6	0.9	0.8	4.6	0.8	0.7	4.6	0.8	0.7
support services	4.6	0.8	0.7	4.5	0.9	0.8	4.5	0.8	0.6	4.5	0.8	0.7
user assistance team	4.7	0.8	0.6	4.6	0.8	0.6	4.6	0.7	0.5	4.6	0.7	0.5
accounts services team	4.7	0.7	0.5	4.7	0.7	0.5	4.7	0.7	0.5	4.7	0.7	0.5
compute resources	4.4	0.9	0.8	4.5	0.9	0.8	4.6	0.8	0.7	4.5	0.8	0.7

1.3 PROBLEM RESOLUTION METRICS

The following operational assessment review metrics were used for problem resolution:

- Average satisfaction ratings for questions on the user survey related to problem resolution are satisfactory or better.
- At least 80% of user problems are addressed (i.e., the problem is resolved or the user is told how the problem will be handled) within 3 business days.

1.3.1 Problem Resolution Metric Summary

In most instances, the OLCF resolves reported problems directly, including identifying and executing the necessary corrective actions. Occasionally, the facility receives problem reports that it is limited in its ability to resolve because of factors beyond the facility's control. In such a scenario, addressing the problem requires OLCF staff to identify and carry out all corrective actions at their disposal for the given situation. For example, if a user reports a suspected bug in a commercial product, prudent measures might be to recreate the issue; open a bug or ticket with the product vendor; provide the vendor necessary information about the issue; provide a workaround to the user, if possible; and track the issue to resolution with the product vendor, which may resolve the issue with a bug fix or workaround acknowledgment.

The OLCF uses Request Tracker software to track queries (i.e., tickets) and ensure response goals are met or exceeded. Users may submit queries via email, the online request form, or by phone. Email is the predominant source of query submittals. The software collates statistics on tickets issued, turnaround times, and other metrics to produce reports. These statistics allow OLCF staff to track patterns and address anomalous behaviors before they have an adverse effect on the work of other users. The OLCF issued 2,404 tickets in response to user queries for CY 2016. The center exceeded the problem-resolution metric and responded to 92% of the queries within 3 business days (Table 1.6).

Tickets are categorized by the most common types. The top three reported categories in 2016 were running jobs, account access, and job compilations (Figure 1.1).

Table 1.6. Problem resolution metric summary

Survey Area	CY 2015		CY 2016	
	Target	Actual	Target	Actual
Percentage of problems addressed in 3 business days	80%	90%	80%	92%
Average of problem resolution ratings	3.5/5.0	4.6/5.0	3.5/5.0	4.6/5.0

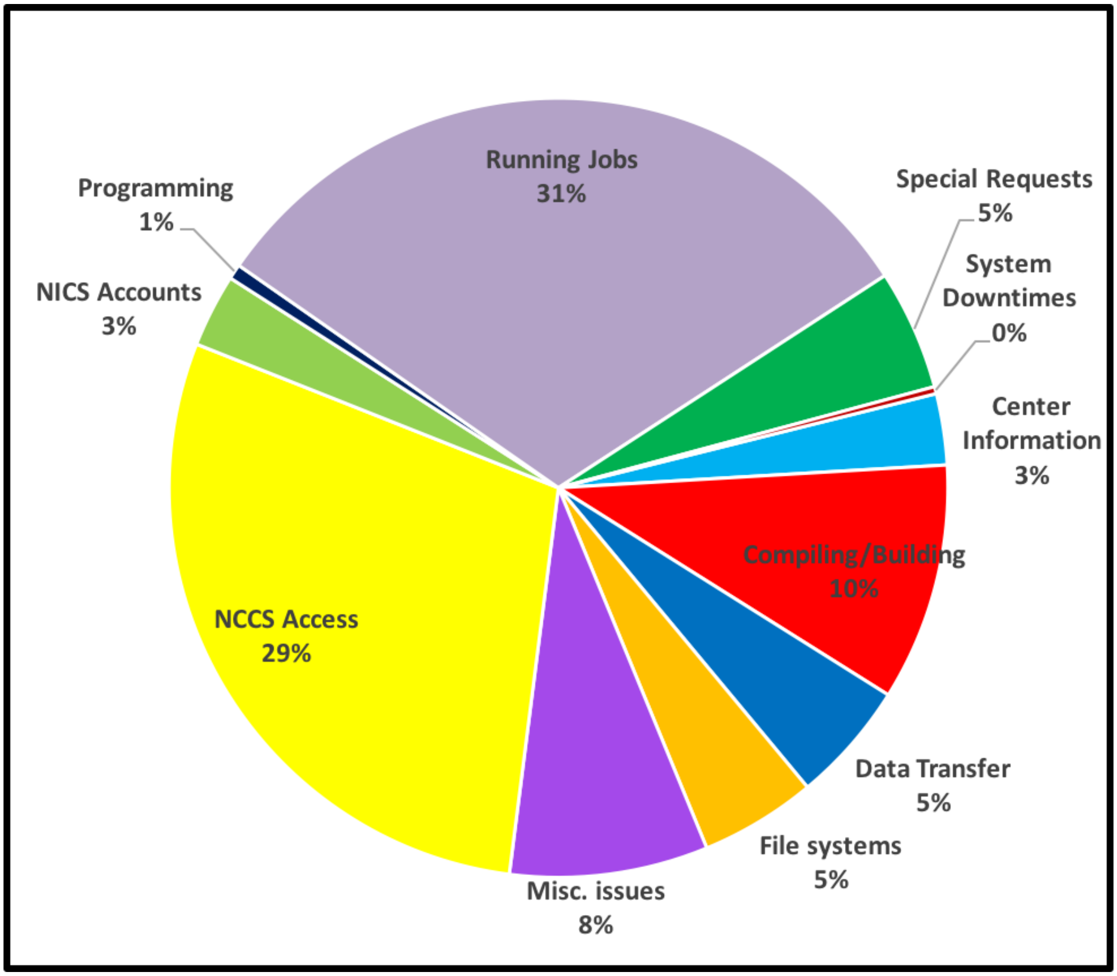

Figure 1.1. Categorization of help desk tickets.

1.4 USER SUPPORT AND OUTREACH

The OAR data requested for user support and outreach includes examples of in-depth collaboration between facility staff and the user community and a summary of training and outreach events conducted during this period (Appendices B–C). The following sections discuss key activities and contributions in the areas the OLCF recognizes as pillars of user support and outreach, including

- a user support staff made up of account management liaisons, User Assistance and Outreach (UAO) analysts, Scientific Computing Group (SciComp) liaisons, data liaisons, and visualization liaisons;
- multiple vehicles to communicate with users, sponsors, and vendors;
- developing and delivering training to current and future users; and
- strong outreach to interface with the next generation of HPC users, the external media, and the public.

1.4.1 User Support

The OLCF recognizes that users of HPC facilities have ranging needs requiring diverse solutions, from immediate, short-term, trouble ticket–oriented support, such as assistance with debugging and optimizing code, to more in-depth support requiring total immersion in and collaboration on projects. The facility provides complementary user support vehicles that include user assistance and outreach staff;

liaisons in respective scientific, data, and visualization areas; and computer scientists who assist on issues surrounding the programming environments and tools. The following sections detail some of the high-level support activities during CY 2016 and the specific OLCF staff resources available to assist users.

1.4.2 User Assistance and Outreach (UAO)

The UAO team addresses user queries, acts as user advocates, covers front-line ticket triage, resolution, and escalation, provides user communications, develops and delivers training and documentation, and installs third-party applications for use on the computational and data resources. The team also manages the OLCF Resource and Allocation Tracking System (RATS), which is the authoritative source for most of the system, user, and project data at the OLCF.

1.4.2.1 Wraprun

Wraprun is a tool developed by the OLCF UAO to assist Titan users with running ensemble jobs that consist of multiple independent tasks run simultaneously. Historically, ensemble jobs have been difficult to run successfully on Titan because they can cause system instability issues when launched with Cray's Aprun job launcher alone. Wraprun eliminates these difficulties. Many projects have made extensive use of the tool. The UAO group has worked closely with these teams to incorporate their feedback and requests into new features. In the past year Wraprun was enhanced to handle MPI-IO tasks specifically for projects wishing to bundle applications using NetCDF. The ability to customize the way Wraprun manages the output from bundled applications was also requested and delivered. Internal feedback and data from the OLCF operations group has also been used to rapidly implement improvements to maximize system stability. The general user experience was improved by adding a feature that permits ensembles to be defined through configuration files in addition to its original command line interface. Wraprun was also given a Python application programming interface (API), allowing users to generate Wraprun ensembles in Python-based workflow management systems.

1.4.2.2 Software

The OLCF User Assistance (UA) group routinely installs software that is not provided by the vendor as part of the programming environment. This effort accounts for many hours of UA time each year, as the OLCF provides multiple production and test-bed resources for use. To make this process more efficient, UA adopted Spack, a community-based scientific package management tool. Spack is primarily developed by Lawrence Livermore National Laboratory (LLNL), with collaborators at the National Energy Research Scientific Computing Center (NERSC), NASA, and EPFL in Switzerland. Switching to a tool that is in use at other HPC centers allows the OLCF to leverage best practices, "software recipes," and solve problems in a common way so that complex scientific software stacks are delivered quicker and in a repeatable fashion. Additionally, the OLCF began migrating to the industry-standard practice of "Continuous Delivery," which provides greater transparency and control over the software stacks available on each of our HPC resources. Automating the delivery of scientific software to researchers helps reduce the delivery time of popular packages, and it helps provide a consistent user experience across a wide range of HPC systems.

1.4.2.3 Resource and Allocation Tracking System (RATS)

The UA team continued to make feature improvements to the OLCF's in-house customer relationship management software, known as RATS, in CY 2016. A total of 13 versions (3.19.0–3.27.1) were deployed to production with minimal downtime. New feature highlights include an auto-email system, project application integrations, new API filtering for reporting, and the addition of fine-grained export control regulations.

To improve the consistency and quality of OLCF communications to end users, the team deployed features to send emails automatically from within RATS. The system can automatically send six distinct emails to PIs or users based on the center's business requirements (e.g., when the end of a certain project is near). Emails can be previewed in RATS, triggered manually if desired, and toggled on and off on a per-project basis. The system also stores a searchable record of all sent emails for future reference. This new feature reduced the amount of time the OLCF accounts team spent crafting, reviewing, and sending hundreds of manual emails each year. It also helped to reduce common errors such as typos and incorrect dates in emails.

The team also deployed integrations between RATS and the popular WordPress plugin called GravityForms, which allows data from form submissions on the OLCF website (e.g., the OLCF Project Application form) to be synced into RATS on a regular, automatic schedule. Synced applications can be viewed, edited, and processed by OLCF accounts staff using several helpful in-app tools. This has reduced the need for previously labor-intensive data entry. Work continues in automating the ingestion of data from other publically accessible OLCF forms.

RATS has a full-featured REST API used for report data gathering and programmatic information dissemination throughout the division. The team added comprehensive filtering features to most of these API endpoints, allowing more complex queries to run against the data with ease. For example, historically complex questions like "which DD projects were enabled at any point in December of 2016?" can be answered quickly and accurately.

Finally, existing "security classifications" in RATS (e.g., category 2) were proving too broad for some use cases. The team added the ability to also associate individual Department of Commerce, DOE, and Department of State regulations to RATS projects, and in turn, to associate prohibited countries with each regulation. Prohibited country information can then be compared with a user's country affiliations, and appropriate warnings are displayed in-app when business needs require it. These new features greatly reduce the chance of adding a user to a project that they are not authorized to access.

1.4.3 Scientific Liaisons

In addition to UAO staff, the OLCF provides certain projects with scientific (1.4.3), data (1.4.4), and visualization (1.4.5) liaisons. These liaisons are a unique OLCF response to HPC problems faced by users. The OLCF actively partners with project PIs and their teams to assist in delivering science outcomes by improving operational scaling and efficiencies, as well as providing computational domain scientific support. The OLCF Scientific Computing group provides certain projects with scientific liaisons. Examples of the impact of scientific liaison collaboration with and support of users are provided in Sections 1.4.3.1, 1.4.3.2, 1.4.3.3.

1.4.3.1 Advancing Models for Multiphase Flow in Porous Medium Systems

The objective of the INCITE project "Advancing Models for Multiphase Flow in Porous Medium Systems," led by James E. McClure is to develop an improved multi-scale description for two-fluid flow in rock and other geologic materials. Lattice Boltzmann simulations are performed within experimentally observed digital rock images to quantify essential aspects of two-fluid flow processes, providing the basis to advance multi-scale theory.

The team's research has broad implications for several research areas, including geologic carbon sequestration, oil and gas recovery, contaminant transport, and fluid flow processes within the critical zone. By building a computational framework and tools for simulating two-fluid flow in porous media, a better understanding of the flow that routinely occurs in natural and engineered geologic systems is obtained. Massively parallel GPU-based simulations of two-fluid flow have been used to directly simulate macroscopically relevant phenomena based on digital rock images. New constitutive laws have been developed to characterize the topologic arrangement of fluids within complex geometric materials. These relationships have been shown to resolve critical deficiencies associated with traditional two-fluid flow models that have been used for decades to describe reservoir-scale transport in geologic systems. Working with the OLCF, McClure developed simulation tools that can model two-phase flow through real-world porous media using data obtained from a 3-D micro-computed tomography (micro-CT) image. With this approach, the behavior of the interfaces has been studied, resulting in a better understanding of the flow and a best paper award from the Society of Core Analysts.

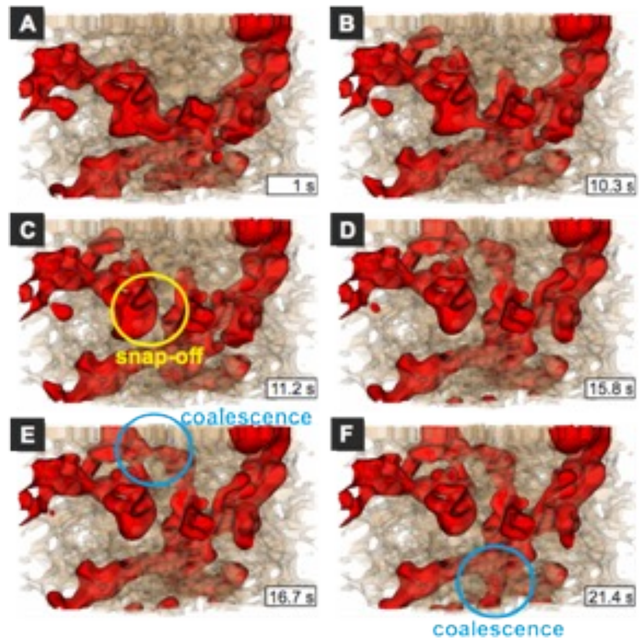

Figure 1.2. Simulations performed on Titan demonstrate that fluid phase configurations change continuously during displacement, as nonwetting phase regions snap-off from and coalesce with each other, rather than flowing in static connected pathways as previously assumed. Properly accounting for topological changes is shown to be important to predict the macroscale flow behavior.

OLCF computational scientist Mark Berrill used guidance from multiscale averaging to develop an in situ analysis framework to continuously analyze the simulation state while it is in memory, taking advantage of idle CPU-cores within each compute node allocated to the simulation. This approach was applied in full-scale production simulations to extract previously inaccessible information tracking the fluid connectivity within the simulation. Ongoing efforts to automate and streamline the process of converting the 3-D micro-CT image data to a 3-D mesh for the simulation will allow future simulations to be conducted in real time with the experimental measurements allowing real-time feedback and comparisons.

1.4.3.2 Large-Scale Coupled-Cluster Calculations of Supramolecular Wires

The INCITE project "Large-Scale Coupled-Cluster Calculations of Supramolecular Wires" led by Poul Jørgensen of the University of Aarhus is focused on the development of linear scaling highly correlated electronic structure methods within the Divide-Expand-Consolidate (DEC) framework of the LSDALTON software.

OLCF computational chemist Dmitry Liakh was instrumental in developing new capabilities in the LSDALTON code as well as resolving multiple outstanding code issues. Liakh coauthored the ScaTeLib library, part of the LSDALTON code that provides basic numeric infrastructure for performing tensor operations in parallel. The computationally intensive portions of the LSDALTON code, namely the coupled-cluster and second-order perturbation theory modules, rely on this library. Furthermore, the library was used to implement the so-called left-transformation coupled-cluster equations necessary for evaluating the wave function response. Since the corresponding equations are solved within the DEC approximation, the electronic response can now be computed for much larger chemical systems. That is, with the help from the OLCF, the LSDALTON code has advanced beyond the ground-state Schrödinger problem (at the coupled-cluster level) and acquired a principal capability to compute analytic energy derivatives and excited states for large chemical systems with controlled error. Additionally, Liakh also implemented a procedure for redistributing tensors in the coupled-cluster module of LSDALTON, which is used to adjust the distributed tensor storage layout between the CCSD and CCSD(T) modules of LSDALTON (different levels of coupled-cluster theory). Finally, Liakh was instrumental in enabling the 2016 Gordon Bell submission based on the LSDALTON DEC-RI-MP2 code that scaled to the full size of Titan. Most notably, Liakh wrote a page counting function that computes the amount of free memory available on the node that is covered by huge pages. This function was crucial for mitigating the severe memory fragmentation problems experienced by LSDALTON on Titan (other memory intensive codes with a dynamic memory allocation pattern will likely experience the memory fragmentation problem as well). By using this function, the LSDALTON code could finally account for the amount of usable memory in terms of huge pages, providing a true free memory estimate and avoiding unexpected crashes during execution.

1.4.3.3 New Capabilities in Nuclear Physics

PI James P. Vary from Iowa State University leads an INCITE project named “Nuclear Structure and Nuclear Reactions,” which describes atomic nuclei from first principles using a collection of quantum many-many methods. These methods are complementary, and each method has its strength in a different area of the nuclear chart. Led by ORNL researcher Gaute Hagen, the NUclear Coupled Cluster Oak Ridge (NUCCOR) group, including OLCF computational scientist Gustav R. Jansen, is spearheading efforts to compute properties of atomic nuclei within the medium-mass range, a range previously only reachable by more phenomenological models. Using model-independent two- and three-body forces derived from quantum chromodynamics, the fundamental field-theory that governs the dynamics of quarks and gluons, the team has gradually increased the size of nuclei that can be described using this approach. This has only been made possible by the simultaneous efforts and accomplishments in nuclear theory, many-body methods, and HPC (Figure 1.3). Working as the interface between nuclear theory and HPC, Jansen made significant contributions by being deeply embedded as an INCITE scientific liaison and one of the lead developers of NUCCOR. He has been responsible for developing many of the computational kernels that are optimized for Titan, especially those related to three-body forces. He has also developed several physics modules and introduced new many-body methods in nuclear physics that greatly increase the reach of first principles calculations. Because of these efforts, the current state-of-the art in nuclear structure has been moved from isotopes of oxygen with twenty-something particles to isotopes of calcium and nickel with 54 and 78 particles, respectively. The expected near-term exponential growth of the number of particles that can be included in first principles calculations holds great promise for nuclear physics research.

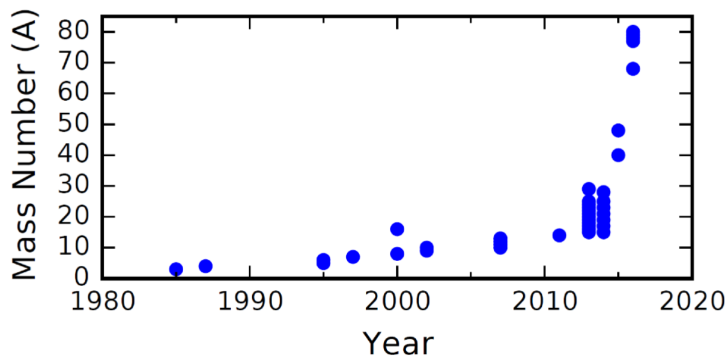

Figure 1.3. The trend of realistic first principles computations of atomic nuclei. The recent advance is due to efficient many-body methods and increasing computational power.

1.4.4 Data Liaisons

Large-scale simulations that execute on leadership-class supercomputers work at such high speeds and resolution that they generate unprecedented amounts of data. The size of these datasets—ranging from a few gigabytes to hundreds of terabytes—makes managing and analyzing the resulting information a challenge.

1.4.4.1 ADIOS Improves TOTAL Workflow

One of the important migration applications in seismic imaging is Reverse Time Migration (RTM). As part of a cooperative research and development agreement with the OLCF, TOTAL performed scaling experiments using RTM on Titan. The application is a hybrid version of RTM that was accelerated using directive-based programming for GPUs. The application generates intermediate data that is written during the "forward phase" of the application and is later read back during the "backward phase." The amount of intermediate data generated by the application is a concern at large scale, as is discussed below.

The dataset processed at large scale is the Society of Exploration Geophysicists Advanced Modeling Program Isotropic Phase I dataset. It consists of 2,793 seismic "shots." Shots can be processed independently of each other. Each shot is processed using 18 MPI processes. One MPI process is spawned per node. Thus, 18 nodes work on one shot. The application writes approximately 9 TB of intermediate data per shot. This data is removed once the shot is processed successfully. A single shot takes approximately 3.5 h to complete on Titan. At full scale, the team used 18,000 GPUs to process 1,000 shots at a time (18,000 GPUs/18 GPUs per shot). These large-scale runs can generate approximately 9 PB of intermediate data on the file system.

The application originally used the write and read POSIX I/O function calls for managing all I/O. Initial scaling experiments to test the performance of the application at a smaller scale showed that the application runtime was dominated by I/O. For example, an experiment to process 500 shots simultaneously in 10 h only managed to complete 175 shots in the allotted time because of the performance of the application's I/O component. Detailed I/O profiling using the Darshan tool revealed that most of the writes were in the range of 10–100 KB size, which is very small and caused too much overhead.

As reported in the 2015 OAR, TOTAL conducted initial I/O studies and showed advantages of using the Adaptable I/O System (ADIOS) over other I/O solutions. ADIOS was added to the new GPU version of RTM in CY 2016. Preliminary experiments on Titan with a single shot showed improvement in the I/O pattern; most writes were in the 10–100 MB size range. The next set of large tests plan to use ADIOS in 2017.

1.4.4.2 In Situ Multi-Scale Visual Analytics for Transformative Extreme Scale Science

OLCF Advanced Data and workflows (ADW) group members Benjamin Hernandez and Jamison Daniel received the best poster award during a Laboratory Directed Research and Development poster session organized by ORNL in September 2016. The title of the poster was "In Situ Multi-Scale Visual Analytics for Transformative Extreme Scale Science," which was led by ORNL Computational Sciences and

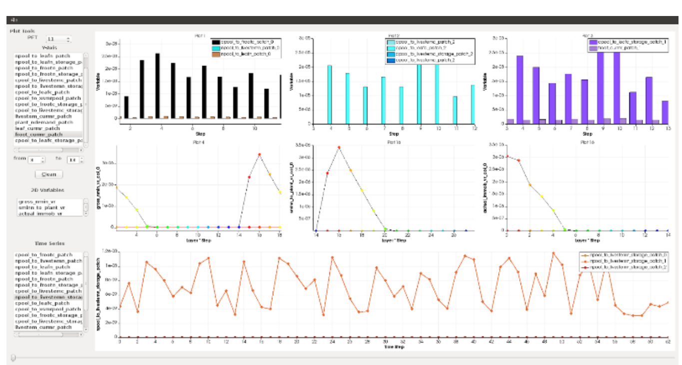

Figure 1.4. Data visualization engine for "In Situ Multi-Scale Visual Analytics for Transformative Extreme Scale Science."

Engineering Division's Chad A. Steed. Benjamin Hernandez extended the ACME Land Model by leveraging an in situ solution. This solution provides a way for researchers to analyze results in near real time, resulting in less data management and far less time to solution as data does not need to be staged in or out before or after analysis. Figure 1.4 shows the coupled data visualization engine that was leveraged for this in situ analysis. This approach successfully used the rendering capabilities of Titan's GPUs and the display capabilities of the Exploratory Visualization Environment for Research in Science and Technology (EVEREST) to view the interactive and running ACME Land Model simulations. This work also resulted in a publication and presentation at the International Congress on Environmental Modelling and Software, in Toulouse, France, on July 10–14, 2016.

1.4.5 Visualization Liaisons

1.4.5.1 Computational Seismology Visualization

Ongoing support for reading and visualization of the ADIOS output files from SPECFEM3D_GLOBE in VisIt assisted Jeroen Tromp's computational seismology project efforts. This support included updating VisIt for ADIOS releases, adding better support for different map projections, and including seismological hotspots.

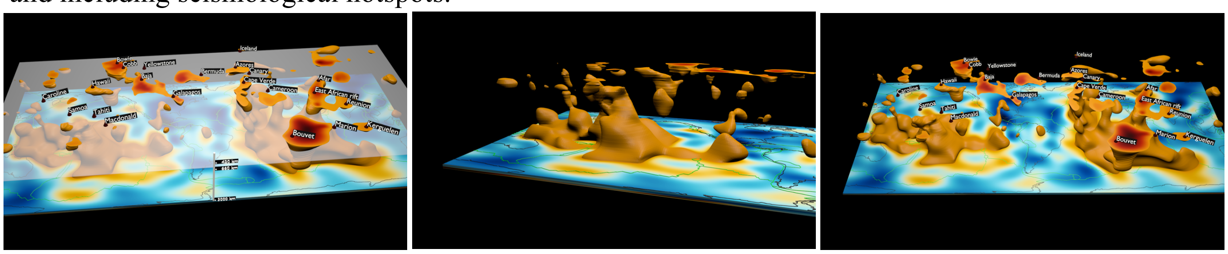

Figure 1.5. Images highlighting the 3-D plume structures, which were shown at the Visualization Showcase at Supercomputing 2016.

The extensive work in CY 2016 with the Tromp team on the analysis and visualization of their first-generation model for global adjoint tomography focused on the visualization and analysis of plume features in key sections of the earth, including the Pacific, Africa, and North America regions. These visualizations helped guide researchers in understanding the structure of plumes originating in the core-mantle boundary and their relationship with seismic hotspots, tectonic plate boundaries, and other geological features (Figure 1.4). This work was reported in the following journal article:

Related publication:

E. Bozdag, D. Peter, M. Lefebvre, D. Komatitsch, J. Tromp, J. Hill, N. Podhorszki, D. Pugmire, "Global Adjoint Tomography: First-Generation Model." *Geophysical Journal International*, 207 (3): 1739–1766 (2016).

1.4.5.2 Fusion Visualizations

Support for XGC1 simulations has focused on several different areas. First, the OLCF continued support for production visualization and analysis using the VisIt parallel visualization tool. In 2016, work included updates to support the latest version of the ADIOS file format, improved support for the visualization of field data, and support for the inclusion of particle data in coordinate systems. These changes have provided easier means to perform analysis that require the flow field of the plasma.

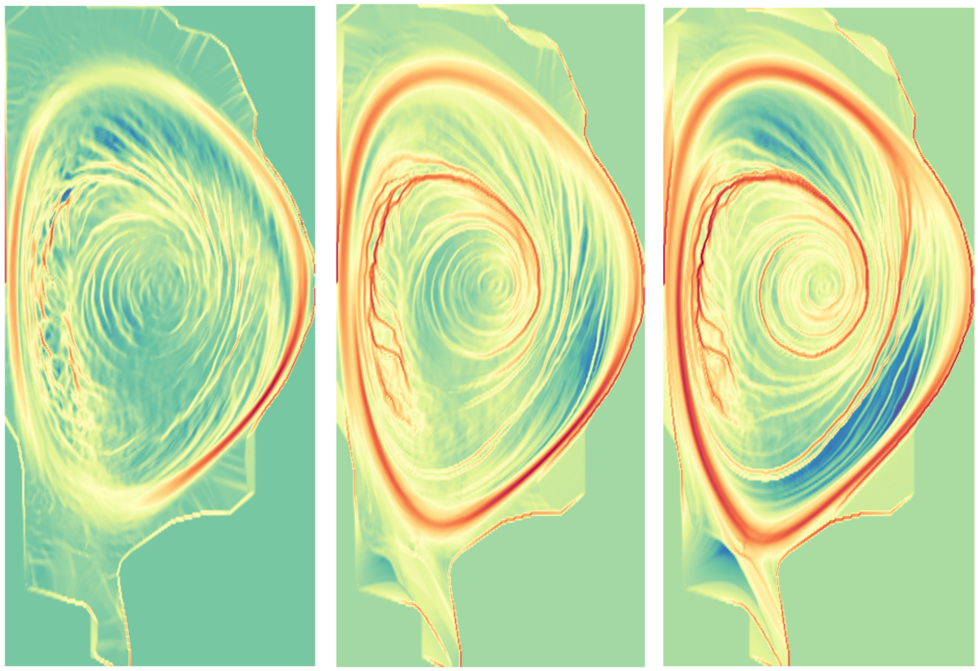

Figure 1.6. In situ slices of a Finite Time Lyapunov Exponents plot using the effective bulk plasma particle velocity vector field. The red values correspond to areas where the flow tends to separate, and blue is where the flow stays together.

Close collaboration with individual XGC1 team members encouraged a better understanding into the workflows, requirements, expectations, and expected usage of visualization. The team members were selected from a wide group of users, including physicists, analysts, performance engineers, computer scientists, and application team members. The goal of the study was to gain a holistic view of how data from a complex simulation are used. Of particular interest was common task identification and in situ processing candidates. Interviews were conducted, results were analyzed, and insights were included in a paper published at a workshop held in conjunction with Supercomputing 2016.

Ongoing support for visualization and processing of the ADIOS output files from XGC1 continued in VisIt. This support includes updating VisIt for ADIOS releases, adding better support for visualization of field and particle data in different coordinate systems. In an effort to understand the bulk behavior of particles in fusion plasmas, the team used the ADIOS data staging mechanism to analyze extremely large numbers of particles and used data reduction techniques to derive approximating fields for bulk velocity. Figure 1.5 was generated using a technique called Finite Time Lyapunov Exponents, which is a technique for the classification of structures within a flow field. This work was recently published in the *Supercomputing Frontiers International* journal.

Related publications:

Kress, J. Churchill, R. M., Klasky, S., Kim, M., Childs, H. and Pugmire, D. "Preparing for In Situ Processing on Upcoming Leading-Edge Supercomputers." *Supercomputing Frontiers International* 3 (4). (2016). doi: 10.14529/jsfi160404.

Kress, J., Pugmire, D., Klasky, S., and Childs, H. "Visualization and Analysis Requirements for In Situ Processing for a Large-Scale Fusion Simulation Code," ISAV 2016, Proceedings of the 2nd Workshop on In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization, Supercomputing 2016, Salt Lake City, UT, Nov. 13, 2016.

1.4.6 OLCF User Group and Executive Board

The OLCF User Group (OUG) consists of all users and meets once a month by conference call and webinar to discuss OLCF news, resources, policies, and timely HPC tutorials and techniques. The OUG Executive Board, consisting of 11 members, meets shortly before or after the monthly call to provide the OLCF with in-depth feedback and guidance on topics such as training, facility resources, and policies. For the 2016–2017 cycle, five users ran for the three open spots on the board. Eighty-four users voted to elect Balint Joo, Joseph Oefelein, and Yifeng Cui to fill the open positions. The current board chair is Katrin Heitmann. The board elected Balint Joo to serve as the vice chair this year, and he will automatically become the chair for the 2017–2018 cycle.

Three members of the board—Hai Ah Nam, David Dixon and Mike Zingale—participated in the OLCF 2016 User Meeting Planning Committee. They were instrumental in helping the OLCF identify and invite speakers to the meeting. They also helped plan the OLCF User Group business meeting, where the OLCF discussed new tools and services that were implemented and solicited feedback on further user needs. Users expressed appreciation for the recently implemented streamlined Globus authentication process for data transfer and interest in training on visualization and data analytics focused topics.

In response to the interest in visualization training, user Mike Zingale gave a talk called "Introduction to Visualization with YT on Titan and Rhea" at the July OUG webinar. He shared his experience with implementing a Python and YT based visualization workflow for his INCITE project. This webinar garnered 35 attendees and is posted on the OLCF training channel, (https://vimeo.com/channels/olcftraining). The video received 24 plays through CY 2016. In keeping with last year's feedback from the user group that the OLCF should provide opportunities for users to share HPC knowledge with each other, user Peter Vincent shared his Python-based linear solver, PyFR, on Titan, which drew 39 attendees for the October webinar. Similarly, user Luigi Capone shared a seminar titled "Portability in Industrial Codes to GPU Hardware: The TITAN Example," which resulted in 30 attendees at the August OUG webinar. For the other monthly webinars, OLCF staff gave short tutorials or seminars on OLCF tools and resources. The OLCF hosted a total of 11 monthly conference calls in 2016 with a total of 803 attendees. In addition, the OLCF hosted the annual User Meeting, which had 151 attendees. More information about the OLCF User Group, including a list of the Executive Board members, can be found at https://www.olcf.ornl.gov/about-olcf/oug/.

1.4.7 Training, Education, and Workshops

Workshops, user conference calls, training events, and seminars are integral components of both user assistance and outreach. Training can obviate difficulties in performing science on such large-scale systems, and training events can serve to engage the public and the user community. In addition to training users to use the resources available at the OLCF, the training program focused on an additional area: software development best practices. The OLCF facilitated 11 user conference calls, four week-long hackathons, one mini-hackathon, and several other workshops and seminars in 2016.

See Appendix B for a complete summary of these events. A few of the notable 2016 events are highlighted in sections 1.4.7.1, 1.4.7.2, 1.4.7.3, and 1.4.7.4.

1.4.7.1 The Interoperable Design of Extreme-Scale Application Software (IDEAS) Project, Argonne, National Energy Research Scientific Computing Center (NERSC), and Oak Ridge Jointly Offer Best Practices Webinars for Software Development

The OLCF, Argonne Leadership Computing Facility (ALCF), NERSC, and the Interoperable Design of Extreme-Scale Application Software (IDEAS) project partnered to produce a series of webinars in 2016—"Best Practices for HPC Software Developers"—to help users of HPC systems carry out their software development more productively. A total of seven webinars were presented in 2016. The IDEAS project focuses on increasing software sustainability and developer productivity for high-performance applications in computational science. IDEAS is a collaboration among researchers at seven DOE laboratories—Argonne National Laboratory (ANL), LLNL, Los Alamos National Laboratory (LANL), ORNL, Pacific Northwest National Laboratory, and Sandia National Laboratories (SNL)—with the Colorado School of Mines. One challenge IDEAS researchers are addressing is the development of an extreme-scale scientific software development kit. The extreme-scale scientific software development kit brings together many widely used DOE numerical libraries into a common configuration and build environment to simplify usage and interoperability for software developers who rely on these libraries. Another element is methodologies for software productivity, which researchers use to adapt, develop, and deploy best practices—drawing on approaches in the software engineering community and translating them into resources that are useful for the HPC computational science community. The IDEAS project partnered with the ASCR facilities to deliver the webinar series, which was presented through a teleconferencing service to encourage participation from the audience. Attendees had opportunities to ask questions, and several sessions included interactive demonstrations. The sessions were recorded and posted online.

1.4.7.2 2016 Hackathons

In 2016, the OLCF and partners expanded efforts to train individuals and organizations from academia, industry, and government in GPU programming for current and next-generation HPC systems.

Twenty-six teams—more than any previous year—attended four hackathon events split between the United States and Europe. Hackathons were held at the Technische Universitat Dresden in collaboration with Forschungszentrum Julich research center in Germany, the University of Delaware, the Swiss National Supercomputing Centre in Switzerland, and the OLCF. The 2016 events totaled more than 200 attendees and attracted a wide range of experienced and new scientific users from domains as varied as biology, materials science, and cyber security. Hackathon teams representing universities such as Duke, Georgia Tech, and Delaware highlighted increased participation from student programmers, and teams from St. Jude's Children Hospital, the National Cancer Institute, and the National Institutes of Health represented growing interest from the medical research community in leadership-class computing.

As in past years, each hackathon gave participants the opportunity to work directly with HPC mentors and vendors, develop new programming skills, and improve the quality of their scientific results in an intensely collaborative environment. The week-long events also provide new and veteran programmers an opportunity to familiarize themselves with cutting-edge, diversified architectures and develop marketable skills.

Previous hackathons are beginning to yield publications as well as new projects to the OLCF. Teams who attended the hackathons have also applied for Center for Accelerated Application Readiness (CAAR), ALCC, and INCITE programs. One particular team, a collaboration of Nek5000 and NekCEM, received an R&D100 award and credit the work at the hackathons contributing to the award. Teams report continued work on their applications well after the event, and much of the hackathon-derivative work has resulted in production-quality applications.

1.4.7.3 2016 OLCF User Meeting

One hundred fifty-one people attended the 2016 OLCF User Meeting, making it the largest user meeting to date. Users and administrators from the OLCF participated in the 3-day event May 24–26 to learn about, share, and discuss the most recent science conducted on the OLCF's Titan supercomputer. The first two days of the meeting focused on the finer points of HPC at the OLCF through user, facility, and how-to talks that covered a wide breadth of topics. Users presented their research results, and staff informed users about current and future facility developments. On the first day, Paul Messina, project director for the Exascale Computing Project (ECP) and senior strategic advisor at the ALCF, gave the keynote talk, "The Path to Capable Exascale Computing," in which he provided an update on the project. Jacqueline Chen, distinguished member of the technical staff at the Combustion Research Facility at SNL, gave the closing keynote address, “Towards Exascale Simulation of Combustion Science and Technology.” Other notable user talks came from the University of Utah’s Martin Berzins, LANL’s Rajan Gupta, and OLCF research scientist Yi Wang. Sessions on the third day were dedicated to the future of HPC and included talks and a panel focused on task-based programming. Four experts, including Berzins, gave presentations on different programming models such as Legion, Uintah, MADNESS, and ParSEC. The OLCF improved the roundtable lunch discussions by pairing the event with a pre-roundtable “slam,” during which each of the OLCF staff members facilitating conversations was given a chance to briefly introduce the topic that his or her individual table would discuss. The attendees then chose which table to sit at for their working lunch. The various topics were selected based on user input during registration. The meeting also coincided with the election of three new OUG board members. As a companion to the user meeting, the OLCF held a day of tutorials with two tracks. One track covered data analytics and visualization tools, and the other focused on programming environments and OLCF tools. Here are a few comments collected at the conclusion of the User Meeting:

“At the 2016 user meeting, I became keenly aware of the commitment, knowledge, and zeal of the staff and gained a deep appreciation for their dedicated efforts to provide a smoothly functioning system that I take for granted as a remote user.”

“This conference has provided a great opportunity for computational fluid dynamics modelers to meet many HPC experts at the OLCF and scientists from different fields.”

1.4.7.4 OLCF Mini-Hackathon

In November, the OLCF hosted its first 3-day mini-hackathon, an extension of the center’s 5-day GPU hackathon program, which began in 2014. The mini-hackathon took place November 1–3 and drew 41 attendees. This mini-hackathon began as previous 5-day hackathons were routinely oversubscribed, limiting attendance. Increased demand made way for an abbreviated hackathon. Unlike the full hackathon, which typically draws teams with some experience using GPUs, this year’s mini-hackathon was geared toward participants with little to no experience accelerating applications for GPUs. Some users worked on their own applications during the event, but others who did not have an application got a feel for what they may need to do to program software to use the GPUs. Fernanda Foertter of the OLCF partnered with Jeff Larkin and Brent Leback of NVIDIA to lead the event’s lectures and hands-on sessions. Instruction focused on programming techniques using OpenACC, an application programming interface designed for accelerated architectures. The survey responses received after the event showed that users were satisfied with the event overall, and based on this feedback, the OLCF plans to host both the 5-day hackathon and the mini-hackathon again next year. Below is one example of feedback received after the event:

“Thank you very much for your presentations, coordination, and help during the OLCF mini-hackathon. This is the type of event that is very useful to HPC users at all levels: new users, advanced users, and those looking for the most up-to-date information. With the OpenACC standard, the use of GPUs and parallel programming in general have become so much more user friendly compared to their predecessors or alternative options. I was very impressed with the quality of the presentations, the responses of the professional staff in answering our questions, and with the OLCF help team. I know hackathon days are a “marathon” and, despite that, you and the help team were always funny, friendly, and ready to help. I found all the presentation material to be very helpful, but of course I need some time to assimilate and put into practice all that I learned at the event.”

1.4.8 Training and Outreach Activities for the Future Members of the HPC Community and the General Public

The OLCF maintains a broad program of collaborations, internships, and fellowships for young researchers. The OLCF brought in 29 students during Summer 2016, whose backgrounds ranged from computer architecture, to mathematics and statistics, to artificial intelligence. Internships at the OLCF provide students with research and learning opportunities as they prepare for careers as scientists, programmers, and engineers. Interns work closely with their staff mentors to perform high-priority research at the center that contributes to its mission. Projects ranged from helping the center prepare for the arrival of Summit, automating report generation, and enhancing the flexibility and performance of the OLCF's data analysis environments

The student interns included Yufang Bao, Aaron Barlow, Drew Cage, Arghya Chatterjee, Swapnil Desai, Oumar Diallo, Jeffrey Graves, Preston Gull, Rachel Harken, Joseph Huber, Radhika Katti, Benjamin Klein, Mark Mudrick, Christopher Muzyn, Sarah Neuwirth, Miki Nolin, Woong Shin, Hyogi Sim, Rina Singh, Michaela Vaitova, Duane Vick, Alexander Walker, James Wynne III, Peter Xenopoulos, Wei Xie, Sisi Xiong, Luna Xu, Lechen Yu, and Shiqi Zhong.

In addition, the center led or participated in several events that introduced the OLCF to the next generation of HPC users and staff. Examples of these events are described below.

1.4.8.1 OLCF Staff Give Students Access to Supercomputers at Diversity Conference

OLCF staff presented for the first time at the 2016 Association for Computing Machinery's Richard Tapia Celebration of Diversity in Computing conference, which aims to bring students and professionals together to promote diversity and create networking opportunities in computational domains. The conference—held this year in Austin, Texas, from September 14 to 17—was open to students, drawing undergraduates and graduate students as well as professionals from academia and industry. The conference is among the most diverse in computing, with the attendees being primarily from underrepresented groups. OLCF staff participated in the event's career fair and led an HPC session that gave attendees temporary access to Metis, a two-cabinet Cray XK7 system with the same architecture as Titan. One of the most important parts of the event this year was the new "HPC Day," which included HPC sessions led by staff from multiple national laboratories. OLCF user support specialists Adam Simpson and Veronica Vergara Larrea led a GPU-accelerated computing hands-on session, during which students became familiar with OpenACC, a programming standard for parallel computing. Simpson gave the students a highly serial code sample and walked them through OpenACC's directives—instructions that tell the program what to do to parallelize a code and offload it to the GPU. Once he explained the directives, he gave the students an exercise and helped them determine where they should include each directive. Approximately 1,000 people attended the 2016 Tapia conference.

1.4.8.2 International Summer School on Atmospheric and Oceanic Sciences (ISSAOS) Meeting Focuses on Supercomputing Concepts

This year the International Summer School on Atmospheric and Oceanic Sciences (ISSAOS) in L'Aquila, Italy, featured a HPC theme for the first time to introduce earth and climate science students from around the globe to supercomputing concepts. Valentine Anantharaj, OLCF computational scientist, contributed to the event's organization by planning HPC topics, delivering lectures, and leading workshops. Since its beginning in 2000, ISSAOS has featured themes such as aerosols and climate change, atmospheric data assimilation, and chaos in geophysical flows. People from various institutions and backgrounds in the climate sciences come to the weeklong summer school to engage in lectures taught by computational scientists representing institutions in North America and Europe. This year's summer school, "Advanced Programming Techniques for the Earth System Science," informed climate science PhD students about computer architectures, HPC programming techniques, and software tools that may benefit their research. A total of 45 students attended the 2016 ISSAOS. During the HPC overview lecture on the history of the field and the world's top computers, Anantharaj emphasized the critical role of HPC services, which are often behind the scenes, that are provided via user assistance and HPC operations. On the second day, he gave a lecture on how to parallelize code using MPI, a standard communication protocol for programming with parallel computers. The lecture was followed by a hands-on session in which students employed techniques discussed during the lecture. Anantharaj's hands-on MPI session gave students access to Metis, a scaled-down version of the OLCF's leadership-class hybrid Cray XK7, Titan.

1.4.8.3 Computational Science Graduate Fellowship (CSGF) Annual Program Review Brings Staff and Students Together

For the sixth year in a row, OLCF staff networked with promising graduate fellows and introduced them to research opportunities during DOE's Computational Science Graduate Fellowship (CSGF) Annual Program Review. The meeting—which took place this year from July 25–28 in Arlington, Virginia—provided CSGF fellows with an outlet for research discussion and one-on-one meetings with DOE laboratory staff. The event also offers interactive HPC training sessions that give first-year fellows a proper introduction to the fundamentals of supercomputing and advanced training sessions on more complicated topics like parallel algorithm design and code optimization. This year a multi-laboratory component of the workshop included presentations from staff members of two new participating organizations—ANL and the Energy Sciences Network—in addition to staff members from ORNL and the National Energy Research Scientific Computing Center, who participated last year. Before the review Veronica Vergara Larrea and Adam Simpson, user support specialists at the OLCF, led a pre-session webinar about basic aspects of HPC, from accessing the system to submitting and running parallel jobs. During the review, they also led a session on GPU-accelerated computing that provided fellows with hands-on experience using OpenACC, a programming standard for parallel computing, as well as access to Metis, a two-cabinet Cray XK7 system with the same architecture as Titan. After students became acquainted with OpenACC, Vergara Larrea and Simpson introduced a smoothed-particle hydrodynamics code, which simulates the behavior of fluids. The students then had the opportunity to use OpenACC to accelerate the smoothed-particle hydrodynamics code using the tools they learned about during the session. The 2016 workshops included four sessions covering intermediate and advanced levels of supercomputing, an improvement from previous years' workshops that focused mainly on introductory topics. In general, the fellows who participated were already knowledgeable about computational science as applied to their particular scientific domains. Consequently, the workshops aimed to expand the fellows' knowledge of parallel computing challenges and opportunities that are applicable across disciplines.

1.4.8.4 Introduce Your Daughter to Code

In July, the OLCF partnered with ORNL's Women in Computing (WiC) Group and hosted an event titled "Introduce Your Daughter to Code," during which girls between the ages of 10–16 ran code on OLCF's supercomputers. The event provided daughters of staff in ORNL's Computing and Computational Sciences Directorate the opportunity to learn about code from programmers in several ORNL departments, including the OLCF. During the first hour, participants used their ages and the numbers of letters in their names to change the specifications for visualizations they created using the program fractalName. Designed by Susheela Singh, an OLCF intern from North Carolina State University, fractalName uses input data to generate images in the form of colorful fractals, repeating geometric patterns that make up larger shapes. The girls displayed their fractals on the visualization wall in the EVEREST facility. The girls and their parents also ran an example application, called Birthday Pi, on Titan. The Birthday Pi program, designed by Katie Schuman, a Liane Russell Distinguished Early Career Fellow, enabled the girls to find their birthday sequences within the first 100,000 digits of pi. The coding experience gave the girls a glimpse into computing, and it gave their parents the opportunity to share insights into their careers.

1.4.8.5 OLCF Staff Participate in Bay Area Maker Faire

In 2006, *Make:* magazine created an event aimed at bringing together its do it yourself–inspired readers to share homemade innovations ranging from life-sized robots to powerful small-scale computers. Such events, dubbed Maker Faires, have grown exponentially, with small events regularly happening in countless locations, and annual meetings taking place in large cities from California to Singapore.

No strangers to science and technology innovations, DOE employees converged on the San Francisco Bay Area for the 22nd annual Bay Area Maker Faire, which ran May 20–22. OLCF staff contributed by taking a modified version of Tiny Titan—a mini parallel computer made of several Raspberry Pi computers that helps educate students about the principles of parallel computing—to inspire a new generation of HPC enthusiasts. Typically, OLCF staff members use an Xbox controller to control the fluid-dynamics simulation running on Tiny Titan. Tiny Titan provides the ability to show students how adding or subtracting processors can affect a simulation. For the Maker Faire, the OLCF recoded the program to work with Leap Motion, which uses a camera and several infrared sensors, allowing students to control the ball in the Tiny Titan simulation with their hands. The demonstration was a hit with both the parents and the kids that visited the DOE exhibit.

1.4.8.6 Tiny Titan

In addition to the Maker Faire, Tiny Titan once again represented the OLCF at regional schools, the USA Science and Engineering Festival in Washington, DC, and the DOE National Science Bowl. The OLCF also received additional reports of other institutions building their own Tiny Titan, including one at Jefferson Laboratory. Tiny Titan was also used again in the overlook area of the OLCF where tours were conducted regularly. The tool continues to help make Titan and the power of supercomputers more understandable to the numerous visitors to the facility each year.

1.4.9 Outreach

The OLCF Outreach team works to engage new and next-generation users and showcases OLCF research through strategic communication activities including tours, highlights, fact sheets, posters, snapshots, the OLCF website, and center publications (Appendix C). The Outreach team was responsible for the creation of 74 highlights—including science, technology, and people features—and for 179 total outreach products in 2016. In addition, the OLCF provides tours to groups of visitors who range from middle-school students to senior-level government officials throughout the year. The center conducted 251 tours of the facility in 2016.

In 2016, the Outreach team continued its initiative to be more intentional and proactive in its amplification methods via social media and other platforms. The net result was that Outreach science highlights were picked up by media outlets—including science journals, trade publications, and, in a few cases, the regular press—a total of 262 times in 2016. OLCF Outreach science highlights were published by *Science Magazine, Science Daily*, *Phys.Org*, *HPCWire*, and other outlets. Outreach highlights were also published on the DOE Office of Science home page and other DOE outlets over 25 times in CY 2016. In all, the Outreach team produced a total of 24 science highlights in 2016. Those highlights touched on a variety of science domains, from geoscience ("A Seismic Mapping Milestone"), to engineering ("Streamlining Accelerated Computing for Industry"), to materials ("The Shape of Melting in Two Dimensions"). In addition, the team completed 50 people and technology feature stories.

The story of science on Titan was the theme for the 2015–2016 OLCF annual report. The report, produced in 2016, carried the theme of achieving great science throughout the volume. In addition to presenting accomplishments from the previous year, the annual report pointed to the future with an update on the completion of Summit. OLCF graphic designer Jason Smith was recognized by the American Inhouse Design Awards Competition for the design of the 2014–2015 OLCF Annual Report, receiving a

“Certificate of Excellence” award in 2016. The report design also received an Award of Merit in the 41st annual “V Awards” competition from the Public Relations Society of America.

The Outreach team continues to leverage social media and online news services to broaden the reach of OLCF stories within the HPC community and among the general public. In October, the team carried out a social media campaign celebrating the fourth anniversary of the Titan supercomputer’s operation. The campaign, called “TitanWeek,” garnered around 90,000 impressions on Twitter and attracted participation from more than 100 organizations and individuals on Twitter and Facebook, including vendors, national labs, and the Secretary of Energy.

The OLCF website received 299,569 page views and a mean survey rating of 4.4/5.0 in 2016. The most visited pages in 2016 were the Titan resource overview page, the OLCF homepage, and the Titan User Guide. Overall, the OLCF user guides were the most highly rated aspect of the website with a mean rating of 4.4/5.0, and 94% of users indicated that they were satisfied or highly satisfied with the user guide. The user guide received 24,887 page views and 20,687 unique page views in 2016.

1.5 LOOKING FORWARD

1.5.1 Application Portability

The portability of scientific and engineering applications is increasingly important to the users of ASCR computing facilities, and it continues to be one of the application development focus areas. Application developers target a wide range of architectures, and because applications have much longer lifespans than computer architecture, they need to be developed for changing architectures. In addition, many of the OLCF’s PIs have allocations at multiple computing facilities and having portable applications greatly facilitates their science campaigns.

Recognizing the responsibility to make applications both architecturally and performance portable among the ASCR computing facilities, Tjerk Straatsma, leader for the OLCF SciComp group; Katie Antypas, NERSC Scientific Computing and Data Services Department Head; and Timothy Williams, ALCF Principal Project Specialist, are working on an initiative to coordinate application readiness activities and developing a strategy to provide guidance and tools encouraging application development that is portable across different architectures. As part of this effort, the three facilities are collaborating on a joint project to assess the impact of portable programming approaches on performance using three small but full applications. The computing intensive parts are being implemented using different models. For the OLCF, computational scientist Arnold Tharrington and postdoctoral associate Ada Sedova are working on the implementation of nonbonded force routines in molecular dynamics simulations using the ARGOS application. This work will be enabled by the OLCF purchase of a single cabinet Cray XC40 system that mirrors the architecture of the ALCF’s Theta system.

Ada Sedova joined the Scientific Computing group in August 2016. She earned her PhD from the Department of Biomedical Sciences of University at Albany for work in computational molecular biophysics, biophysical chemistry, and structural biology, and her masters of arts degree in mathematics from the Department of Mathematics and Statistics at the same institution. Following her doctoral studies, she worked with DNA electrochemistry and biosensor development for 1 year, including experimental electrochemistry and computational fluid dynamics (CFD) simulations, at the University at Albany’s Department of Chemistry. Ada is working on the molecular dynamics portion of the portability project.

1.5.2 Application Readiness and Early Science

OLCF’s CAAR is a partnership of the SciComp group, scientific application teams, vendor partners, and tools developers with the goal of readying a set of applications for the Summit architecture. The suite of CAAR applications covers a broad range of scientific disciplines and employs a range of programming models and software designs. In addition to developing highly efficient implementations for Summit, the teams are working toward performance portability across different architectures to serve the user base.

The applications that are part of the CAAR program are summarized in Table 1.7. More details can be found on the CAAR webpage at https://www.olcf.ornl.gov/caar.

Table 1.7. Applications in the Center for Accelerated Application Readiness (CAAR)

Application	Principal investigator	CAAR liaison	Scientific discipline
ACME	David Bader Lawrence Livermore National Laboratory	Matthew Norman	Climate science
DIRAC	Prof. Lucas Visscher Free University of Amsterdam	Dmitry Liakh	Relativistic chemistry
FLASH	Bronson Messer Oak Ridge National Laboratory	Bronson Messer	Astrophysics
GTC	Zhihong Lin University of California–Irvine	Wayne Joubert	Plasma physics
HACC	Salman Habib Argonne National Laboratory	Bronson Messer	Cosmology
LS-DALTON	Prof. Poul Jørgensen Aarhus University	Dmitry Liakh	Chemistry
NAMD	Prof. Klaus Schulten University of Illinois–Urbana- Champaign	*vacant*	Biophysics
NUCCOR	Gaute Hagen Oak Ridge National Laboratory	Gustav Jansen	Nuclear physics
NWCHEM	Karol Kowalski Pacific Northwest National Laboratory	Dmitry Liakh	Chemistry
QMCPACK	Paul Kent Oak Ridge National Laboratory	Ying Way Li	Materials science
RAPTOR	Joseph Oefelein Sandia National Laboratories	Ramanan Sankaran	Combustion
SPECFEM	Prof. Jeroen Tromp Princeton University	Judy Hill	Seismology
XGC	C. S. Chang Princeton Plasma Physics Laboratory	Ed D'Azevedo	Plasma physics

1.5.3 Computational Scientists for Energy, the Environment, and National Security (CSEEN) Postdoctoral Program

DOE recognizes the need to train and retain computational scientists in a broad range of disciplines that support DOE and the nation's critical mission needs to maintain the US competitive advantage in high-performance and data-intensive scientific computing. Considering the ever-increasing capability of high-end computer architectures, there is a continued and increased need to ensure a well-trained computational science workforce in academia and industry and at the national laboratories. In recognition of this need, DOE proposed that ASCR establish a postdoctoral training program at its user facilities, including the OLCF, ALCF, and NERSC, for future Computational Scientists for Energy, the Environment, and National Security (CSEEN). The objectives of this program are (1) to help ensure an adequate supply of scientists and engineers who are appropriately trained to meet national workforce needs, including those of DOE, for high-end computational science and engineering, with skills relevant to both exascale and data-intensive computing; (2) to make ASCR facilities available, through limited-term appointments, for applied work on authentic problems with highly productive work teams and increasingly cross-disciplinary training; and (3) to raise the visibility of careers in computational science and engineering to build the next generation of leaders in computational science.

The OLCF CSEEN Postdoctoral program seeks to provide opportunities to bridge the experience gap between the need to address domain science challenges and the need to develop high-performance software development expertise. One of the focus areas is to provide the skills required to port, develop, and use software suites on the leadership computing resources at the OLCF. The software development activities occur in conjunction with a CAAR project. This model offers the greatest potential for scientific breakthroughs through computing and provides ample opportunity to publish in domain scientific literature. This approach will ensure the postdoctoral trainees continue to build their reputations in their chosen science communities. Participants in the CSEEN Postdoctoral program are encouraged to attend tutorials, training workshops, and training courses on select computer science topics. One of the most important outcomes for the postdoctoral trainee is the opportunity to publish and present research accomplishments. In 2016, the CSEEN Postdoctoral program at the OLCF supported seven trainees:

Stephen Abbott joined the SciComp group in November 2015. He obtained a PhD in physics in September 2015 from the University of New Hampshire, where he studied reconnecting magnetic instabilities in fusion plasmas. He will assist in the development of and conduct research with particle-based plasma models, particularly the XGC gyrokinetic particle-in-cell simulation code that is being prepared for Summit under the CAAR.

Yangkang Chen joined the SciComp group in April 2016, received his PhD degree in geophysics from the University of Texas at Austin, where he worked with Prof. Sergey Fomel on developing computational methods for processing massive seismic data that are used to create high-resolution images of the subsurface oil and gas reservoirs. Yangkang's focus at ORNL is to develop workflows for efficient massively parallel implementation of seismological methodologies on current state-of-the-art accelerated computer architectures, and he will be responsible for applying the workflows in solving large scientific challenge problems such as inverting the global geological structure of the earth, using the SpecFEM CAAR application.

Amelia Fitzsimmons joined the group in March 2016. She has a background in computational chemistry and is working with the DIRAC code as part of the CAAR project. Her current research interests include studying interactions between heavy elements that require relativistic computational treatment and biological and inorganic systems.

Kalyana Gottiparthi joined SciComp in September 2015. He earned a Bachelor of Technology degree in aerospace engineering from Indian Institute of Technology–Kharagpur in 2007 and received a PhD in aerospace engineering from Georgia Institute of Technology in 2015. During his graduate studies, he investigated complex multiphase turbulent post-detonation and post-explosion flows using HPC resources. He also developed and implemented massively parallel routines for multiphase flows. He is a recipient of the Institute Silver Medal and J. C. Ghosh Memorial prize from the Indian Institute of Technology–Kharagpur. As a distinguished Postdoctoral Associate, he will be performing high-fidelity simulations using the Raptor large eddy simulation (LES) code that is one of the CAAR applications. The code will be improved for scalability on Titan, and his simulations will showcase the physics and the performance developments in the code.

Tom Papatheodore joined the SciComp group in August 2015. Just before this appointment, he received a PhD from the University of Tennessee (UT), where he studied computational astrophysics. His research within SciComp focuses on supporting a CAAR project intended to port the astrophysics simulation code FLASH to Summit. He has subsequently joined NVIDIA.

Micah Schuster joined SciComp in August 2015 and is working on the NUCCOR CAAR project. He finished a PhD in computational science from Claremont Graduate University in 2015. His dissertation focused on nuclear structure physics. Over the course of his research, he traveled to LLNL for 3 years to work with members of the Computational Nuclear Physics Group and the Institute for Scientific Computing Research. During that time, he added new physics to existing codes and modified them to run more efficiently on the Sierra cluster at LLNL.

Andreas Tillack joined the group in October 2016. He received his PhD in chemistry from the University of Washington. Andreas holds a master's degree in physics from Humboldt University of Berlin (Germany). He is working with the QMCPACK CAAR team on materials science applications.

Business Results

HIGH PERFORMANCE COMPUTING FACILITY
2016 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2017

2. BUSINESS RESULTS

CHARGE QUESTION 2: Is the facility maximizing the use of its HPC systems and other resources consistent with its mission?

OLCF RESPONSE: Yes. The OLCF provides a series of highly capable and reliable systems for the user community. The 2016 reporting period includes full CY production periods for the following HPC resources: the Cray XK7 (Titan), the Cray XC30 (Eos), the Lustre file systems (Spider II), and the archival storage system (HPSS). The effectiveness of these resources is demonstrated by the business result metrics, which were met or exceeded in all cases. The OLCF team managed policies and job-scheduling priorities that maximized access to these production systems. In 2016, the OLCF once again delivered all of the compute hours committed to the three major allocation programs: INCITE, ALCC, and DD. OLCF computational and data resources remain critical to scientific research of production simulations across many scientific domains.

2.1 BUSINESS RESULTS SUMMARY

Business results measure the performance of the OLCF against a series of operational parameters. The two operational metrics relevant to the OLCF's business results are resource availability and the capability utilization of the HPC resources. The OLCF additionally describes resource utilization as a reported number, not a metric.

2.2 CRAY XK7 (TITAN) RESOURCE SUMMARY

The OLCF upgraded the existing Cray Jaguar from a model XT5 to a model XK7, releasing it to production on May 31, 2013. The resulting system contains 18,688 NVIDIA K20X (Kepler) accelerators, in which each existing AMD Opteron connects to an NVIDIA Kepler to form a CPU-GPU pair. The completed XK7 system, which has more than 27 petaflops of peak computational capacity, is named Titan.

2.3 CRAY XC30 (EOS) RESOURCE SUMMARY

Eos is a four-cabinet Cray XC30. The system has 736 Intel Xeon E5-2670 compute nodes and 47.6 TB of memory and provides the OLCF user community with a substantive large-memory-per-node computing platform. The Eos nodes are connected by Cray's Aries interconnect in a network topology called Dragonfly. All INCITE users are automatically granted access to the XC30.

2.4 LUSTRE FILE SYSTEMS (SPIDER II) RESOURCE SUMMARY

In October 2013, the OLCF released Spider II, its next-generation Lustre parallel file system, to production. Spider II contains two instantiations of the /atlas file system, with an aggregate capacity of more than 30 PB and block-level performance of more than 1.3 TB/s. The Spider II file system is the default high-performance parallel file system for all compute resources.

2.5 DATA ANALYSIS AND VISUALIZATION CLUSTER (RHEA) RESOURCE SUMMARY

Rhea is a 512-node large memory data analytics Linux cluster. The primary purpose of Rhea is to provide a conduit for large-scale scientific discovery through pre- and post-processing of simulation data generated on Titan. Users with accounts on INCITE- or ALCC-supported projects are automatically given accounts on Rhea. DD projects may also request access to Rhea. Each of Rhea's nodes contains two 8-core 2.0 GHz Intel Xeon processors with hyperthreading and 128 GB of main memory (upgraded in 2015 from 64 GB). Rhea offers nine additional heterogeneous nodes, each of which boasts 1 TB of main memory and 2 NVIDIA Tesla K80 (Kepler GK210) GPUs. Rhea is connected to the OLCF's 30+ PB high-performance Lustre file system, Spider II.

2.6 HIGH PERFORMANCE STORAGE SYSTEM (HPSS) RESOURCE SUMMARY

The OLCF provides a long-term storage archive system based on the HPSS software product co-developed by IBM, LANL, SNL, LLNL, Lawrence Berkeley National Laboratory (LBNL), and ORNL. The ORNL HPSS instance is currently over 60 PB in size and provides up to 200 Gb/s of read and write performance. The archive has ingested over 225 TB in a single day several times in the last year; the previous daily maximum was just over 150 TB/day.

The archive is built from hardware from Dell, Hewlett Packard, Brocade, NetApp, DataDirect Networks, and Oracle. An 18 PB disk cache allows burst rates into the archive at up to 200 Gb/s; there is 26 Gb/s of read and write bandwidth to the archive via 120 Oracle T10K series tape drives. There are six Oracle SL8500 tape libraries for tape archival storage that each contain 10,100 slots; the archive's maximum capacity is over 500 PB using these libraries.

2.7 VISUALIZATION RESOURCE SUMMARY

The EVEREST facility has three computing systems and two separate state-of-the-art visualization display walls. The primary display wall spans 30.5 ft × 8.5 ft and consists of eighteen 1920 × 1080 stereoscopic Barco projection displays arranged in a 6 × 3 configuration. The secondary display wall contains sixteen 1920 × 1080 planar displays arranged in a 4 × 4 configuration, providing a standard 16:9 aspect ratio. The stereoscopic capabilities allow the user to experience binocular depth perception. An array of sequentially pulsed infrared LED cameras record the physical position and orientation of the user, and the resolution density provides an optimal solution for human visual acuity. These combined technologies, along with OLCF staff expertise, allow scientists to analyze complex scientific datasets in an immersive environment and communicate abstract concepts in an intuitive visual format.

2.8 OLCF COMPUTATIONAL AND DATA RESOURCE SUMMARY

The OLCF provided the Titan and Eos computational resources and the Spider II and HPSS data resources in 2016 (Table 2.1). Supporting systems such as EVEREST, Rhea, and data transfer nodes were also offered. Metrics for these supporting systems are not provided.

2.8.1 OLCF HPC Resource Production Schedule

The OLCF production computational systems entered production according to the schedule in Table 2.2. This includes historical data associated with the Cray XT5, the very small overlap in December 2011 beginning with the introduction of the Cray XK6, and the series of Cray XK systems available in 2012 and 2013.

Table 2.1. OLCF production computer systems, 2016

System	Access	Type	CPU	GPU	Computational description			
					Nodes	Node configuration	Memory configuration	Interconnect
Titan	Full production	Cray XK7	2.2 GHz AMD Opteron 6274 (16 core)	732 MHz NVIDIA K20X (Kepler)	18,688	16-core SMP[a] CPU + 14 SM[b] GPU (hosted)	32 GB DDR3-1600 and 6 GB GDDR5 per node; 598,016 GB DDR3 and 112,128 GB GDDR5 aggregate	Gemini (Torus)
Eos	Full production	Cray XC30	2.6 GHz Intel E5-2670 (8 core)	None	736	2 × 8-core SMP	64 GB DDR3—1,600 per node; 47,104 GB DDR3 aggregate	Aries (Dragonfly)

[a] SMP = symmetric multiprocessing
[b] SM = streaming multiprocessor

Table 2.2. OLCF HPC system production dates, 2008–present

System	Type	Production date[a]	Performance end date[b]	Notes
Spider II	Lustre parallel file system	October 3, 2013	—	Delivered as two separate file systems, /atlas1 and /atlas2. 30+ PB capacity
Eos	Cray XC30	October 3, 2013	—	Production with 736 Intel E5, 2,670 nodes.
Titan	Cray XK7	May 31, 2013	—	Production with 18,688 hybrid CPU-GPU nodes (AMD Opteron 6274/NVIDIA K20X)
JaguarPF	Cray XK6	September 18, 2012	October 7, 2012	Production at 240,000 cores until September 18, when partition size was reduced to 120,000 AMD Opteron cores. Additional Kepler installation. TitanDev access terminated
JaguarPF	Cray XK6	February 13, 2012	September 12, 2012	Full production until September 12, when partition size was reduced to 240,000 AMD Opteron cores. Beginning of Kepler installation
JaguarPF	Cray XK6	February 2, 2012	February 13, 2012	Stability test. Restricted user access. 299,008 AMD Opteron 6274 cores. Includes 960-node Fermi-equipped partition
JaguarPF	Cray XK6	January 5, 2012	February 1, 2012	Acceptance. No general access 299,008 AMD Opteron cores
JaguarPF	Cray XK6	December 12, 2011	January 4, 2012	142,848 AMD Opteron cores
JaguarPF	Cray XT5	October 17, 2011	December 11, 2011	117,120 AMD Opteron cores
JaguarPF	Cray XT5	October 10, 2011	October 16, 2011	162,240 AMD Opteron cores
JaguarPF	Cray XT5	September 25, 2009	October 9, 2011	224,256 AMD Opteron cores
JaguarPF	Cray XT5	August 19, 2008	July 28, 2009	151,000 AMD Opteron cores

[a] The production date used for computing statistics is either the initial production date or the production date of the last substantive upgrade to the computational resource.
[b] The performance end date is the last calendar day that user jobs were allowed to execute on that partition.

2.8.2 Business Results Snapshot

Business results are provided for the OLCF computational resources, the HPSS archive system, and the external Lustre file systems (Tables 2.3–2.6).

Table 2.3. OLCF business results summary for Titan

	Measurement	2015 target	2015 actual	2016 target	2016 actual
Cray XK7 (Titan)	Scheduled availability	95%	99.41%	95%	99.64%
	Overall availability	90%	97.01%	90%	97.03%
	MTTI[a] (h)	NAM[c]	326.86	NAM	473.52
	MTTF[b] (h)	NAM	1,088.67	NAM	1,750.73
	Total usage	NAM	90%	NAM	90%
	Core-hours used*	NAM	4,287,795,259	NAM	4,323,608,405
	Core-hours available	NAM	4,764,524,288	NAM	4,778,502,912
	Capability usage				
	INCITE projects	NAM	71.31%	NAM	62%
	All projects	35%	64.19%	35%	59.01%

[a] MTTI = Mean time to interrupt.
[b] MTTF = Mean time to failure.
[c] NAM = Not a metric. No defined metric or target exists for this system. Data provided as reference only.
* Does not include usage recorded during an outage.

Table 2.4. OLCF business results summary for Eos

	Measurement	2015 target	2015 actual	2016 target	2016 actual
Cray XC30 (Eos)	Scheduled availability	NAM[c]	99.67%	NAM	99.89%
	Overall availability	NAM	97.90%	NAM	97.97%
	MTTI[a] (h)	NAM	476.47	NAM	661.97
	MTTF[b] (h)	NAM	2,182.90	NAM	2,924.91

[a] MTTI = Mean time to interrupt.
[b] MTTF = Mean time to failure.
[c] NAM = Not a metric. No defined metric or target exists for this system. Data provided as reference only.

Table 2.5. OLCF business results summary for HPSS

	Measurement	2015 target	2015 actual	2016 target	2016 actual
HPSS	Scheduled availability	95%	99.94%	95%	99.89%
	Overall availability	90%	98%	90%	98.54%
	MTTI[a] (h)	NAM[c]	343.40	NAM	376.33
	MTTF[b] (h)	NAM	1,459.13	NAM	974.93

[a] MTTI = Mean time to interrupt.
[b] MTTF = Mean time to failure.
[c] NAM = Not a metric. No defined metric or target exists for this system. Data provided as reference only.

Table 2.6. OLCF business results summary for Spider II, the external Lustre file system

	Measurement	2015 target	2015 actual	2016 target	2016 actual
/atlas1	Scheduled availability	95%	99.77%	95%	99.92%
	Overall availability	90%	98.9%	90%	98.2%
	MTTI[a] (h)	NAM[c]	618.84	NAM	616.12
	MTTF[b] (h)	NAM	1,456.69	NAM	1,462.79
/atlas2	Scheduled availability	95%	99.76%	95%	99.92%
	Overall availability	90%	98.88%	90%	98.2%
	MTTI (h)	NAM	481.23	NAM	718.81
	MTTF (h)	NAM	1,092.44	NAM	2,194.18

[a] MTTI = Mean time to interrupt.
[b] MTTF = Mean time to failure.
[c] NAM = Not a metric. No defined metric or target exists for this system. Data provided as reference only.

For a period of 1 year following either system acceptance or a major system upgrade, the scheduled availability (SA) target for an HPC compute resource is at least 85%, and the overall availability (OA) target is at least 80%. For year 2, the SA target for an HPC compute resource increases to at least 90% and the OA target increases to at least 85%. For year 3 through the end of life for the associated compute resource, the SA target for an HPC compute resource increases to 95%, and the OA target increases to 90%. Consequently, SA targets are described as 85%/90%/95%, and OA targets are described as 80%/85%/90%.

For a period of 1 year following either system acceptance or a major system upgrade, the SA target for an external file system is at least 90%, and the OA target is at least 85%. For year 2 through the end of life of the asset, the SA target for an external file system increases to at least 95%, and the OA target increases to at least 90%. SA targets are thus described as 90%/95%. OA targets are thus described as 85%/90%.

The Spider II, Titan, and Eos systems all celebrated their 3 year production anniversaries in 2016. The reported results for each system measure are for all of CY 2016 and intentionally do not reflect the partial results to their respective production anniversaries. In all cases, the OLCF results exceeded the most stringent year 3 targets for the accompanying metrics.

Because an outage that could define the SA, OA, mean time to interrupt (MTTI), or mean time to failure (MTTF) may occur outside the reporting period, the data reflected here artificially assume calculation boundaries of 00:00 on January 1, 2016, and January 1, 2017.

2.9 RESOURCE AVAILABILITY

Details of the definitions and formulas describing SA, OA, MTTI, and MTTF are provided in Appendix D.

2.9.1 Scheduled Availability

The scheduled availability is described by equation (1). The OLCF has exceeded the SA targets for the facility's computational resources for 2015 and 2016 (Table 2.7).

$$SA = \left(\frac{time\ in\ period - time\ unavailable\ due\ to\ outages\ in\ period}{time\ in\ period - time\ unavailable\ due\ to\ scheduled\ outages\ in\ period} \right) * 100 \qquad (1)$$

Table 2.7. OLCF business results summary: Scheduled availability

	System	2015 target	2015 actual	2016 target	2016 actual
Scheduled availability	Cray XK7	95%	99.41%	95%	99.64%
	Cray XC30	NAM[a]	99.67%	NAM	99.89%
	HPSS	95%	99.94%	95%	99.89%
	/atlas1	95%	99.77%	95%	99.92%
	/atlas2	95%	99.76%	95%	99.92%

[a] NAM = Not a metric. No defined metric or target exists for this system. Data provided as reference only.

2.9.1.1 Assessing Impacts to Scheduled Availability

The operational posture for the Cray XK7 system contains a regularly scheduled weekly preventative maintenance period. Preventative maintenance is exercised only with the concurrence of the Cray hardware and software teams and the OLCF HPC Operations group. Typical preventative maintenance activities include software updates, application of field notices, and hardware maintenance to replace failed components. Without concurrence, the systems are maintained under normal operating conditions.

In 2016, OLCF staff executed scheduled maintenance on the Cray XK7 a total of 17 times, associated with hardware maintenance, software upgrades, field notice patching, Lustre software stack testing, Lustre upgrades, security patching, planned facility power maintenance, and retroactive hardware component replacements. Four unscheduled outages were reported in 2016, including a maintenance period that ran longer than anticipated, critical security patching, and a small number of hardware or software failures that could not be recovered. Similarly, OLCF performed scheduled maintenance on Eos 12 times in 2016, with two unscheduled outages.

2.9.2 Overall Availability

The overall availability of OLCF resources is derived using equation (2).

$$OA = \left(\frac{time\ in\ period - time\ unavailable\ due\ to\ outages\ in\ period}{time\ in\ period} \right) * 100 \tag{2}$$

As shown in Table 2.8, the OLCF exceeded the OA targets of the facility's resources for 2015 and 2016.

Table 2.8. OLCF business results summary: Overall availability

	System	2015 target	2015 actual	2016 target	2016 actual
Overall Availability	Cray XK7	90%	97.01%	90%	97.03%
	Cray XC30	NAM[a]	97.9%	NAM	97.97%
	HPSS	90%	98%	90%	98.54%
	/atlas1	90%	98.9%	90%	98.2%
	/atlas2	90%	98.88%	90%	98.2%

[a] NAM = Not a metric. No defined metric or target exists for this system. Data provided as reference only.

2.9.3 Mean Time to Interrupt (MTTI)

MTTI for OLCF resources is derived by equation (3), and a summary is shown in Table 2.9.

$$MTTI = \left(\frac{time\ in\ period - (duration\ of\ scheduled\ outages + duration\ of\ unscheduled\ outages)}{number\ of\ scheduled\ outages + number\ of\ unscheduled\ outages + 1} \right) \quad (3)$$

The MTTI summary is shown in Table 2.9.

Table 2.9. OLCF business results summary: Mean time to interrupt (MTTI)

	System	2015 actual	2016 actual
MTTI (h)	Cray XK7	326.86	473.52
	Cray XC30	476.47	661.97
	HPSS	343.40	376.33
	/atlas1	618.84	616.12
	/atlas2	481.23	718.81

MTTI is not a metric. The data is provided as reference only.

2.9.4 Mean Time to Failure (MTTF)

The MTTF is derived from equation (4), and a summary is provided in Table 2.10.

$$MTTF = \frac{time\ in\ period - (duration\ of\ unscheduled\ outages)}{number\ of\ unscheduled\ outages + 1} \quad (4)$$

Table 2.10. OLCF business results summary: Mean time to failure (MTTF)

	System	2015 actual	2016 actual
MTTF (h)	Cray XK7	1,088.67	1,750.73
	Cray XC30	2,182.9	2,924.91
	HPSS	1,459.13	974.93
	/atlas1	1,456.69	1,462.79
	/atlas2	1,092.44	2,194.18

MTTF is not a metric. The data is provided as reference only.

2.10 RESOURCE UTILIZATION

2016 Operational Assessment Guidance

The Facility reports Total System Utilization for each HPC computational system as agreed upon with the Program Manager. This is a reported number, not a metric.

The numbers that are reported for the Cray XK7 resource are Titan core-hours, which are comprised of 16 AMD Opteron core-hours and 14 NVIDIA Kepler SM-hours per Titan node-hour. The OLCF refers to the combination of these traditional core-hours and SM-hours as “Titan core-hours” to denote they are the product of a hybrid node architecture. System production requires the use of node-hours, which is an aggregate of all CPU and GPU resources comprising a single node. The use of node-hours impacts all scheduling and accounting activities. Users describe all job submission activity in node-hours as the smallest unit.

2.10.1 Resource Utilization Snapshot

For the Cray XK7 for the operational assessment period January 1–December 31, 2016, 4,323,608,405 Titan core-hours were used outside of outage periods from an available 4,778,502,912 Titan core-hours. The total system utilization for the Cray XK7 was 90%.

2.10.2 Total System Utilization

2016 Operational Assessment Guidance

The percent of time that the system's computational nodes run user jobs. No adjustment is made to exclude any user group, including staff and vendors.

$$SU = \left(\frac{\text{core hours used in period}}{\text{core hours available in period}} \right) * 100 \tag{5}$$

The measurement period is for 2016, irrespective of the prescribed allocation period of any single program. As an example, the INCITE allocation period follows a CY schedule. The ALCC program follows an allocation cycle that runs for 12 months, beginning July 1 of each year. System utilization for 2016 was 90%, which marks the fourth year that Titan has achieved 90% or higher utilization.

The OLCF tracks the consumption of Titan node-hours by job. By extension, this provides a method for tracking Titan core-hours by job. This method is extended to track the consumption of Titan core-hours by program, project, user, and system with high fidelity. Figure 2.1 summarizes the Cray XK7 utilization by month and by program for all of 2016. Figure 2.1 represents the three major OLCF user programs and does not include consumed core-hours from staff or vendor projects.

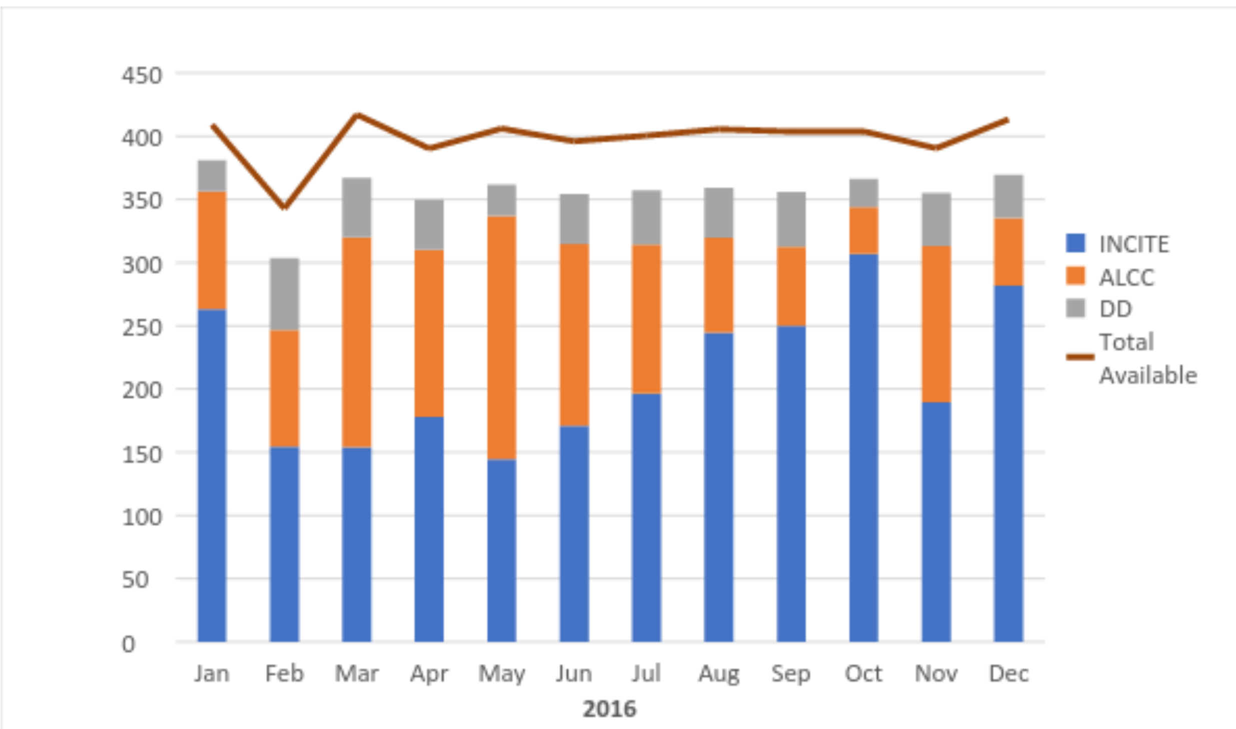

Figure 2.1. 2016 XK7 resource utilization—Titan core-hours by program.

2.10.2.1 Performance of the Allocated Programs

All allocation programs, including INCITE, ALCC, and DD, are aggressively monitored to ensure that projects within these allocation programs maintain appropriate consumption rates. The 2016 INCITE allocation program was the largest program in 2016, with a commitment for 2.25 billion Titan core-hours. The consumption of these allocation programs is shown in Table 2.11.

Non-renewed INCITE projects from 2016 continued running through January 2017 under the OLCF's 13th Month policy. This policy to permit an additional, final, month for completion was recognized as a best practice during a previous OAR review. It also serves to maintain high utilization while new projects establish a more predictable consumption routine. ALCC projects from the 2015 allocation period (ending June 30, 2016) were also granted extensions where appropriate.

Table 2.11. The 2016 allocated program performance on the OLCF resources

Program	Allocation	Hours consumed	Percent of total
INCITE[a]	2,250,000,000	2,675,682,679	60.42%
Titan	—	2,599,021,135	—
Eos	—	76,661,544	
ALCC	—	1,290,126,180	29.13%
ALCC_2016[b]	1,170,000,000	858,107,632	—
ALCC_2017[c]	986,500,000	432,018,548	—
DD	—	456,647,202	10.31%
NOAA[2]	87,500,000	5,813,966	0.13%
Total		4,428,270,027	100%

[a] Includes 13th month usage from January 2017
[b] Number of hours consumed in CY 2016 for the ALCC_2016 program
[c] Number of hours consumed in CY 2016 for the ALCC_2017 program

2.11 CAPABILITY UTILIZATION

Capability usage defines the minimum number of nodes allocated to a particular job on OLCF computing resources. To be classified as a capability job, any single job must use at least 20% of the available nodes of the largest system (Titan). The metric for capability utilization describes the aggregate number of node-hours delivered by capability jobs. The metric for CY 2016 was 35%, and this metric will remain until Titan is retired. The OLCF Resource Utilization Council uses queue policy on the Cray systems to support delivery of this metric target, prioritizing capability jobs with 24 h wall clock times in the queue.

The OLCF continues to exceed expectations for capability usage of its HPC resources (Table 2.12). Keys to the growth of capability usage include the liaison role provided by SciComp members, who work hand-in-hand with users to port, tune, and scale code; and the OLCF support of the application readiness efforts (i.e., CAAR), which actively engage with code developers to promote application portability, suitability to hybrid node systems, and performance. The OLCF aggressively prioritizes capability jobs in the scheduling system.

[2] NOAA = National Oceanic and Atmospheric Administration; period of performance reported: January 1, 2016–December 31, 2016.

Table 2.12. OLCF capability usage on the Cray XK7 system

Leadership usage	CY 2015 target	CY 2015 actual	CY 2016 target	CY 2016 actual
INCITE	NAM[a]	71.31%	NAM	62.00%
ALCC	NAM	60.34%	NAM	61.32%
All projects	35%	64.19%	35%	59.01%

[a] NAM = Not a metric. No defined metric or target exists for this system. Data provided as reference only.

The average consumption of hours by capability jobs, 59.01%, was once again well above the 2016 target of 35%. This consumption varies modestly during the year and is affected by factors including system availability and the progress of the various research projects. To promote the execution of capability jobs, the OLCF provides queue prioritization for all jobs that use 20% or more of the nodes and further boosts the very largest of these jobs, which use >60% (11,250) of the nodes, through aging boosts. The OLCF assesses job data in 10% "bins" to understand the job size distribution. Further, by assessing the aggregate bins, 20%–60% and >60%, the OLCF can assess the impact of queue policy on delivered node-hours.

Figure 2.2 shows the yearly average capability usage for each program, which describes the ratio of compute hours delivered by capability jobs to the compute hours delivered by non-capability jobs.

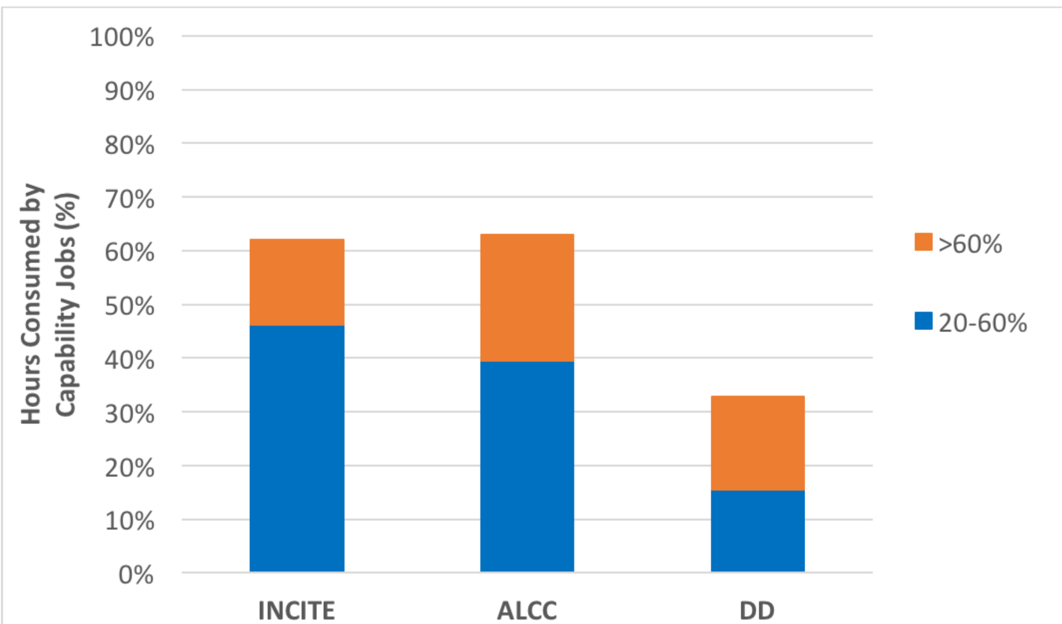

Figure 2.2. Capability usage by job size bins and project type.

2.12 GPU USAGE

Titan's heterogeneous architecture provides a key capability to users and allows them to exploit a hybrid compute node that contains both a CPU and the NVIDIA Kepler GPU. Hybrid nodes provide researchers with diverse architecture that is well suited for certain operations. As such, the use of this diverse architecture is optional and is exercised in different ways by research teams.

In 2016, the OLCF continued tracking GPU usage through Cray's Resource Utilization Reporting tool. Table 2.13 shows the GPU-enabled and CPU-only hours used and percentage breakdowns of each of the three primary allocation programs at the OLCF (INCITE, ALCC, and DD). As shown, the INCITE program uses the most GPU-enabled time on Titan. The INCITE and ALCC programs split CPU-only and GPU-enabled usage, totaling close to 50% for each. The DD program totaled 40% usage for GPU-enabled compute hours. When compared with CY 2015, the INCITE program produced roughly the same usage. INCITE CY 2015 used 60% GPU-enabled and 40% CPU-only hours, whereas the CY 2016 INCITE program used 54% GPU-enabled and 46% CPU-only hours. The ALCC program virtually splits time using both the GPUs and CPUs, with a slight edge toward CPU computing, similar to CY 2015. The DD program showed the least preference for GPU-enabled computing with an almost 3:2 CPU-to-GPU computing ratio. This ratio showed an improvement over the 3:1 ratio that was reported in the 2015 OLCF OAR. In general, these usage patterns match the expectations for each of the allocation programs. The INCITE computational readiness review criteria provide valuable insight into the proposed use of GPUs when allocating time and projects. The ALCC program does not require computational readiness reviews for GPU usage, and the DD program supports projects that may be in the beginning phases of porting code to GPUs. For most months in CY 2016, GPU-enabled INCITE applications were consistently responsible for more than half of the delivered hours to those projects and averaged 56.96% for all user programs in 2016.

Table 2.13. 2016 GPU-enabled and CPU-only usage by program

Program	Percentage	Hours
INCITE (GPU-enabled)	53.76	1,397,293,292
INCITE (CPU-only)	46.24	1,201,727,843
ALCC (GPU-enabled)	48.65	627,660,549
ALCC (CPU-only)	51.35	662,465,631
DD (GPU-enabled)	40.89	186,736,390
DD (CPU-only)	59.11	269,910,812

Approximately 51% of all delivered compute time on Titan was by GPU-enabled applications in CY 2016. Figure 2.3 shows the percentage of GPU-enabled compute time by month.

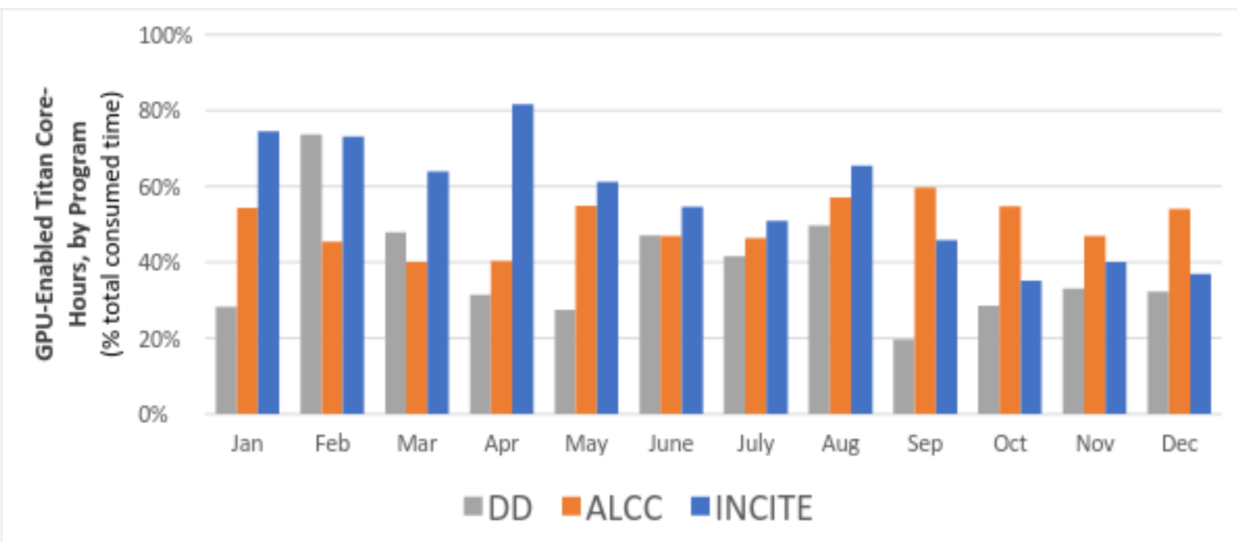

Figure 2.3. GPU-enabled percentage of compute time for the DD, ALCC, and INCITE user programs.

2.13 FUNCTIONAL PARTITIONING MULTI-TASKING FRAMEWORK ACCELERATES SCIENTIFIC DISCOVERY

OLCF staff previously developed and reported the functional partitioning (FP) runtime library to exploit GPU-accelerated architectures for nontraditional compute tasks such as post-processing or data analysis. Once integrated with the FP library, an application can launch data analysis tasks as FP processes on the unused cores. The FP runtime library was transitioned to production on Titan in 2016. This was accomplished through the development of a new suite of APIs, a shared memory transport to communicate information between the main application simulation, and the FP tasks and language-specific wrappers and bindings (e.g., for FORTRAN applications). OLCF staff Scott Atchley, Ross Miller, and Sudharshan S. Vazhkudai partnered with AMBER developer Pratul Agarwal to demonstrate the positive impact of the OLCF-developed FP library.

Since proteins and enzymes can sample many conformations, it is common to run large ensembles of individual molecular dynamics simulations to identify any new diversity. Typically, these simulations are performed "blindly." That is, they are run for a fixed number of time steps or a fixed amount of wall time, and then the results are analyzed to determine if new diversity was sampled. Typically, the post-processing tasks occur after the fixed length ensemble simulation completes and often results in inefficient use of allocated resources. By integrating FP calls into the AMBER molecular dynamics code, the team launched an ensemble simulation that performed the initial analysis while the simulation was still running.

The production ensemble used 1,000 nodes on Titan to simulate 1,000 samples of the enzyme dihydrofolate reductase. The individual simulations were configured to write their results every 1,000 time steps, and after each new set of results were written, the team used FP to execute a basic root-mean-square (RMS) deviation analysis on those results. The results were communicated from the application simulation to the FP task using the communication mechanism inherent in the FP library, which is set up during the FP initiation phase. The AMBER application performs most of the computation on the GPU, which leaves the CPU cores free for this in situ approach. The team selected a predefined cut-off value for the RMS deviation. Once reached, the sample did not need to be simulated any further.

Without FP, this ensemble would have been set up to simulate around 5 million time steps, which would have taken approximately 3 h. By leveraging FP this normally 3 h ensemble ended in less than 18 min.

Strategic Results

HIGH PERFORMANCE COMPUTING FACILITY
2016 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2017

3. STRATEGIC RESULTS

CHARGE QUESTION 3: *Are the processes for science engagement and outreach effective and do these processes enable scientific achievements consistent with the DOE strategic goals?*

OLCF RESPONSE: Yes. The OLCF continues to enable high-impact science results through access to the leadership-class systems and support resources. The allocation mechanisms are robust and effective.

OLCF projects and user programs are advancing DOE's mission to ensure America's security and prosperity by addressing its energy, environmental, and nuclear challenges through transformative science and technology solutions. The selected accomplishments described in this section serve to highlight how the OLCF is advancing two strategic objectives of DOE's Strategic Plan Goal 1, "Science and Energy: Advance foundational science, innovate energy technologies, and inform data driven policies that enhance U.S. economic growth and job creation, energy security, and environmental quality," as stated in the *US Department of Energy Strategic Plan: 2014–2018* (March 2014):

- Strategic Objective 2—Support a more economically competitive, environmentally responsible, secure and resilient US energy infrastructure.
- Strategic Objective 3—Deliver the scientific discoveries and major scientific tools that transform our understanding of nature and strengthen the connection between advances in fundamental science and technology innovation.

3.1 SCIENCE OUTPUT

2016 Operational Assessment Guidance

The Facility tracks and reports the number of refereed publications written annually based on using (at least in part) the Facility's resources. For the LCFs, tracking is done for a period of five years following the project's use of the Facility. This number may include publications in press or accepted, but not submitted or in preparation. This is a reported number, not a metric. In addition, the Facility may report other publications where appropriate.

3.1.1 OLCF Publications Report

In 2016, 405 publications resulting from the use of OLCF resources were published, as identified in the data collection completed on February 15, 2017, a 17% increase in publications yield for the year.[3,4] A list of 2015–2016 publications is available on the OLCF website.[5] In 2016, users and OLCF staff jointly

[3] In this document, "year" refers to the calendar year unless it carries the prefix "FY" indicating the fiscal year.
[4] 341 publications were reported in the 2015 OLCF OAR.
[5] https://www.olcf.ornl.gov/leadership-science/publications/

authored 53publications, a significant increase over 31 jointly authored publications reported for 2015. OLCF users published 312 without OLCF staff coauthorship, and OLCF staff published 40 without user coauthorship.

Sponsor guidance allows accepted and in press publications to be reported. However, the OLCF only tabulates publications appearing in print in the year under review (e.g., 2016). The OLCF continues to search for publications after the OAR is submitted to DOE each year. The number of publications reported in previous OARs will be updated annually in the current report. Table 3.1 includes the updated, verified, and validated publication count for the 2012–2016 period, showing continued growth in both total publication count and publications in journals with high impact factors, year-over-year.

Table 3.1 Summary of unique OLCF publications for 2012–2016

Year	Unique, confirmed OLCF publications	High-impact publications with JIF* > 7
2016	405	59
2015	347	41
2014	293	42
2013	357	30
2012	331	43

*JIF = Journal impact factor

3.2 SCIENTIFIC ACCOMPLISHMENTS

The OLCF advances DOE's science and engineering enterprise by fostering robust scientific engagement with its users through the INCITE liaison program, the user assistance program, and OLCF DD program outreach. The following subsections provide brief summaries of selected scientific and engineering accomplishments, as well as resources for obtaining additional information. While they cannot capture the full scope and scale of achievements enabled at the OLCF in 2016, these accomplishments advance the state-of-the-art in science and engineering R&D across diverse disciplines and are advancing DOE's science programs toward their targeted outcomes and mission goals. As an additional indication of the breadth of achievements, OLCF users published many breakthrough publications in high-impact journals in 2016, including one in *Nature*, one in *Chemical Reviews*, one in *Science*, four in *Nature Physics*, two in *Nature Climate Change*, two in *Nano Letters*, three in *ACS Nano*, five in the *Journal of the American Chemical Society*, one in *Nature Chemical Biology*, one in *Nature Geoscience*, one in *Nature Energy*, one in *Nature Plants*, one in *Advanced Functional Materials*, seven in *Nature Communications*, two in the *Proceedings of the National Academy of Sciences*, two in *Physical Review X*, and thirteen in *Physical Review Letters*. Altogether in 2016, OLCF users published 59 papers in journals with a journal impact factor greater than seven. Also of relevance is the large breadth of journal and conference titles in which OLCF publications appear in 2016. More than 180 unique journal and conference titles are represented in this list. Three PhD dissertations were published in 2016 that acknowledged use of OLCF resources. By all measures, 2016 set a new record for publication productivity by OLCF users and staff.

3.2.1 Streamlining Accelerated Computing for Industry: Peter Vincent, Imperial College, Director's Discretionary (DD) Program

Objective: To modernize industry-standard CFD and accurately resolve unsteady turbulent flow problems—such as low-pressure jet turbine cascades—at fine-grained scales through the development and deployment of PyFR, an open-source Python-based application designed by Peter Vincent, from Imperial College, to leverage accelerated computing architectures.

Impact: As supercomputers have increased in size and scale, CFD applications used in industries as diverse as aerospace, car manufacturing, and clean energy have struggled to keep pace, limiting their effectiveness in optimizing product design and supplanting physical testing. Furthermore, many CFD codes have not yet been adapted for accelerated computing architectures like Titan's.

The open-source PyFR application developed by Vincent's team combines the highly accurate flux reconstruction numerical scheme with a highly flexible, portable, and scalable code implementation. By scaling PyFR on a leadership-class accelerated supercomputer, Vincent's team is tackling unsteady airflow patterns in jet engines and providing engineers with an unprecedented tool to solve long-standing design problems while also demonstrating PyFR's potential for industry.

Accomplishments: On Titan, Vincent's team ran a simulation of five turbine blades, scaling the simulation up to 18,000 GPUs (Figure 3.1). The team's highest-performing run contained 195 billion degrees of freedom—or independent variables—and operated at a sustained speed of 13.7 petaflops, or 13.7 quadrillion calculations per second.

In recognition of PyFR's performance, Vincent's team was named a 2016 finalist for the Association of Computer Machinery's Gordon Bell Prize. The team will continue its investigation of low-pressure turbines in 2017 under a newly awarded INCITE allocation, capping a successful new-user outreach project in the OLCF's DD program.

OLCF science engagement: OLCF's Ben Hernandez created a visualization workflow to process and transmit simulation data in real-time between ORNL and Imperial College, speeding up the extraction of useful information.

Ninety-percent of Vincent's allocation used GPUs, with 37% of the team's simulation runs requiring more than 60% of Titan. Fifty-seven percent of the project used 20%–60% of system resources, while 6% of the work used less than 20% of the machine.

Figure 3.1. A simulation of flow over five jet engine low-pressure turbine blades. Credit: Vincent Lab, Imperial College.

Related publication:

P. Vincent, F. Witherden, B. Vermeire, J. Seok Park, and A. Iyer. "Towards Green Aviation with Python at Petascale." In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, p. 1. IEEE Press, 2016.

Online story: J. Hines, "Streamlining Accelerated Computing for Industry." *OLCF News*, August 23, 2016.

3.2.2 A Seismic Mapping Milestone: Jeroen Tromp, Princeton University, INCITE

Objective: To create a detailed 3-D picture of the earth's interior, imaging the entire globe from the surface to the core-mantle boundary. The team uses a combination of earthquake data and robust numerical methods to iteratively improve its world map of large-scale subsurface features.

Impact: The team's high-fidelity simulations add context to ongoing debates related to earth's geologic history and dynamics, bringing prominent features like tectonic plates, magma plumes, and hotspots into focus. Furthermore, promising early returns of the team's project demonstrate the value of adjoint-state tomographic methods for global-scale seismic modeling.

Accomplishments: Using Titan, Tromp and his collaborators completed their first-generation model, assimilating the seismic wave data recorded from 253 earthquakes and performing an iterative full-waveform inversion technique to produce a global map. The model is approaching continental-scale resolution, particularly in regions with dense data coverage.

Ongoing work dedicated to assimilating additional earthquake events into the model, automating workflow processes, and adjusting model parameters is enhancing the accuracy and scientific value of the team's model further.

OLCF science engagement: Collaboration with OLCF staff enabled the efficient use and accessibility of project data. Norbert Podhorszki worked with the team to create a new file format called Adaptable Seismic Data Format to improve data flexibility and movement on HPC resources. Judy Hill collaborated on a new automated workflow processes that holds the potential to drastically reduce the time it takes to perform a model update from a month to a few days. David Pugmire implemented in situ visualization tools, allowing team members to easily check their work remotely and avoid costly file transfers.

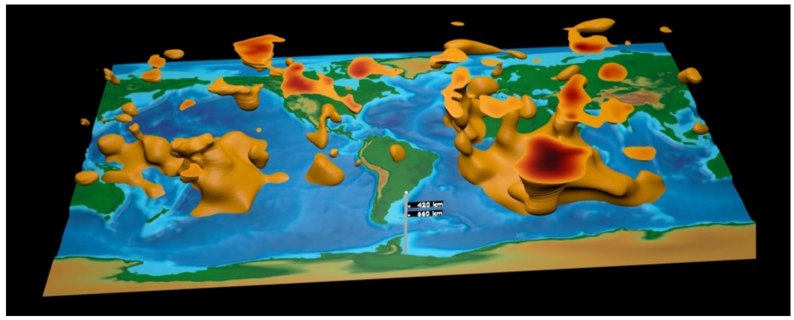

Figure 3.2. Simulations performed on Titan enabled the first global tomographic model of interior based on adjoint state methods. Credit: David Pugmire, ORNL.

Tromp's team began its 3-year INCITE project in 2015. In the second year of the project, the team utilized GPUs for 95% of its runs, with 48% of those simulations requiring more than 60% of Titan. Forty-two percent of the work used 20%–60% of the machine, and 10% of the project used 20% or less of system resources.

In the project's first year, the team utilized GPUs for 91% of its work, with 85% of those runs requiring more than 60% of Titan. Four percent of the project used 20%–60% of the machine, while 12% of runs used 20% or less of the system.

Related publications:

E. Bozdağ, D. Peter, M. Lefebvre, D. Komatitsch, J. Tromp, J. Hill, N. Podhorszki, and D. Pugmire. "Global Adjoint Tomography: First-Generation Model." *Geophysical Journal International* 207, no. 3: 1739–66 (2016). doi: 10.1093/gji/ggw356.

L. Krischer,J. Smith,W. Lei,M. Lefebvre,Y. Ruan, E. Sales de Andrade, N. Podhorszki, E. Bozdağ, and J. Tromp, "An Adaptable Seismic Data Format." *Geophysical Journal International* 207, no. 2: 1003–11 (2016). doi: 10.1093/gji/ggw319.

D. Komatitsch, Z. Xie, E. Bozdağ, E. Sales de Andrade, D. Peter, Q. Liu, and J. Tromp, "Anelastic Sensitivity Kernels with Parsimonious Storage for Adjoint Tomography and Full Waveform Inversion." *Geophysical Journal International* 206, no. 3: 1467–78 (2016). doi: 10.1093/gji/ggw224.

Online story: OLCF Staff Writer. "A Seismic Mapping Milestone." *OLCF News* (forthcoming).

3.2.3 The Shape of Melting in Two Dimensions: Sharon Glotzer, University of Michigan, INCITE

Objective: To understand the fundamental processes that influence phase transitions at the molecular level, with a focus on how particles assemble, grow, and repeat. The team uses scaled-up hard particle simulations to study emergent phenomena and the role of shape in self-assembly.

Impact: The capacity to predict and control crystallization in nanoparticle-based systems is a key capability sought in a wide range of industries, from pharmaceuticals to chocolate production. This interest extends to 2-D materials, such as thin films used in solar panels and batteries. Studying hard particle interactions in 2-D and 3-D could provide insight into these systems and the mechanism behind familiar, yet insufficiently understood, phase transitions, such as melting.

Accomplishments: Using Titan, Glotzer's team ran a series of hard particle simulations to study melting in 2-D systems. Specifically, the team explored how particle shape affects the physics of a 2-D solid-to-fluid melting transition.

The team simulated 11 different shape systems, ranging from triangles to 14-sided polygons, of up to 1 million particles. Each system was simulated at 21 different densities. The lowest densities represented a fluid state and the highest densities a solid state.

The researchers identified three distinct melting scenarios dependent on the shape of the systems' polygons. The most significant scenario emerged from hexagon systems, which Glotzer's team found to perfectly follow the phase transition described by a well-known theoretical model of 2-D melting called the KTHNY theory. In this scenario, the particles' shift from a solid state to an intermediate hexatic state and from a hexatic state to a fluid state in a perfect continuous phase transition process.

This work paves the way for Glotzer's team to tackle pressing phase transition problems in 3-D, such as how fluid particles crystallize into complex colloids—mixtures consisting of particles suspended throughout another substance.

OLCF science engagement: Access to Titan allowed the team to simulate 2-D hard polygon systems large enough to answer outstanding questions about the presence of a hexatic phase. Running its GPU-accelerated Monte Carlo code HOOMD-blue, the team greatly benefited from Titan's hybrid architecture, leveraging 64 GPUs for each individual hard particle system within large ensemble simulations necessary to achieve reasonable time to solution.

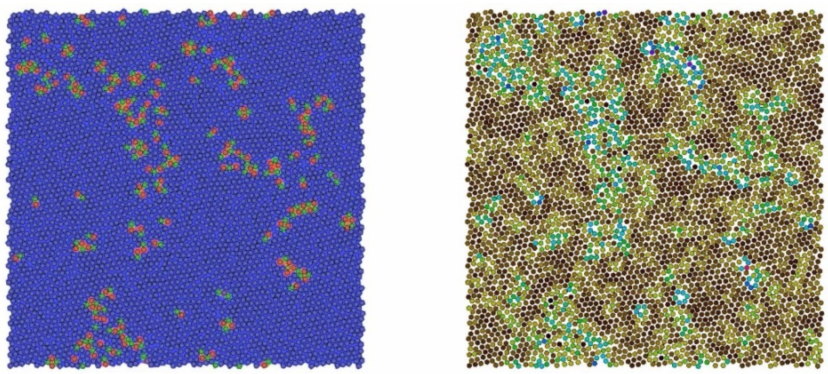

Figure 3.3. Representative snapshot depicting the hexatic phase transition of a 2-D hard particle system of hexagons and green particle pairs indicate defects in the system (*left*) and yellow particles depict how defects migrate (*right*). Credit: Joshua A. Anderson and James Antonaglia, University of Michigan.

Glotzer's team began its 2-year INCITE project in 2015. In the project's second year, 34% of simulations used less than 20% of Titan. Fifty-nine percent of runs used between 20%–60% of the machine, and 7% of the work used more than 60% of the system. In the project's first year, 43% of its simulations used less than 20% of Titan. Fifty-six percent of runs used between 20%–60% of the machine, and 1% of the work used more than 60%.

Related publication:

J. A. Anderson, J. Antonaglia, J. A. Millan, M. Engel, and S. C. Glotzer. "Shape and Symmetry Determine Two-Dimensional Melting Transitions of Hard Regular Polygons," December 26, 2016. arXiv: 1606.00687.

J. A. Anderson, M. E. Irrgang, and S. C. Glotzer. "Scalable Metropolis Monte Carlo for Simulation of Hard Shapes." *Computer Physics Communications* 204 pp. 21–30 (2016). doi: 10.1016/j.cpc.2016.02.02.

Online story: J. Hines. "The Shape of Melting in Two Dimensions" *OLCF News*, January 24, 2017.

3.2.4 Modeling Topographical Changes in Multiphase Flows: James McClure, Virginia Tech, INCITE

Objective: The team uses x-ray microtomography data in conjunction with GPU-based simulations to study multiphase flows in porous media in the earth's subsurface. By using GPU-accelerated computing, the team evaluates the macroscale relationships between fluid reservoirs and porous rock in the subsurface. Access to leadership-class computing resources also enables the team to advance the fundamental understanding of multiphase fluid transport in the earth's subsurface.

Impact: The team's research has broad implications for several research areas, including subsurface contamination remediation, oil recovery, and carbon sequestration. In addition to other experimental advances, building the computational framework for general subsurface flow phenomena will allow geophysicists to analyze interactions below the surface that would be impossible to measure or observe with other methods.

Accomplishments: In its multiyear INCITE award, the McClure team created a computational framework to study complex subsurface interactions. Building on that success, the team incorporated micro-computed tomography (micro-CT) imaging data into its simulations this year. By using micro-CT data, the team can directly visualize the movement of fluids in underground reservoir rocks and other geologic materials.

As a result, the team is directly comparing imaging data to the simulations being run on OLCF resources. This allows it to track oil clusters underground in both space and time. During these simulations, the team found that "disconnected" oil—meaning oil that is not directly connected to a larger, flowing oil reserve, often thought to be immobile—actually does move and contribute to flow. This work was recently awarded best paper at the 31st annual International Symposium of the Society of Core Analysts—a professional organization dedicated to highlighting scientific research into rock and core analysis.

OLCF science engagement: OLCF computational scientist Mark Berrill created a workflow to run in situ data analysis, which allows the team to perform data analysis tasks while running simulations. The team structured its code to run primarily on Titan's GPUs, and Berrill helped the team use the CPUs for data analysis, ultimately reducing the need for extensive data transfer. Berrill also increased parallelism while running on OLCF resources. These initiatives helped enable the team perform additional calculations and expand the scope of its proposed goals for the year.

The team is in the third year of its INCITE project, and the bulk of its leadership runs occured in the second year. This year, 91% of team's runs used less than 20% of Titan's nodes, and the remaining 9% used between 20%–60% of Titan's nodes. In the prior year, 94% of the team's runs were at capability scale, using 20–60% of the machine. The remaining 6% used less than 20% of Titan.

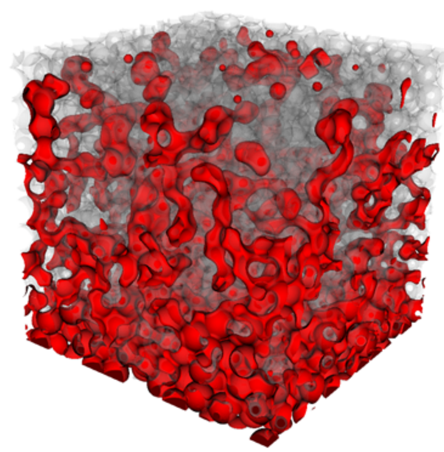

Figure 3.4. Two-fluid displacement process simulated on Titan using the lattice Boltzmann method to solve flow equations within experimentally imaged rock geometry. Credit: James McClure.

Related publications:

R. T. Armstrong, J. E. McClure M. A. Berrill, M. Rucker, S. Schluter, and S. Berg (2016) "Beyond Darcy's Law: The Role of Phase Topology and Ganglion Dynamics for Two Fluid Flow." *Physical Review E* 94, 043113 (2016). doi: 10.1103/PhysRevE.94.043113.

R. T. Armstrong, J. E. McClure, M. A. Berrill, M. Rucker, S. Schluter, and S. Berg. "Flow Regimes during Immiscible Displacement." Winner of the best paper award at the International Symposium of the Society of Core Analysts, Snow Mass, CO, August 21–26, 2016.

J. E. McClure, M. A. Berrill, W. G. Gray, and C. T. Miller. "Tracking Interface and Common Curve Dynamics for Two-Fluid Flow in Porous Media." *Journal of Fluid Mechanics* 796: 211–32 (2016).

Online stories: E. Gedenek. "Researchers Mine Information from Next-Generation Subsurface Flow Simulations." *OLCF News*, September 9, 2015.

3.2.5 First Simulations of Magnetism at the Atomic Level for a Real Nanoparticle: Markus Eisenbach, Oak Ridge National Laboratory, INCITE

Objective: To correlate the atomic structure and chemical order and disorder of an iron-platinum (FePt) nanoparticle with its magnetic properties. Using highly accurate experimental data provided by advanced imaging reconstruction techniques, this computational study simulated—for the first time—the atomic-level magnetic properties in a real nanoparticle, rather than a representative structure.

Impact: This research on the FePt nanoparticle could significantly impact the development of next-generation magnetic storage media. Magnetic storage media fabricated at the nanoscale could have a large storage capacity, but understanding the magnetic properties within a single nanoparticle—the foundational building block of this system—is critical to success. The computational approach used in this project could also be applied to any magnetic material for a range of applications if given the same quality and resolution of experimental data.

Accomplishments: ORNL and OLCF researchers joined colleagues from the University of California, Los Angeles and LBNL in combining state-of-the-art experimental imaging and analysis techniques using atomic resolution electron tomography with ORNL's supercomputing resources and award-winning quantum mechanical electronic structure code, LSMS, to model the properties of strongly magnetic regions of an FePt nanoparticle.

Typically, magnetic properties are projected at the atomic level using an idealized crystal structure for a material, which is the averaged representation of the atomic coordinates and chemical order based on x-ray crystallography results. However, real materials have atomic-scale defects that may affect performance and cannot be captured using averaged representations.

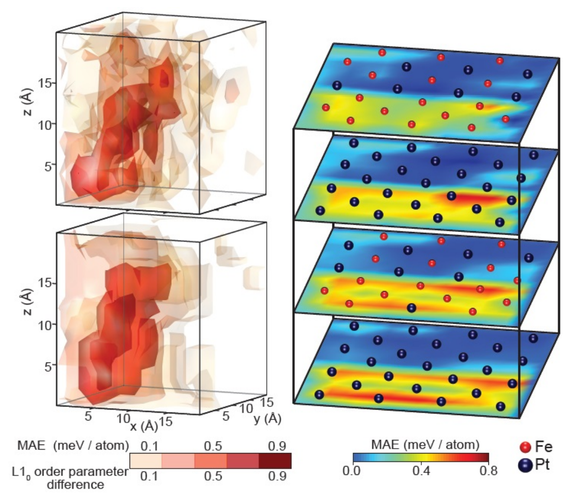

Figure 3.5. For the first time, researchers have simulated local magnetic anisotropy at the atomic level in a magnetic material based on experimental data from a single iron-platinum nanoparticle. Credit: Markus Eisenbach.

With unprecedented experimental data on the structure of an FePt nanoparticle provided by researchers at the University of California, Los Angeles and LBNL, ORNL researchers used the LSMS code on Titan to further determine the magnetic anisotropy of more than 1,300 atoms from regions of the nanoparticle. Magnetic anisotropy is the directional strength of an atom's magnetism and an important property for storing information with magnetic devices. LSMS results demonstrated how magnetic anisotropy changes between grain (regional) boundaries within the nanoparticle. In future applications, this kind of information could be used to optimize nanoparticle fabrication.

OLCF Science Engagement: Markus Eisenbach and Paul Kent of ORNL utilized an 8.6-times speedup of the LSMS code on Titan's GPUs to simulate a supercell of FePt atoms. Eisenbach and Kent also worked closely with Renat Sabirianov of the University of Nebraska at Omaha to simulate local magnetism in 32 atoms, which validated the larger LSMS simulations. In the team's INCITE project, 22% of its simulations used less than 20% of Titan and 78% used 20%–60% of the machine.

Related Publications:

Y. Yang, C-C. Chen, M. C. Scott, C. Ophus, R. Xu, A. Pryor Jr., L. Wu, F. Sun, W. Theis, J. Zhou, M. Eisenbach, P. R. C. Kent, R. F. Sabirianov, H. Zeng, P. Ercius, and J. Miao, "Deciphering Chemical Order/Disorder and Material Properties at the Single-Atom Level." *Nature* 542: 75–79 (2017). doi: 10.1038/nature21042.

3.3 DIRECTOR'S DISCRETIONARY (DD) PROGRAM

2016 Operational Assessment Guidance

The Facility should describe how the Director's Discretionary reserve is allocated and list the awarded projects, showing the PI name, organization, hours awarded, and project title.

The OLCF primarily allocates time on leadership resources through the INCITE program and through the facility's DD program. The OLCF seeks to enable scientific productivity via capability computing through both programs. Accordingly, a set of criteria are considered in making allocations, including the strategic impact of the expected scientific results and the degree that awardees can effectively use leadership resources. Further, through the ALCC program, the ASCR office allocates up to 30% of the facility's resources.

The goals of the DD program are threefold:

1. To enable users to prepare for leadership computing competitions such as, INCITE and ALCC (e.g., to improve and document application computational readiness)
2. To broaden the community of researchers capable of using leadership computing by enabling new and nontraditional research topics
3. To support R&D partnerships, both internal and external to ORNL, to advance DOE and ORNL strategic agendas

These goals are aligned particularly well with three of the OLCF's four missions:

1. To enable high-impact, grand-challenge science and engineering that could not otherwise be performed without the leadership-class computational and data resources
2. To enable fundamentally new methods of scientific discovery by building stronger collaborations with experimental facilities as well as DOE offices that have large compute and data science challenges
3. To educate and train the next-generation workforce in the application of leadership computing to solve the most challenging scientific and engineering problems

R&D partnerships are aligned with DOE and ORNL strategic agendas. They may be entirely new areas with respect to HPC, or they may be areas in need of nurturing. Example projects are those associated with the ORNL Laboratory Directed Research and Development program; the ECP; programmatic science areas (fusion, materials, chemistry, climate, nuclear physics, nuclear engineering, and bioenergy science and technology); and key academic partnerships (e.g., the UT-ORNL Joint Institute for Computational Sciences). Examples of strategic partners in the DD program include the Consortium for Advanced Simulation of Light Water Reactors; the Critical Materials Institute hub led by Ames National Laboratory; Accelerated Climate Model for Energy program; Center for Nanophase Materials

Sciences; and large experimental facilities such as the Spallation Neutron Source and the ATLAS (DOE Office of High Energy Physics) and ALICE (DOE Nuclear Physics program) experiments at CERN demonstrating at scale the PanDA workflow management system to achieve the integration of Titan into the Worldwide Large Hadron Collider Computing Grid. A science achievement highlight from the ECP-funded Cancer Moonshot project is summarized in Section 3.3.1. Also included in this broad category are projects that come to the OLCF through the Accelerating Competitiveness through Computational ExceLlence (ACCEL) Industrial HPC Partnerships program, which provides opportunities for industrial researchers to access the leadership systems to carry out research that would not otherwise be possible. See Section 3.4 for more information about ACCEL.

The OLCF DD program also supports a variety of data projects that require data storage and bandwidth capabilities but few compute resources (Section 4.2). Ongoing data projects include the Earth System Grid Federation, the BigPanDA@Titan operational demonstration, an operational demonstration of the Portal for Data Analysis Services for Cosmological Simulations, and the Majorana Demonstrator Secondary Data Archive. In addition, infrastructure software, such as frameworks, libraries, and application tools, and research support areas for next-generation operating systems, performance tools, and debugging environments are often developed by DD projects.

The Resource Utilization Council makes the final decision on DD applications, using written reviews from subject matter experts. Consistent with our integration of the OLCF and CADES capabilities, as described in the Executive Summary, the council is also managing discretionary allocations on resources (e.g., Metis, a two-cabinet Cray XK7 system, and Percival, a Cray XC40 with Intel KNL processors) for performance portability research. The actual DD project lifetime is specified upon award: allocations are typically for 1 year or less. However, projects may request 3 month extensions, or renewals up to an additional 12 months. The average size of a DD award is roughly 3 million Titan core-hours, but awards can range from tens of thousands to 12 million hours or more.

In 2016, the OLCF DD program participants used 10.31% of the total allocable resource planned for these DD program goals, consuming 456,647,202 Titan core hours.

3.3.1 Accelerating Cancer Research with Deep Learning, Director's Discretionary (DD) Program

The development and maturation of automated data tools for cancer research, part of the objectives outlined in the White House's Cancer Moonshot initiative, could give medical researchers and policymakers an unprecedented view of the US cancer population at a level of detail typically only obtained for clinical trial patients, which has been historically less than 5% of the overall cancer population. Using the Titan supercomputer, a team led by ORNL's Georgia Tourassi is making progress toward automating data tools by employing deep learning techniques to extract useful information from text-based cancer pathology reports. Thus far, the team has established deep learning's advantages in multitask learning, using nearly 2,000 cancer pathology reports to train a neural network to identify a cancer's primary site (e.g., kidney) and laterality (e.g., left or right kidney) from text. In another study, Tourassi's team deployed deep learning to match the cancer's origin to a corresponding topological code, a more specific classification than primary site or laterality. The promising performance trends measured in these studies will guide the team as they scale up deep learning algorithms to tackle larger datasets and move toward less human supervision. This discretionary project has contributed outcomes for the feasibility of the recently initiated ECP effort, the CANcer Distributed Learning Environment (CANDLE) project, led by ANL's Rick Stevens with Georgia Tourassi as a Co-PI.

Related publications:

H.-J. Yoon, A. Ramanathan, and G. Tourassi. "Multi-Task Deep Neural Networks for Automated Extraction of Primary Site and Laterality Information from Cancer Pathology Reports." In A. P., Y. Manolopoulos, L. Iliadis, A. Roy, M. Vellasco, eds. Advances in Big Data. INNS 2016. Advances in Intelligent Systems and Computing, 529 (2016). Cham, Switzerland: Springer. doi: 10.1007/978-3-319-47898-2_21.

J. X. Qiu, H.-J. Yoon, and G. Tourassi. "Deep Learning for Automated Extraction of Primary Sites from Cancer Pathology Reports." Under review.

3.4 INDUSTRIAL HPC PARTNERSHIPS PROGRAM: ACCELERATING COMPETITIVENESS THROUGH COMPUTATIONAL EXCELLENCE (ACCEL)

ACCEL continued to achieve its mission by assisting large and small companies with solving complex science and engineering problems that exceed their internal computing capabilities. After 8 years, program growth continues, albeit more slowly as it approaches a steady state of operations. Project results affirm that ACCEL is furthering DOE's goal of improving the knowledge base of the science and engineering community, as well as increasing the community of researchers able to use next-generation leadership computing resources.

Forty-three industrial projects were underway during 2016, which represented 13.4% of the total number of projects provided to external user programs (i.e., INCITE, ALCC, and DD). These projects used 242,569,881 h, representing approximately 5.6% of the total hours that Titan delivered in 2016. In 2016, 29% of the industrial project hours were allocated through INCITE, 61% through ALCC, and 10% through the OLCF DD program.

Of 43 projects, 16 were new. These new firms received awards via INCITE (1 project), ALCC (6 projects), and the DD program (9 projects). Two of the new ALCC projects were awardees of DOE's Advanced Manufacturing Office HPC4Manufacturing program. This program selects projects through a competitive call for proposals and links companies to national laboratory computational science experts and leadership computing resources to apply modeling, simulation and data analysis to advance innovation in energy-efficient manufacturing and clean energy technologies.

Two firms that received DD awards were new to ACCEL, and both were small businesses, demonstrating again that large-scale industrial problems and computational expertise are not just resident in large companies. Pinnacle Engines Inc., in collaboration with experts at ORNL's National Transportation Research Center and software firm Convergent Science, is using OLCF resources to dramatically improve the performance of a new engine for light-duty passenger vehicle applications. This new engine is expected to be competitive with best-in-class diesel engines on fuel economy and emissions. High Performance Imaging LLC is using Titan to develop algorithmic innovations using GPUs that significantly reduce the computational requirements of model-based iterative reconstruction, a technique that dramatically enhances the image quality of computer tomography (CT) applications but which can be computationally prohibitive. Through this work, model-based iterative reconstruction will become a more accessible solution for a much wider class of problems in medical imaging, security, manufacturing, and scientific imaging.

3.4.1 United Technologies Research Center Researchers Exploit GPUs and Perform Important Cost-Benefit Analysis

To better understand the complex fluid dynamics and chemical reactions happening inside a jet engine combustor, a research team at United Technologies Research Center teamed up with experts at the OLCF to improve computational models to more accurately simulate jet engine conditions. By identifying better ways to model chemical processes happening inside a jet engine combustor, United Technologies Research Center researchers hoped to develop cleaner, safer, and more efficient aircraft engines.

Using a DD award on Titan, the team achieved its two main goals. First, they successfully ported UTRC's Large Eddy Simulation with Linear Eddy code to run on Titan's GPUs, improving the chemical kinetics and giving it a five-fold speedup over CPU-only methods. Second, they performed a first ever cost-benefit analysis of a well understood, but rarely used, method of LES calculation called explicitly filtered LES. They compared this mathematically rigorous, but computationally intensive methodology to simulate combustion, with the more commonly used and less computationally intensive—but less accurate—implicitly filtered LES approach.

The team found that the gain in accuracy in explicitly filtered simulations was not worth the cost of the additional computing resources needed to achieve it, an important finding as UTRC considers its investment in HPC resources to meet its business objectives. Predictive simulation capabilities, coupled with smart testing and experimentation, are key to improving current designs and developing next-generation combustors with increased efficiencies and reduced pollutants. This exercise allowed UTRC to explore and evaluate a potentially accurate modeling methodology and conduct a cost-benefit analysis. For now, UTRC will continue to use the implicitly filtered LES approach coupled with its deep modeling expertise. However, in completing this analysis, they built the foundation for future higher-fidelity LES methods.

Related publications:

P. A. Cocks, V. Sankaran, and M. Soteriou. "Towards Predictive Reacting Flow LES." 52nd Aerospace Sciences Meeting, AIAA SciTech Forum, (AIAA 2014-0826), http://dx.doi.org/10.2514/6.2014-0826.

4. INNOVATION

CHARGE QUESTION 4: Have innovations been implemented that have improved the facility's operations?

OLCF RESPONSE: Yes. The OLCF actively pursues innovations that can enhance facility operations. Through collaborations with users, other facilities, and vendors, many of these innovations are disseminated and adopted across the country.

Since the facility's inception in 2004, OLCF staff have provided leadership in the HPC community, spearheading the creation and development of tools and policies necessary for computing and computational science. In 2016, the OLCF pursued strategic organizational and technological innovations centered around leadership computing and data challenges. It is not possible to highlight all the innovative work carried out by the OLCF. Instead, this section will focus on several key strategic areas of operations in 2016: innovative data technologies and support approaches, experimental computing and testbeds, and operational best practice development and adoption.

4.1 DATA TECHNOLOGIES AND SUPPORT APPROACHES

4.1.1 Big Data Analytics: Deep Learning

Deep Learning (DL) at the OLCF has been initiated through the DD user program. DL enables applications to learn levels of representation and abstraction that apply to large datasets and is the fastest growing subfield in machine learning. Although initially conceived several years ago, DL has become more achievable with rapid advances in modern GPU technology and larger source datasets, and is increasingly used by domain and computational scientists across the country. To this end, the ADW group introduced these capabilities to OLCF users. State-of-the-art DL libraries (e.g., Caffe, Torch, and Theano) were deployed as easily accessible module packages on Rhea and Titan. Early experiences revealed several other scalable DL libraries, such as Tensorflow, CNTK, and MXnet, that use multiple GPUs for faster DL model training are needed and are on the deployment roadmap.

The OLCF took the initiative to procure, deploy, and utilize the newly released NVIDIA DGX-1 platform, which is the world's first purpose-built system for DL and artificial intelligence analytics using high-speed NVLink technology to enable scalable (multi-GPU) training of DL models. The DGX-1 was deployed in the CADES open computing environment as a tool for preparing DL applications for Summit, which has architectural similarities to the DGX-1.

ADW group data scientist John Harney initiated and established a triweekly ORNL DL User Group meeting. The user group is an organization of DL experts, enthusiasts, and domain scientists that fosters discussions and collaborations across groups and topics at ORNL. Activities include seminars about ongoing research at ORNL, tutorials for developing models, and hackathons to solve common problems. The user group has also strengthened relationships with vendors—such as NVIDIA—to foster a collaborative environment.

4.1.2 Analytics at Scale on Titan

Large-scale simulations running on leadership-class supercomputers, such as Titan and DOE's experimental facilities, are generating unprecedented amounts of data. The ability to analyze the resulting information has become a significant challenge. Mike Matheson from the OLCF's ADW Group and George Ostrouchov from the ORNL Mathematics Division's Scientific Data Group (SDG) worked together to scale R—the most commonly used data analytics software in academia and a rising programming language in HPC—to the OLCF's Rhea, Eos, and Titan systems. Standard R users typically employ the software to analyze smaller datasets on a single workstation. However, with the Programming with Big Data (pbdR) extensions to R, users can analyze large amounts of data and scale to thousands of processors, yielding improvements of an order of magnitude or better when analyzing data.

In 2016, difficulty in computing a principal-component analysis on a huge matrix was encountered by the Center for Nanophase Materials Science (CNMS) at ORNL because of the limitations of a single workstation. The interdisciplinary team combined the expertise of scientists who posed a data-driven science problem in microscopy with the expertise of data workflow scientists, like Dale Stansberry who understand how to orchestrate and optimize workflow with big data tools in an HPC environment. Two approaches were taken: the Apache Spark cloud-based framework and the pbdR HPC framework. The initial Spark solution required the CNMS datasets to be converted to a different framework and required days of computation on less than eight nodes. However, the pbdR solution read the CNMS data directly and was computed on Titan in less than 1 min. The pbdR approach demonstrated that large datasets could be analyzed efficiently using the power of a supercomputer—thousands of cores, potentially GPUs, and parallel I/O reading the format written by the CNMS experiment. It was shown that for larger datasets, 30 thousand cores could be employed, and scalability would likely continue to larger core counts. More importantly, it demonstrated that experimental facilities of the future could achieve near real-time analysis of an experiment.

Other approaches, such as machine learning, were investigated on simulation data. Support Vector Machine was used as a data classifier and various clustering techniques were utilized. These approaches were run at up to 128 thousand cores on Titan. The clustering approach was used to analyze a time dependent dataset from a rotating-stratified turbulence simulation. The clustering approaches were evaluated to assess viability of automatically detecting the complex vortex network at small scales.

4.1.3 Scalable Data Transfer: Adaptable I/O System (ADIOS) Software

The ADIOS software received two new releases in 2016. Primary features included a new Query API with three different implementations of query engines; ZFP, the first lossy compression technique as data transformation; temporal aggregation in memory to boost the performance of frequent, small I/O; a recovery tool for damaged datasets; better staging for I/O; and an improved Python interface.

ADIOS 1.10 was released in July 2016 and was used for staging diagnostic data from the XGC fusion application running on 16 thousand compute nodes of Titan. The staging approach reduced the overhead of collecting frequent diagnostic information compared to direct file I/O from 1.5% to 0.07% (publication pending). The release also included the bprecover utility that was used for a multiday leadership-class simulation of XGC. During the simulation, the job experienced a failure that subsequently damaged the diagnostic file, which contained 500 output steps gathered over 2.5 days. The bprecover utility was used to recover the damaged diagnostic file just before the corruption. This functionality is made possible by the ADIOS-BP file format and the duplication of metadata in the dataset.

Additionally, the ability to query large datasets was released in ADIOS 1.10 with the deployment of the Query API. While there are different techniques for indexing and searching data, ADIOS supports the FastBit tool developed at LBNL and ALACRITY developed at North Carolina State University. It also includes a simple and fast indexing and querying mechanism based on minimum and maximum statistics.

The ADIOS 1.11 release brought two very important features into production in November 2016. The ZFP lossy compression method, developed by LLNL, significantly reduces datasets and gives researchers control of the accuracy of the retained datasets. A new data transformation method in ADIOS now uses ZFP to compress data. Data reduction is a very important aspect of I/O because the size and bandwidth of file systems is continually declining compared with the capability of new compute resources. The second feature introduced was temporal aggregation of data in memory. Frequent, small output still causes large I/O overhead despite the great performance ADIOS provides. Temporal aggregation is data buffering that spans multiple output steps, reducing the I/O overhead of multiple disk accesses to one. This feature works with all file-based output methods of ADIOS, which makes it ideal for analyzing diagnostic data.

4.1.4 Constellation Digital Object Identifier (DOI) Data Service

The OLCF's Constellation DOI project makes it possible for researchers to obtain a digital object identifier (DOI) to catalog and publish scientific data artifacts for open access. The DOI workflow includes submission, review, approval, publication, and dissemination of DOI data and metadata. The Office of Scientific and Technical Information serves as the registration authority for OLCF DOIs, generating and providing specific DOI values when requested, and disseminates DOI metadata through DataCite. The Constellation DOI portal provides open access for data discovery and download by interested data consumers. This project supports the increasing desire to publicly host federally funded open research results.

The Constellation DOI portal features a distributed architecture consisting of a web-application front-end supported by a mixture of custom and off-the-shelf data management services. The web-application communicates with back-end services through a proxy layer, where service requests are translated into a message-based API supported by ZeroMQ and Google Protocol Buffers. Custom back-end services include a federated identity service for authentication and session management; several data-transfer daemons to move data to and from user file systems, archival storage, and file servers; and a DOI management service. The DOI management service persists system data via a commercial RDBMS and coordinates workflows associated with DOI access, creation, review, archiving, publication, and download. The DOI data is archived and curated in the OLCF HPSS in a center managed space for long-term storage.

In December, two OLCF researchers were invited to use the Constellation DOI portal to request and acquire DOI's for published datasets. Arvind Ramanathan of the ORNL Computer Science and Engineering Division published a dataset that contains almost 8,000 images generated by bench testing reconstruction for the BigNeuron effort, a community effort to simulate a single neuron and an essential unsolved challenge in brain science. Joseph Kennedy of the ORNL Climate Change Institute published a benchmark dataset for regression testing with the Community Ice Sheet Model, supporting comparison of standard modeling tools such as the Land Ice Verification and Validation kit. These publications demonstrate a new OLCF data service capability to support scientists in publicizing results, fostering and supporting collaboration and scientific discovery.

4.2 EXPERIMENTAL COMPUTING THROUGH TESTBEDS

The OLCF, in partnership with CADES, began deployment of experimental and early access systems to pursue evaluations of novel architectures. These evaluations are critical for advancing adequate hardware and software offerings for use within DOE Leadership Computing Facilities. In 2016 the following systems were purchased, deployed, or operated at the OLCF:

- Crest—Crest is an 8 node IBM very early access prototype of pre-Summit architecture. Each node includes one Power8 CPU and four Nvidia K40x GPUs.
- Summitdev—This IBM early access Summit system is similar to Crest but has updated technology that is much closer to the Summit system. Summitdev is a 54-node system, and each node contains a Power8+ CPU and four NVIDIA P100 GPUs.

- Percival—This is a Cray XC40 system that represents a single cabinet clone of the Theta system at the ALCF. The primary goal of Percival is to provide a local platform for OLCF researchers to access diverse architectures with a focus on application portability.
- NVIDIA DGX-1—A single node consisting of eight GP100 NVIDIA GPUs and a single Intel Xeon CPU leveraging NVlink technology. The architecture is custom packaged by NVIDIA and is targeted for DL and artificial intelligence workloads.
- Cray Urika-GX—Consists of 28 compute nodes, each with two 18-core Intel Broadwell CPUs, 4 TB of hard disk drive storage, one 800 GB solid-state drive, and 256 GB of DDR4 RAM. This architecture combines the previous Cray Urika GD and XA hardware lines into an optimal system for analytics.
- Cavium ARM System—The Cavium ARM System is comprised of a 16-node testbed consisting two socket, 48-core 64-bit ARM CPUs providing access to prerelease server class ARM architectures for library and application development.

4.3 OPERATIONS

4.3.1 Data Transfer Service Upgrades to Improve Efficiency and Performance

In 2016, the OLCF began a project that would result in an upgrade to the data transfer services offered by the center. At its core, the completed project will provide increased performance and speed for data transfers to and from the OLCF. A set of highly specialized servers, DTNs, facilitate data transfers in three forms: interactive nodes, scheduled nodes, and GridFTP nodes. GridFTP is a Globus Toolkit protocol that provides highly tuned and parallel data transfers to researchers using wide-area networks. For the first time, Globus/GridFTP services will also be provided for access to the HPSS archive. Interactive nodes allow standard Unix-based utilities to be used. These utilities are typically serial in nature, and the scheduled class of nodes offers a cluster-style approach to data transfer using a workload management system that allocates the transfer across a set of nodes.

The previous DTN cluster specialized for Globus transfers were specifically tuned for wide-area networks transfers when they were first deployed in 2008. However, new use cases and capabilities allowed users to take advantage of these nodes for transfers within the OLCF as well. This increased usage coupled with recent upgrades to the HPSS necessitated the need for an upgraded DTN cluster to meet the timely demands of OLCF users.

The new DTN upgrades include a major hardware refresh, which improves the DTNs' Ethernet capabilities from 10-gigabit to 40-gigabit connectivity, enabling faster transfers without any changes required by the user. Adding 44 new DTNs, managed as a cluster utilizing diskless imaging and partitioned to act independently of one another, will create fewer opportunities for oversubscription in the system through close monitoring of the utilization of each node type. System administrators are able to quickly reprovision a node to another type when necessary. Additionally, a 100-gigabit path to ORNL's network was added, which allows for more parallel transfers for external traffic. When deployed, the new DTN cluster should be able to facilitate petabyte per week transfers to other DOE labs without adversely impacting other active transfers. The new DTN cluster was prepared for production at the end of CY 2016 and was put into production on February 7, 2017.

4.3.2 Workforce Development through Undergraduate Data Center Course

As reported in the 2015 OLCF OAR, the OLCF partnered with local ORNL facility and data center experts, the UT, and a strategic customer to offer a first of its kind data center design and management course. The course teaches critical data center concepts including electrical engineering, mechanical engineering, computer science, computer engineering, facilities management, and day-to-day operations. The class is led by data center experts and professionals from UT and ORNL, who serve as guest lecturers for the semester. Students are charged with responding to a request for proposals, requiring them to have professionally addressed all conditions and objectives in the solicitation. The request for proposals calls for a comprehensive solution to the allocation of existing space for the computational needs of the customer(s). Through their responses, the students are required to address key components when designing and operating a data center.

To date, this course has proved several internship opportunities for previous and current students and has even led to a full-time hire. Systems and Reliability Engineer John Gutman currently works on ensuring the OLCF's data center resources are running smoothly and securely. Before accepting this position, Gutman interned at ORNL while taking the data center course and was able to parlay that experience into a job.

In 2016, the guest lecture schedule was optimized to provide students more relevant experience and foundational knowledge. One specific enhancement was the engagement of data center professionals outside of ORNL. Realizing that the ORNL data centers are very mission oriented, the class organizers wanted to offer a slightly different perspective to the students. Of note was a lecture given by Keith Gray, HPC manager for British Petroleum. Gray spoke about how British Petroleum data center needs differ from those of a center like the OLCF and how having a reliable data center is essential to BP's corporate health.

4.3.3 HPC Repurposed

The OLCF was presented with a unique opportunity when the National Institute for Computational Sciences (NICS) sought to retire the four cabinet XC30 HPC resource named Darter. Creatively, the OLCF leveraged the fundamental infrastructure of that system (cabinet framing, electrical and mechanical systems, networking equipment, etc.) to build newer XC40 systems. The first system, Percival, is a single cabinet XC40 that mirrors the specifications of the Theta HPC resource from the ALCF. The mission of Percival is to provide a platform where OLCF users, who are currently porting code to the Summit node architecture, can also discover analogous ways to port code to Intel accelerators. This focus on portability is shared between the ASCR user facilities, and the deployment of Percival represents the first dedicated portability platform available for use.

The second result of this creative repurposing was the ability to deploy a new system for the Biological and Environmental Research Atmospheric Radiation Measurement Climate Research Facility. This new system marks the first dedicated Biological and Environmental Research–funded compute system managed by the NCCS. With this new capability, researchers will be given a computational boost in performance and support with the NCCS managing the system. This new system will also interact between operational enclaves and will serve as one of the key driving factors in the CADES and OLCF/NCCS integration activities scheduled to start in CY 2017.

By leveraging key partnerships with NICS and Cray, the OLCF saved an average of one million dollars per cabinet when compared to purchasing each system brand new. With the partnership, the OLCF will allow NICS to use up to ten percent of the Percival resource for their purposes.

4.3.4 XALT Implementation

In 2016, the OLCF completed the transition of the Automatic Library Tracking Database (ALTD) to XALT on all computational resources. XALT replaced the previously used ALTD utility, which is the production library tracking tool of choice at several HPC facilities around the world. XALT is a lightweight software tool that is designed to run on any Linux cluster, workstation, or high-end supercomputer to track executable information and linkage of static shared and dynamically linked libraries. By moving to XALT, the OLCF has a better understanding of the software that researchers need and are using on the systems. XALT provides a modular installation and support stance that is flexible for facilities who prefer to digest the data in or out of real-time. By moving to XALT the OLCF remains consistent with the broader HPC community in meeting and anticipating the needs of computational researchers.

5. RISK MANAGEMENT

CHARGE QUESTION 5: Is the facility effectively managing risk?

OLCF RESPONSE: Yes, the OLCF has a very successful history of anticipating, analyzing, rating, and retiring both project- and operations-based risks. The OLCF risk management approach is modeled after the Project Management Institute's best practices. Risks are tracked and, when appropriate, retired, reclassified, or mitigated. A change history is maintained for historical reference.

The major risks for the OLCF in CY 2016 are listed and described in this section. Planned mitigations and implementations are included in the subsequent descriptions. As of this writing, the OLCF has two high-priority operational risks:

1. Decreasing availability of spare parts as Titan ages. The OLCF addressed this risk during this reporting period by partnering with Cray and NVIDIA to analyze historical failures, which resulted in a proactive GPU replacement strategy to maintain Titan's high output.
2. Significant facility investment required for an exascale system. The OLCF continues to examine the ever-changing requirements for an exascale system's operations and to evaluate options for delivering those requirements on schedule for a potential system delivery.

More information on these risks can be found in subsequent sections of this report.

5.1 RISK MANAGEMENT SUMMARY

The OLCF's Risk Management Plan describes a regular, rigorous, proactive, and highly successful review process. The Risk Management Plan is reviewed at least annually and is updated as necessary. The plan covers both OLCF operations and its various projects. Each project execution plan refers to the main Risk Management Plan but may incorporate project specific adjustments. Risks are tracked in a risk registry database application that is capable of tracking project and operational risks separately.

Weekly operations and project meetings are held and risks are continually assessed and monitored. Specific risk meetings are held monthly and are attended by the federal project director, facility management, OLCF group leaders, and others subject matter experts and risk owners. When assessing risks, the OLCF management team focuses its attention on the high and moderate risks as well as any low risks within the impact horizons associated with the risk. Trigger conditions and impact dates are recorded in the risk notes narrative section of the register. Risk owners are proactive in tracking trigger conditions and impact horizons for their risks and bringing appropriate management attention to those risks, regardless of the risk-rating level.

The OLCF reports a change summary of affected risks to the DOE program office as part of its monthly operations report. At the time of this writing, 26 active entries are in the OLCF operations risk register that fall into two categories: risks for the entire facility and risks for a specific portion of the facility. Facility-wide risks are concerned with such things as safety, funding and expenditures, and staffing. More focused risks are concerned with reliability, availability, and use of the system or its components (e.g., the computing platforms, power and cooling, storage, networks, software, and user support). In addition to operational risks, at the time of this report, there are also 44 tracked risks for the OLCF-4 project.

Costs for handling risks are integrated in the budgeting exercises for the entire facility. For operations, the costs of risk mitigation are accepted, and residual risk values are estimated by expert opinion and are accommodated as much as possible in management reserves. This reserve is continually reevaluated throughout the year.

5.2 MAJOR RISKS TRACKED IN 2016

Table 5.1 contains the major risks tracked for OLCF operations in 2016. The full OLCF operations risk register is available on request. The selected risks are all rated medium or high in impact.

Table 5.1. 2016 OLCF major risks

Risk ID /description	Probability/ impact	Action	Status
361: Scientists decline to port to heterogeneous architectures	Medium/medium	Mitigating	This risk was not triggered. There was a lower number of applications to the INCITE 2017 program, but there was a significant increase in ALCC requests. Additionally, the 13 CAAR applications signify the move by application scientists to the heterogeneous architectures.
406: System cyber security failures	Low/high	Mitigating	The OLCF continues to see a rise in the quantity of cyber-security attacks against the computer resources. This increase does not directly correlate to higher success rates as the OLCF employs various techniques to repel these attacks, such as proactive patching for zero-day exploits, formal review of cyber security plans, a two-factor authentication requirement for system access, and a multifactor authentication (MFA)4 requirement for privileged access to OLCF resources.
721: Lustre metadata performance continues to impact applications	Medium/medium	Mitigating	Lustre 2.8 was deployed, which enabled the deployment of multiple metadata servers. The OLCF deployed these new metadata servers and storage in 2016 that eliminated bottlenecks in the original hardware design. The new systems provide solid-state drive storage for the metadata, which had a significant positive impact on metadata performance.
906: Programming environment tools may be insufficient	Low/medium	Mitigating	The OLCF leverages the expertise of the NCCS Software Tools Group to evaluate gaps in the programming environment and tools. The OLCF purchased the Forge tool, which includes DDT and MAP, from Allinea to address gaps identified by the Software Tools Group.
917: Robust support will not be available to ensure portability of restructured applications	Medium/medium	Mitigating	The OLCF deployed multiple compilers which maximizes the exposure of multiple levels of concurrency in user applications. The OLCF involvement in the standards bodies such as the OpenACC consortium continues to assist in mitigating this risk.

Table 5.1. (continued)

Risk ID /description	Probability/ impact	Action	Status
948: Significant facility investment required for exascale systems	Medium/high	Accept	Anticipated facility requirements for the ECP are extremely demanding. The deployment of the OLCF-4 project includes the addition of 20 MW of technical load, which should support pre-exascale plans. The push to exascale will require additional uplifts to the current facility.
997: Problems with system reliability, diagnosis, and recovery in large hybrid systems may arise	High/low	Mitigating	In addressing risk 1154 ("Lack of available spare parts causes issues as Titan ages"), the OLCF triggered this risk. Diagnosing the failure mechanism was a significant effort that impacted long-running, large jobs on Titan.
1006: Inability to acquire sufficient staff	Medium/low	Accept	The OLCF reduced the probability of encountering this risk to medium in 2015. The same status was maintained for 2016. The number of open positions carried over from CY 2016 is lower than the threshold determined to trigger this risk (10%).
1063: Programming environment and tools may be inadequate for future architectures	Medium/medium	Mitigating	In response to the gaps identified in addressing risk 906 ("Programming environment tools may be insufficient"), the OLCF deployed MAP from Allinea, and continues to engage with user communities and standards organizations to address feedback received from the OLCF user community.
1142: OLCF cost increases because fewer computer room customers to distribute maintenance and operation costs among	Low/high	Mitigating	In 2016, the data center customer base remained static, and slightly fewer systems were deployed. As 2017 begins, new customers and projects are anticipated. The cleanout of and facility upgrades to an existing data center, projected for FY 2018, will significantly increase the available data center space.
1145: Changes from external project managers cause development impacts to HPSS	Medium/medium	Mitigating	IBM has continued to push for items that are not on the development roadmap to support requests of potential customers.
1154: Decreasing availability of spare parts as Titan ages	High/high	Mitigating	At this time the OLCF has a sufficient cache of spare parts to satisfy the repair demands for the remainder of Titan's operational lifetime. The number of spare parts was developed using a predictive model from Cray and NVIDIA based on observed failure rates.

5.3 RISKS THAT OCCURRED DURING THE CURRENT YEAR AND THE EFFECTIVENESS OF THEIR MITIGATION

The following risks were encountered and effectively mitigated in 2016. A short summary of the status and impact of the risk on the operations of the OLCF is included.

Risk No. 997	Problems with reliability, diagnosis, and recovery in a large hybrid system may arise		
Risk owner	James H. Rogers		
Status	Mitigating		
Probability	High		
Impact	***Cost:*** High	***Schedule:*** High	***Scope/Tech:*** Medium
Mitigations	Mitigations for this risk include continued development of diagnostic tools that can provide hardware and system administrators with the tools and mechanisms to effectively diagnose state and failure conditions for the GPU.		
Triggers	Intelligence on actual or likely problems will cause this to occur.		

The OLCF encountered this risk during the reporting period when the failure of GPU components in Titan began to impact user jobs. Investigation of the specific failures led to the discovery of a failed component on the GPU SXM module. The results of this investigation triggered Risk No. 1154. Components were proactively replaced with new parts that are not susceptible to the discovered issue. The OLCF worked with vendor partners to determine a suitable strategy to satisfy the operational lifetime and requirements of Titan.

Risk No. 1154	Decreasing availability of spare parts as Titan ages		
Risk owner	Don E. Maxwell		
Status	Mitigate		
Probability	High		
Impact	***Cost:*** High	***Schedule:*** Medium	***Scope/Tech:*** Low
Mitigations	We will stockpile more spare parts and monitor failure rates to stay ahead of issues.		
Triggers	Measured stability trends and a burn rate of allocated hours at a rate that will not meet annual commitments will cause this. Additional insight is available by examining failures in relation to the number of nodes that are allocated to a particular job.		

When the 2016 Titan issue was fully understood, the supply of parts was not sufficient to meet the operational lifetime and requirements of Titan. The OLCF engaged the appropriate vendor partners and developed a strategy to fulfill the operational requirements of Titan.

5.4 RISKS RETIRED DURING THE CURRENT YEAR

The following risks were retired from OLCF operations during this reporting period.

Risk No. 1156	Lack of support for HSI/HTAR tools		
Retirement date	2016-12-31		
Risk owner	Kevin G. Thach		
Status	Retired		
Probability	Low		
Impact	***Cost:*** Low	***Schedule:*** Low	***Scope/Tech:*** Low
Mitigations	OLCF evaluated alternative mechanisms to move data in and out of HPSS (e.g., HPSS/Globus DSI). We do have a source for the HSI/HTAR utilities, so their usage could be extended (without official support) while transitioning to other tools. IBM is offering to support HTAR, and NERSC and LBNL are considering supporting HSI in repair-only mode.		
Retirement comments	After the announcement that NERSC and Gleicher Enterprises will collaborate on building developers and a support organization for the HSI/HTAR tools at LBNL, this risk is no longer active for OLCF operations.		

Risk No. 1130	Reduced cooling tower performance
Retirement date	2016-08-02
Risk owner	Bart A. Hammontree
Status	Retired
Probability	Low
Impact	***Cost:*** Low ***Schedule:*** Low ***Scope/Tech:*** Low
Mitigations	We intend to monitor local wet bulb conditions and use local data to size new cooling towers. The location of the new towers will be carefully evaluated to ensure it does not negatively impact the existing towers.
Retirement comments	With the construction of the OLCF-4 cooling tower substantially complete, the monitoring continues to show there will be no negative impact on OLCF operations from the locations of the cooling towers.

5.5 NEW OR RECHARACTERIZED RISKS SINCE LAST REVIEW

5.5.1 Recharacterized Risks

The following risks had changes in their status or impact in the reporting period.

Risk No. 406	System cyber security failures
Risk owner	Kevin G. Thach
Status	Accept → Mitigating
Probability	Low
Impact	***Cost:*** High ***Schedule:*** High ***Scope/Tech:*** Medium
Mitigations	We have developed a cyber security plan that implements protection to the moderate level. The plan includes two-factor authentication and periodic formal tests and reviews. We have a response plan to react to incidents, some of which is official use only.
Triggers	Intrusion detection system alerts, manual log analysis indicating unauthorized access, or a combination of both would trigger this event. Reports of odd system behavior from system administrators or users are also a trigger.

Risk No. 1142	OLCF cost increases because there are fewer computer room customers to distribute maintenance and operation costs among
Risk owner	Stephen McNally
Status	Mitigating
Probability	Medium → Low
Impact	***Cost:*** High ***Schedule:*** Low ***Scope/Tech:*** Low
Mitigations	Maintenance costs could be reduced by eliminating the equipment that needs to be maintained. It is also possible this risk could be shared with ORNL Facilities and Operations Directorate as they might be willing to contribute to the maintenance costs if this is in the best interest of the laboratory as a whole. OLCF may be forced to decommission portions of the existing central energy plant.
Triggers	Loss of any substantive customer in the data center

5.5.2 New Risks in This Reporting Period

The following risks were created and tracked during CY 2016. They are included with their risk creation date, mitigations, and triggers.

Risk No. 1154	Decreasing availability of spare parts as Titan ages
Risk creation date	2016-01-08
Risk owner	Don E. Maxwell
Status	Mitigating—**Current**

Probability	High
Impact	***Cost:*** Medium ***Schedule:*** Low ***Scope/Tech:*** High
Mitigations	We will stockpile more spare parts and monitor failure rates to stay ahead of issues.
Triggers	Measured stability trends and a burn rate of allocated hours at a rate that will not meet annual commitments will trigger this. Additional insight is available by examining failures in relation to the number of nodes that are allocated to a particular job.

Risk No. 1155	Extended life of Atlas file system causes increase in parts failures, possible downtime, and possible loss of data
Risk creation date	2016-03-08
Risk owner	Kevin G. Thach
Status	Mitigating—**Current**
Probability	Low
Impact	***Cost:*** Medium ***Schedule:*** Low ***Scope/Tech:*** Medium
Mitigations	Support contracts are extended through CY 2019 to cover all hardware. The OLCF is engaging with their software support to extend that support as well.
Triggers	Increases in applications failures resulting from the storage system, increased parts failures, and decreases in file system performance are all triggering events.

5.6 MAJOR RISKS FOR NEXT YEAR

The following risks are the top OLCF risks for CY 2017 with mitigation and trigger descriptions.

Risk No. 948	Significant facility investment required for an exascale system
Risk owner	James H. Rogers
Status	Accept
Probability	Medium
Impact	***Cost:*** High ***Schedule:*** High ***Scope/Tech:*** Medium
Mitigations	ORNL has a plan to house the exascale system in building 5600 by moving other systems out of the building. However, the much-preferred approach would be to construct a new building that is designed for exascale from the beginning. The Office of Management and Budget (OMB) has rejected third-party financing as a method of building such a facility, so this will require a congressional line item funding.
Triggers	Intelligence on the size and power requirements of proposed systems.

Rep. No. 1145	Changes from external project managers cause development impacts to HPSS
Risk owner	Sudharshan S. Vazhkudai
Status	Mitigate
Probability	Medium
Impact	***Cost:*** Medium ***Schedule:*** Medium ***Scope/Tech:*** Low
Mitigations	Each lab gets a representative to the technical council, which helps set direction for HPSS development. OLCF works closely with the council to ensure our users' needs are met. However, IBM has trumped established development plans before, causing disruptions in schedules.
Triggers	Input from the technical council changing project direction will trigger this disruption.

Rep. No. 1154	Decreasing availability of spare parts as Titan ages
Risk owner	Don E. Maxwell
Status	Mitigating
Probability	High
Impact	***Cost:*** Medium ***Schedule:*** Low ***Scope/Tech:*** High
Mitigations	We will stockpile more spare parts and monitor failure rates to stay ahead of issues.

Triggers	Measured stability trends and a burn rate of allocated hours at a rate that will not meet annual commitments will trigger this. Additional insight is available by examining failures in relation to the number of nodes that are allocated to a particular job.

5.6.1 Risks That Will Transition from the OLCF-4 Project into Operations

These risks will transition from the OLCF-4 project into the operations risk registers, but not until CY2018. They are listed with their mitigations and triggers.

Risk No. 1163	Nontraditional computer room may cost more or have unexpected requirements
Risk owner	Bart A. Hammontree
Status	Mitigating
Probability	Low
Impact	***Cost:*** High ***Schedule:*** Medium ***Scope/Tech:*** Low
Mitigations	The OLCF-4 project used an outside architectural firm to design the mechanical and electrical systems inside the computer room. Additionally, a LIDAR 3-D scan of the room was used to build a complete model before construction to view any potential problems.
Triggers	The project has worked closely to understand the requirements of the overhead utilities and continues to engage with subcontractors as the design progresses.

Risk No. 1164	Nontraditional computer room construction created new challenges for safety of operations and maintenance
Risk owner	Bart A. Hammontree
Status	Mitigating
Probability	Low
Impact	***Cost:*** Low ***Schedule:*** Low ***Scope/Tech:*** Low
Mitigations	Working closely with maintenance personnel to fully understand procedures will mitigate this risk. A culture of safety will refine any issues that are encountered.
Triggers	Any increase in reportable safety incidents to DOE will trigger this risk.

Site Office Safety Metrics

HIGH PERFORMANCE COMPUTING FACILITY
2016 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2017

6. SITE OFFICE SAFETY METRICS

CHARGE QUESTION 6: Has the site implemented measures for safety of staff and the public that are appropriate for HPC/networking facilities?

OLCF RESPONSE: Yes.

6.1 SUMMARY

ORNL is committed to operating under the DOE safety regulations specified in 10 CFR 851,[6] "Worker Safety and Health Program" as well as applicable DOE orders and standards. These safety requirements are incorporated into the ORNL contract as required compliance documents. To implement these safety requirements in a consistent manner across ORNL, UT-Battelle LLC deploys an online procedure management system referred to as the Standards-Based Management System (SBMS). Within SBMS, there are work control requirements that describe the processes to be used in ORNL operations and R&D activities to implement integrated safety management functions and principles.

A key feature of the integrated safety management process is the development and implementation of project-specific research safety summaries (RSS), which define the scope of work, identify and analyze hazards, and establish safety procedures. Each RSS is reviewed and approved by the ORNL Safety Services Division, line managers, and research staff. An RSS provides the means for ORNL management and staff to plan and conduct research in a safe manner. It is used to control work, train participants, and provide information about operations and emergency services if needed. Under a work control review system, work plans are also written before maintenance work is allowed to proceed, to ensure work is conducted safely. Safety specifications are written into the service contracts and undergo a review by the authority with jurisdiction before new construction and service subcontractors are allowed to begin work.

Safety assessments are conducted for RSSs, work plans, and subcontracts, as well as inspections of job sites throughout each year. Lessons learned, safety snapshots, safety talks, and management assessments are conducted and recorded in the Assessment and Commitment Tracking System. The tracking system documents the completion of the ORNL integrated safety management process and provides a means for analysis. The DOE ORNL Site Office participates in the field implementation and documentation of all operational safety reviews, and it partners with the ORNL Offices of Institutional Planning, Integrated Performance Management, and Safety Services Division on independent safety management system assessments.

The culture of safety at ORNL is reflected in these processes, which seek to reduce and prevent injuries to personnel and potential exposure to hazards associated with operation of the facility. The OLCF works closely with the site office and Dan Hoag, the Federal Project Director, who solicited the

---

[6] 10 CFR 851 outlines the requirements for a worker safety/health program to ensure DOE contractors and their workers operate a safe workplace. Additionally, 10 CFR 851 establishes procedures for investigating if a violation of a requirement has occurred, for determining the nature and extent of any such violation, and for imposing an appropriate remedy.

following feedback from the site office staff in the Operations and Oversight Division about OLCF's safety culture.

- A review of the monthly safety charts and the total recordable cases and days away, restricted, or transferred (DART) summary documents indicated that overall FY 2015 total recordable cases rate increased slightly, and DART rates have decreased when compared with FY 2015. Operations of the OLCF in the NCCS remained safe, efficient, and effective as there were zero total recordable cases and DARTs in FY 2016.
- UT-Battelle emphasized the subjects of material handling and electrical safety because of incidents and recommendations from DOE. UT-Battelle lowered the incident rate of material handling injuries by incorporating rules that provided lifting restrictions in some scenarios, and UT-Battelle is in the final stages of finalizing procedures that will specifically address this subject matter and provide guidance to all employees. The Electrical Safety Program continues to be enhanced by implementing an arc flash analysis procedure, creating roles, responsibilities, authorities, and accountabilities for division electrical safety officers, enhancing the lockout-tagout and electrical training, and reviewing electrical equipment across the site to evaluate the risk of using equipment that is not inspected by a National Recognized Testing Laboratory.
- The OLCF-4 project Health and Safety Plan and the Hazard Analysis were reviewed 2016.
- A sound level survey was conducted in the Titan data center to determine areas where noise exceeded 85 dBA. The survey was conducted using noise dosimetry on selected employees to determine exposure during normal work day. The survey results revealed that requiring all ORNL employees to undergo mandatory hearing conservation training was unnecessary. This training is now optional and is available to any employee who requests enrollment in the program. The removal of this mandatory hearing conservation program will result in a small cost savings without incurring any increased risk of hearing loss or injury to employees. The requirement to wear hearing protection in the data center remains unchanged.

The OLCF strives to exceed the minimum requirements by requiring limited scope hazard analyses and by employing award-winning professionals throughout the facility. Each contractor is required to submit an activity hazard analysis to UT-Battelle personnel for review so UT-Battelle can be aware of the scope of work, ensure hazards and controls have been appropriately addressed, and understand the plan for completing the job scope. This process keeps UT-Battelle aware of the projects being completed and holds contractors accountable for recognizing hazards and implementing mitigating controls.

Cyber Security

HIGH PERFORMANCE COMPUTING FACILITY
2016 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2017

7. CYBER SECURITY

CHARGE QUESTION 7: Have innovations been implemented to improve the facilities' cyber security posture? Does the site have a valid cyber security plan and authority to operate?

OLCF RESPONSE: Yes, the most recent OLCF authority to operate was granted on November 11, 2016. The current authority to operate expires on June 30, 2017.

7.1 SUMMARY

All information technology systems operating for the federal government must have certification and accreditation to operate. This involves the development and approval of a policy and the implementation of a continuous monitoring program to confirm the policy is effectively put into practice. The OLCF has the authority to operate for 7 months under the ORNL certification and accreditation package approved by DOE on November 11, 2016. This package is an interim authority to operate while ORNL restructures several networks, which will be tested at the beginning of CY 2017, and site-wide reaccreditation for a full year is expected at that time. The ORNL certification and accreditation package currently uses *Recommended Security Controls for Federal Information Systems and Organizations* (National Institute of Standards and Technology Special Publication 800-53, revision 3, 2009) as a guideline for security controls. The OLCF is accredited at the moderate level of controls for protecting the confidentiality and integrity of user and system information (Federal Information Processing Standards Publication 199), which authorizes the facility to process sensitive, proprietary, and export-controlled data.

In the future, cyber security planning will become more complex as the center continues its mission to produce great science. As the facility moves forward, the OLCF is very proactive, viewing its cyber security plans as dynamic documentation to which it will preemptively respond and modify as the needs of the facility change to provide an appropriately secure environment.

7.2 MULTIFACTOR AUTHENTICATION (MFA) LEVEL OF ASSURANCE 4 (LOA4)

The OMB set a Level of Assurance 4 (LOA4) multifactor authentication goal during a Cyber Sprint activity shortly after the Office of Personnel Management was breached in July 2015. That breach resulted in the release of 21.5 million records of current and former government employees, including their names, dates of birth, and social security numbers. The DOE multifactor implementation project took the goal set by the Office of Personnel Management requiring multifactor authentication LOA4 (MFA4) and set aggressive implementation targets that also required 100% of privileged users and 85% of standard users to use LOA4 credentials to log into interactive sessions as of September 30, 2016. NCCS previously used RSA LOA3 credentials, which consisted of a PIN plus a one-time passcode generated by an RSA token, to authenticate OLCF users. The following steps were necessary to upgrade to LOA4 credentials, which consist of DOE-issued personal identity verification (PIV) badges:

- Obtain an LOA3 exception for standard users
- Identify NCCS privileged users
- Guide privileged users through the DOE PIV badge application process
- Issue applicable users X.509 certificates after verifying their DOE PIV badge

- Convert X.509 certificates into Secure Socket Shell (SSH) public keys
- Deploy SSH Secure Socket Shell server configuration to enforce use of MFA4 by privileged users
- Deploy the derived X.509 SSH Secure Socket Shell public keys
- Create Lightweight Directory Access Protocol groups to house privileged users
- Develop a privileged user configuration auditing process

Additional OLCF internal goals included:

1. Obtain LOA3 exceptions for standard users
2. Achieve compliance ahead of the September 30, 2016, implementation deadline
3. Use standard NCCS configuration management systems to deploy changes
4. Ensure continuity of operations during implementation
5. Minimize cost to the organization
6. Achieve productivity gains for systems administrators by using PIV badges as single-use sign-on credentials

All project goals were met or exceeded. The NCCS-OLCF implementation team was able to successfully meet the MFA4 access requirements well ahead of the mandated deadline and at significantly less cost than other similar organizations. Approximately two full-time employee's worth of effort was used during FY 2016 to secure approximately 4,000 NCSS computing systems and achieve MFA4 compliance. Additionally, systems administrators are able to use the new LOA4 credential more easily than RSA tokens in a single sign-on fashion that does not compromise security, which achieves the stretch goal of making productivity gains while adhering to the new OMB mandate. Overall, the NCCS-OLCF team was able to take a very difficult and challenging set of requirements, and find a unique, creative solution that also enhances productivity. Changes were implemented efficiently, thoroughly, and in a deliberate, methodical way that did not impact production operations.

7.3 DEVELOPMENT OF A PUBLIC KEY–BASED PASSWORD DISTRIBUTION SYSTEM

Recognizing the need for a more sophisticated and manageable method for maintaining and distributing root passwords to administrative staff, a new public key-based password distribution system was developed by the HPC Operations Cyber Security staff. Several third-party password management solutions exist, but few are designed for enterprise systems administrators sharing and participating in on-call rotations. Some rely on unencrypted-in-memory password stores, but passwords could be exposed in the event the central password repository was compromised. Others encrypt passwords at rest, but require administrators to manage complex GPG public key infrastructure and keyrings. With the current use of MFA4, DOE is now distributing PIV cards with cryptographically signed certificate and private key pairs to administrators and other security-significant users. This new distribution system offers the NCCS and OLCF a mechanism to use this public key infrastructure to securely encrypt passwords for other recipients. This new system is cross platform (i.e., it works on OSX and Linux systems) and provides the following critical features: encryption at rest, password distribution and organization based on definable hierarchies, password creation time stamps, password history and change logs, distributed backup capabilities, MFA4 credential-based encryption and decryption, and must be based on common smart card encryption and decryption tools, such as OpenSC and OpenSSL.

The new distribution system was adopted by OLCF systems administration staff to replace a very antiquated and difficult to manage secured paper notebook checkout procedure for maintaining root passwords for OLCF systems. This innovation makes it much easier to track who has knowledge of root passwords for NCCS systems and deal with changing passwords when staff leave or no longer require access. New features and advanced functionality are planned and are being worked on as time allows, and ultimately there are plans to make the utility available to other national labs and the greater open-source community.

Summary of the Proposed Metric Values

HIGH PERFORMANCE COMPUTING FACILITY
2016 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2017

8. ACTUAL AND PROPOSED METRIC VALUES

CHARGE QUESTION 8: Are the performance metrics used for the review year and proposed for future years sufficient and reasonable for assessing Operational performance?

OLCF RESPONSE: Yes. The OLCF works closely with the DOE program manager to develop and update metrics and to target values that reflect the expectations of the stakeholders in delivering a leadership-class HPC resource.

8.1 SUMMARY

Table 8.1 provides a summary of the metrics and actual data for the current reporting period and proposed metrics and targets for the subsequent 2 years.

Table 8.1. OLCF metrics and actual data for 2016 and proposed metrics and target for 2017 (2015 data provided for context)

2016 metric and target	2015 target	2015 actual	2016 target	2016 actual	2017 target
Are the processes for supporting the customers, resolving problems, and outreach effective?					
Customer Metric 1: Customer Satisfaction					
Overall score on the OLCF user survey. Target: Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample. (Annual)	The OLCF exceeded the metric target: 4.6/5.0.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample. (Annual)	The OLCF exceeded the metric target: 4.6/5.0.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample
Improvement on results that scored below satisfactory in the previous period. Target: Results will show improvement on at least one-half of questions that scored below satisfactory (3.5) in the previous period.	Results will show improvement on at least one-half of the questions that scored below satisfactory (3.5) in the previous period. (Annual)	The OLCF exceeded the metric target: No question scored below satisfactory (3.5/5.0) on the 2015 survey.	Results will show improvement on at least one-half of the questions that scored below satisfactory (3.5) in the previous period. (Annual)	The OLCF exceeded the metric target: No question scored below satisfactory (3.5/5.0) on the 2015 survey.	Results will show improvement on at least one-half of the questions that scored below satisfactory (3.5) in the previous period.
Customer Metric 2: Problem Resolution					
OLCF survey results related to problem resolution. Target: Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample. (Annual)	The OLCF exceeded the metric target: 4.6/5.0.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample. (Annual)	The OLCF exceeded the metric target: 4.6/5.0.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample
OLCF user problem resolution time period. Target: 80% of OLCF user problems will be addressed within 3 business days, by either resolving the problem or informing the user how the problem will be resolved.	Eighty percent of OLCF user problems will be addressed within 3 business days, by either resolving the problem or informing the user how the problem will be resolved. (Monthly)	The OLCF exceeded the metric target: 92%.	Eighty percent of OLCF user problems will be addressed within 3 business days, by either resolving the problem or informing the user how the problem will be resolved. (Monthly)	The OLCF exceeded the metric target: 92%.	Eighty percent of OLCF user problems will be addressed within 3 business days, by either resolving the problem or informing the user how the problem will be resolved.

Table 8.1. (continued)

2016 metric and target	2015 target	2015 actual	2016 target	2016 actual	2017 target
Customer Metric 3: User Support					
Average of user support ratings. Target: Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample. (Annual)	The OLCF exceeded the metric target: 4.6/5.0.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample. (Annual)	The OLCF exceeded the metric target: 4.6/5.0.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample
Is the facility maximizing the use of its HPC systems and other resources consistent with its mission?					
Business Metric 1: System availability (for a period of one year following a major system upgrade, the targeted scheduled availability is 85% and overall availability is 80%)[a]					
Scheduled Availability. 2015 Targets: Titan: 95%; HPSS: 95%; external file systems: 95%.	Titan: 95%; HPSS 95%; external file systems 95% (Monthly)	The OLCF exceeded the metric target. Titan: 99.41%; HPSS: 99.94%; /atlas0: 99.77%; /atlas1: 99.76%.	Titan: 95%; HPSS 95%; external file systems 95% (Monthly)	The OLCF exceeded the metric target. Titan: 99.64%; HPSS: 99.89%; /atlas0: 99.92%; /atlas1: 99.92%.	Titan: 95%; HPSS 95%; external file systems 95%.
Overall Availability. 2015 Targets: Titan: 90%; HPSS 90%; external file systems 90%.	Titan: 90%; HPSS 90%; external file systems: existing, 90% (Monthly)	The OLCF exceeded the metric target: Titan: 97.01%; HPSS: 98%; /atlas0: 98.9%; /atlas1: 98.88%.	Titan: 90%; HPSS 90%; external file systems: existing, 90% (Monthly)	The OLCF exceeded the metric target: Titan: 97.03%; HPSS: 98.54%; /atlas0: 98.2%; /atlas1: 98.2%.	Titan: 90%; HPSS 90%; external file systems: existing, 90%.
Business Metric 2: Capability Usage					
OLCF will report on capability usage. Target: In the first year of production, at least 30% of the consumed node-hours will be from jobs requesting 20% or more of the available compute nodes. In subsequent years, this increases from 30 to 35%.	In subsequent years, at least 35% of the consumed node-hours will be from jobs requesting 20% or more of the available compute nodes. (Monthly)	The capability usage was 64.19%. The OLCF exceeded the metric target.	In subsequent years, at least 35% of the consumed node-hours will be from jobs requesting 20% or more of the available compute nodes. (Monthly)	The capability usage was 59.01%. The OLCF exceeded the metric target.	In subsequent years, at least 35% of the consumed node-hours will be from jobs requesting 20% or more of the available compute nodes.
OLCF will report GPU usage (reference only, no target)	N/A[b]	53.13%	N/A[b]	56.96	N/A[b]

[a] The Cray XK7, Titan, went into production on May 31, 2013. The external file system, Spider II, went into production on October 3, 2013.
[b] Not applicable.

APPENDIX A. RESPONSES TO RECOMMENDATIONS FROM THE 2015 OPERATIONAL ASSESSMENT REVIEW

In March 2016, the OLCF presented the 2015 operational activities of the center to the DOE sponsor. During this review, the review committee identified one recommendation.

Recommendation

OLCF created a clear and useful formula to define %gpu_activity, which the center is able to measure. However, values of %gpu_activity were not reported as such in the document. A table like 2.13 for GPU activity percentage would be interesting, including the overall usage.

Response

GPU_activity values were included in the 2015 OLCF OAR in Figure 2.4. However, a table similar to table 2.13 from the 2015 OLCF OAR is provided below for convenience. An initial hypothesis was that the GPU_activity data would show usage efficiencies of the GPUs. While reviewing the data, it became clear that the data is not particularly helpful at the macro analysis level. Given that this measure reports GPU SM usage and does not take into account memory or I/O usages, is it short sighted and assumes that SM usage is the one and only measure that is relevant when evaluating GPU usage effectiveness. As such, the OLCF decided to collect and monitor this data, but that it was non-essential to track for the operation of Titan. Rather, the GPU_activity data provides a reasonable high-level view of the overall activity which is useful for internal planning and tracking purposes. This is especially helpful when OLCF staff partner with research teams who are interested in optimizing their applications. It is also worth noting that while the industry remains very focused on effective usage of GPUs, we do not currently measure usage efficiency on other system components such as disk, cpu, or memory. For these reasons, the OLCF has decided to continue to track and measure GPU_activity for internal purposes but will not include it as a standard reportable metric.

Table A.1. GPU_activity by core hours.

	GPU_activity (in millions)				GPU_enabled (in millions)			
	INCITE	**ALCC**	**DD**	**Total Hours**	**DD**	**ALCC**	**INCITE**	**Total Hours**
Jan	43.51	24.72	2.50	70.73	6.95	90.51	61.51	158.97
Feb	42.48	10.76	2.17	55.41	7.55	37.77	91.84	137.17
Mar	38.10	18.58	1.01	57.69	3.45	60.32	115.68	179.45
Apr	34.90	8.35	4.76	48.00	12.86	31.38	137.79	182.03
May	35.79	11.50	2.07	49.37	8.65	52.49	126.78	187.93
Jun	24.27	13.55	3.37	41.19	10.00	31.95	122.77	164.72
Jul	24.94	13.06	3.37	41.37	7.97	18.22	126.05	152.25
Aug	60.25	31.54	5.89	97.68	17.50	52.60	115.82	185.92
Sep	39.21	61.23	3.12	103.55	6.14	91.08	136.37	233.59
Oct	51.98	26.12	4.10	82.20	12.36	65.57	168.73	246.66
Nov	47.51	12.12	4.30	63.94	13.17	29.13	150.82	193.12
Dec	54.09	8.80	2.23	65.13	4.27	13.86	184.89	203.02
Total	497.04	240.33	38.89	776.26	110.88	574.88	1,539.06	2,224.83

The overall assessment comments of the OLCF for CY 2015 are as follows:

The OLCF did an exceptional job this year – deploying Titan while still managing to deliver value to its users. They did a great job addressing issues with the machine. OLCF continues to exhibit effective management strategies for providing scientific capabilities to a varied and demanding user community.

This report provides a sharp and detailed description of its planning process, risk management strategies, and overall operations. We expect that OLCF will continue to have direct and significant impact on innovation and driving discoveries.

Excellent. In addition to successfully rolling out a new HPC system, the OLCF facility continues to innovate and provide an effective platform for science done on the cutting edge of HPC. The OLCF has done a great job in operating the facility in 2013. In spite of major issues with the main computational resource (Titan), the facility managed to reach the availability and utilization metrics. The center has also done an excellent job in getting users prepared to use the novel architecture of the system at scale and provides very good response to user issues and problems. The facility has also done a good job in outreach and industrial partnerships. The innovative work done at the center over the 2013 period was impressive and will not only assist facility staff with managing the systems effectively, but also improve the user experience overall.

APPENDIX B. TRAINING, WORKSHOPS, AND SEMINARS

Event type	Description	Date	No. Participants
Monthly User Conference Calls	2016 OLCF User Conference Call: How to OLCF	1/27/16	46
Monthly User Conference Calls	2016 OLCF User Conference Call: Check Point Advisory Tool and Darshan	2/24/16	41
Monthly User Conference Calls	2016 OLCF User Conference Call: Allinea Performance Report Tools and DTN Queues	3/23/16	36
Monthly User Conference Calls	2016 OLCF User Conference Call: Lustre Tips and Performance Portability of XGC	4/27/16	40
Monthly User Conference Calls	2016 OLCF User Conference Call: Publication Tracking	6/29/16	26
Monthly User Conference Calls	2016 OLCF User Conference Call: Introduction to YT Visualization on Titan/Rhea	7/27/16	35
Monthly User Conference Calls	2016 OLCF User Conference Call: How to OLCF	8/8/16	13
Monthly User Conference Calls	2016 OLCF User Conference Call: Development of Strategy and Procedure for Assessing Portability in Industrial Codes to GPUs Hardware: The TITAN Example	8/31/16	30
Monthly User Conference Calls	2016 OLCF User Conference Call: Managing Python on OLCF Resources	9/28/16	51
Monthly User Conference Calls	2016 OLCF User Conference Call: PyFR on Titan	10/26/16	39
Monthly User Conference Calls	2016 OLCF User Conference Call: Getting Started with ADIOS	12/14/16	28
Seminar Series	A Framework to Analyze the Propagation of Transient Faults in HPC Applications	1/11/16	—
Seminar Series	Quantum Computing Systems and Software	1/13/16	—
Seminar Series	BEAM: A Computational Workflow System Enabling Scalable in Silico & Empirical Exploration of Materials Science Data in the DOE HPC Cloud	1/13/16	—
Seminar Series	GPU Computing Update and Roadmap	1/20/16	—
Seminar Series	Computational Studies of Novel Radon Containing Molecules	1/22/16	—
Seminar Series	Achieving High Reliability and Efficiency in Maintaining Large-Scale Storage Systems through Optimal Resource Provisioning and Data Placement	1/25/16	—
Seminar Series	Campaign Storage and MarFS: So You Want to Store a 1 PB file, or 50 Million 20 MB Files?	1/29/16	—

Event type	Description	Date	No. Participants
Seminar Series	IBM Watson Cognitive Services: Application to the Sciences	2/17/16	—
Seminar Series	Application-Level Fault Tolerance Using Proposal from the MPI Forum's Fault Tolerance Working Group	3/2/16	—
Seminar Series	Substituting Disk Failure Avoidance for Redundancy in Wide Area Fault Tolerant Storage Systems	3/24/16	—
Seminar Series	Interview: Predicting Arctic Sea Ice Variability with a Data-Driven Forecasting Method	4/25/16	—
Seminar Series	Biomolecular Structure and Dynamics: Probing Behavior of Oligonucleotides and Small Proteins in Solution State, Solid Phase, and at the Tethered Surface	5/18/16	—
Seminar Series	On the Implications of Large-Scale Manycores and NoCs for Exascale Prof. Frank Mueller, North Carolina State University	1/11/16	—
Seminar Series	Code Integration Between the BISON Fuel Performance and PROTEUS Neutronics Applications	1/13/16	—
Seminar Series	"Oh Rats!"—The NCCS Resource and Allocation Tracking System	1/13/16	—
Seminar Series	Integrated Modeling and Simulation with Eclipse ICE and Its Applicability to Neutron Science	1/20/16	—
Seminar Series	The Divide-Expand-Consolidate (DEC) Scheme in Large-Scale Molecular Science Computations	5/19/16	—
Seminar Series	Towards Large Scale Relativistic Coupled Cluster Applications	7/13/16	—
Seminar Series	Towards Large Scale ab initio 4c and 2c EPR Calculations and More (MC-srDFT)	7/14/16	—
Seminar Series	IBM's TrueNorth Neuomoprhic Computing Chip	7/20/16	—
Seminar Series	Coupled Electron and Nuclear Dynamics	7/22/16	—
Seminar Series	White Dwarf Mergers as Astrophysical Explosion Progenitors	8/1/16	—
Seminar Series	Multiscale Modeling Order in Condensed Matter: A Guide Towards Optimized Electro-Optic Materials	8/23/16	—
Seminar Series	Data I/O Middleware for Isolating Development Complexities and Enabling Real-Time Optimization	10/3/16	—
Seminar Series	Future of HDF	10/12/16	—
Seminar Series	Kanor: A Language for Declarative Communication	10/28/16	—

Event type	Description	Date	No. Participants
Seminar Series	Quantum Chemistry Calculations of Energetic and Spectroscopic Properties of p-, d-, and f-block Molecules	11/4/16	—
Seminar Series	DASH; A C++ PGAS Library for Distributed Data Structures and Parallel Algorithms	11/9/16	—
Seminar Series	Rise to the Sun: A Computational/ Observational Perspective on the Solar Dynamo Problems	11/29/16	—
Seminar Series	Towards Large Scale Simulations of Turbulent Reacting Flow Problems Involving Moving Boundaries	12/5/16	—
Seminar Series	Applications of Quantum Chemistry in Molecular Biophysics	12/12/16	—
Seminar Series	Numerical Modeling of a Binary Alloy Droplet Deposition Process	12/15/16	—
IDEAS/Facility Webinar Series	IDEAS/ASCR Seminar Series (Session 1): What All Codes Should Do: Overview of Best Practices in HPC Software Development	5/4/16	324
IDEAS/Facility Webinar Series	IDEAS/ASCR Seminar Series (Session 2): Developing, Configuring, Building, and Deploying HPC Software Developing	5/18/16	125
IDEAS/Facility Webinar Series	IDEAS/ASCR Seminar Series (Session 3): Distributed Version Control and Continuous Integrated Testing	6/2/16	98
IDEAS/Facility Webinar Series	IDEAS/ASCR Seminar Series (Session 4): Testing and Documenting Your Code	6/15/16	77
IDEAS/Facility Webinar Series	IDEAS/ASCR Seminar Series (Session 5): How ithe HPC Environment is Different from the Desktop (and Why)	7/14/16	54
IDEAS/Facility Webinar Series	IDEAS/ASCR Seminar Series (Session 6): An Introduction to High-Performance Parallel I/O	7/28/16	64
IDEAS/Facility Webinar Series	IDEAS/ASCR Seminar Series (Session 7): Basic Performance Analysis and Optimization—An Ant Farm Approach	8/9/16	61
Workshop/Training/ Meeting	Parallel Demo	2/4/16	15
Workshop/Training/ Meeting	2016 GPU Hackathons—TU-Dresden	2/29/16–3/4/16	34
Workshop/Training/ Meeting	INCITE Proposal Writing Webinar	4/13/16	49
Workshop/Training/ Meeting	Eclipse User Group and Developers Meeting	4/14/16	31
Workshop/Training/ Meeting	2016 GPU Hackathons—UDELHACK	5/2/16–5/6/16	63
Workshop/Training/ Meeting	INCITE Proposal Writing Webinar	5/19/16	76

Event type	Description	Date	No. Participants
Workshop/Training/ Meeting	2016 OLCF User Meeting: Closing in on the Summit	5/24/16–5/26/16	151
Workshop/Training/ Meeting	HPC Debugging and Profiling with Allinea Forge: Getting Started	6/24/16	34
Workshop/Training/ Meeting	INCITE Computational Readiness Kickoff Call	7/7/16	49
Workshop/Training/ Meeting	Supercomputing in a Nutshell	7/14/16	—
Workshop/Training/ Meeting	CPU Computing Hands-On Session at the CSGF Annual Review	7/27/16	—
Workshop/Training/ Meeting	HPC Debugging and Profiling with Allinea Forge: Getting Started	8/19/16	14
Workshop/Training/ Meeting	HPC Day at Tapia 2016: GPU Accelerated Computing	9/17/16	—
Workshop/Training/ Meeting	2016 GPU Hackathons—OLCFHack	10/17/16–10/21/16	68
Workshop/Training/ Meeting	Mini-GPU Hackathon	11/1/16–11/3/16	45
Workshop/Training/ Meeting	2016 GPU Hackathons: EuroHack	7/4/16–8/16/16	40

APPENDIX C. OUTREACH PRODUCTS

Date	Product	Title
4/18/16	Blog Post	New Users Tap Titan for Another Year of World-Class Supercomputing Research
11/1/16	Fact Sheets	INCITE Fact Sheets
2/3/16	Graphic	ACCEL Boeing Graphic
3/1/16	Graphic	IMPACT Materials
5/2/16	Graphic	DOE Slides ASCR Program Manager
5/13/16	Graphic	Astro Code Concept Art
5/16/16	Graphic	OLCF User Meeting Materials
7/7/16	Graphic	Smoky Mountain Conference Graphics
7/11/16	Graphic	Awards Night Materials
8/1/16	Graphic	Women in Computing Poster
8/21/16	Graphic	OLCF Slide Show (Static Version)
9/6/16	Graphic	OLCF Slide Show (Animated Video)
9/13/16	Graphic	Summit Logo Contest
9/28/16	Graphic	Titan Week Graphics
1/5/16	Highlight	Fighting Fire with FireFOAM
1/5/16	Highlight	OLCF Delivers New Job Launching Tool
1/5/16	Highlight	Tool to Optimize New Codes
1/5/16	Highlight	BigNeuron Hackathon Branches Out at OLCF
1/26/16	Highlight	Titan Targets Tumors
1/26/16	Highlight	ACCEL Program Continues to Experience Growth, Success
1/26/16	Highlight	On the Road to Summit: More Power
1/26/16	Highlight	OLCF Team Works to Facilitate Data Transfer
2/15/16	Highlight	Biofuel Researchers Employ Titan to Probe “Lignin Shield”
2/15/16	Highlight	Titan Probes Depths of Biofuel’s Biggest Barrier
2/15/16	Highlight	OLCF Training Video Channel Now Open
2/15/16	Highlight	Staff Report Successes, Lessons from Titan-Era Early Applications
3/7/16	Highlight	Boeing Catches Caution from the Wind
3/7/16	Highlight	Illuminating the Universe’s Ignition
3/7/16	Highlight	OLCF-Developed Tool Offers a Window into Center Operations
3/7/16	Highlight	OLCF Storage Upgrades Paying Off for Users
3/28/16	Highlight	Broadening the Bilayer
3/28/16	Highlight	OLCF Group to Offer Spark On-Demand Data Analysis
3/28/16	Highlight	New Software to Enhance Performance for OLCF Users
3/28/16	Highlight	Intra-Lab Collaboration Automates OLCF Publication Tracking System with Cobra Tool
4/18/16	Highlight	Chemistry Consortium Uses Titan to Understand Actinides
4/18/16	Highlight	Scientific Computing Group Helps Projects Prepare for Next-Generation Supercomputers
4/18/16	Highlight	OLCF Provides Leadership, Expertise for Dresden GPU Hackathon
4/18/16	Highlight	OLCF-MDF Data Collaboration Advances 3-D Printing
4/18/16	Highlight	New Users Tap Titan for Another Year of World-Class Supercomputing Research

Date	Product	Title
5/9/16	Highlight	Titan Shines Light on High-Temperature Superconductor Pathway
5/9/16	Highlight	Titan Goes on Tour
5/9/16	Highlight	Eclipse Software Workshop Brings Together Experts, Enthusiasts
5/9/16	Highlight	OLCF Updates Lustre Community on New Features at Annual Workshop
5/30/16	Highlight	Better Combustion for Power Generation
5/30/16	Highlight	OLCF-Total Partnership Bolsters Simulation Strength for Oil and Gas Industry
5/30/16	Highlight	X-Stack Projects Use Titan's GPUs for Demos
5/30/16	Highlight	OLCF, ALCF, NERSC Co-Host HPC Software Webinar Series
6/14/16	Highlight	Opening Neurotransmission's Gatekeepers
6/14/16	Highlight	Tiny Titan Travels Again
6/14/16	Highlight	At GPU Technology Conference, OLCF Staff Teach Best Practices for Portability, OpenACC
6/14/16	Highlight	Accelerating Codes at University of Delaware Hackathon
6/14/16	Highlight	Graphic Design USA, PRSA Laud 2014–15 OLCF Annual Report
7/5/16	Highlight	Physics Researchers Question Calcium-52's Magic
7/5/16	Highlight	One Billion Processor Hours Awarded to 22 Projects through ALCC
7/5/16	Highlight	Wells Brings HPC Expertise to Seismology Meeting
7/5/16	Highlight	Big Success at OLCF's Largest User Meeting
7/5/16	Highlight	OLCF Expands Data Analytics Capability with Popular Programming Language
7/26/16	Highlight	Fundamental Fission Modeling Finds a Foothold
7/26/16	Highlight	OLCF Experts Lead and Learn at Cray User Group Meeting
7/26/16	Highlight	OLCF Staff Inspire Innovation at Bay Area Maker Faire
7/26/16	Highlight	Summer Interns' Innovations Benefit OLCF and Beyond
7/26/16	Highlight	Oracle's Large Tape Users Group Elects OLCF System Administrator as President
8/17/16	Highlight	Delving into the Dark Sky
8/17/16	Highlight	FM Global Researchers Say "ADIOS" to Bottlenecks at OpenFOAM Conference
8/17/16	Highlight	Directorate's Daughters Run Code on Titan at WiC Event
8/17/16	Highlight	OLCF to Play Key Role in New International Accelerated Computing Collaboration
9/5/16	Highlight	Streamlining Accelerated Computing for Industry
9/5/16	Highlight	OLCF Dives into Deep Learning
9/5/16	Highlight	CSGF Annual Program Review Brings Staff and Students Together
9/5/16	Highlight	OLCF Staff Engage at International Supercomputing Conference
9/26/16	Highlight	Physicists Quench Their Thirst for Modeling Superfluids
9/26/16	Highlight	OLCF-Fermilab Collaboration Gives ADIOS a Boost
9/26/16	Highlight	OLCF Team Resolves Performance Bottleneck in OpenACC Code
9/26/16	Highlight	#TitanWeek to Celebrate Titan Supercomputer's 4th "Birthday"
10/17/16	Highlight	Unraveling the Science Behind Biomass Breakdown
10/17/16	Highlight	ORiGAMI Folds in Open-Source Functionality, Named a Finalist for R&D 100 Awards
10/17/16	Highlight	OLCF User Support Specialist Receives YWCA Award Nomination
10/17/16	Highlight	OLCF Contributes to ORNL–UT Data Center Boot Camp
10/17/16	Highlight	ISSOAS Summer School
11/7/16	Highlight	Accelerated Cancer Research with Deep Learning

Date	Product	Title
11/7/16	Highlight	October Hackathon Brings New Teams, GPU Success
11/7/16	Highlight	User Notifications Streamlined Over Last Decade
11/7/16	Highlight	OLCF Staff Members Lead HPC Sessions at 2016 Tapia Conference
11/7/16	Highlight	Experts Converge on ORNL for US-Japan Fusion Workshop
12/5/16	Highlight	United Technologies Research Center
12/5/16	Highlight	ORNL Celebrates Women at Grace Hopper Conference
12/5/16	Highlight	SC Wrap-Up
12/5/16	Highlight	OLCF Data Management Course
6/30/16	Poster	ALCC Poster
7/1/16	Poster	OLCF Recruiting Poster
7/10/16	Poster	CSGF Poster
10/3/16	Poster	SC16 Poster #1 (Science/Summit)
10/3/16	Poster	SC16 Poster #2 Industry
12/5/16	Poster	INCITE 2017 Projects Poster
1/5/16	PPT Slide	OLCF Delivers New Job Launching Tool
1/5/16	PPT Slide	Tool to Optimize New Codes
1/5/16	PPT Slide	BigNeuron Hackathon Branches Out at OLCF
1/26/16	PPT Slide	ACCEL Program Continues to Experience Growth, Success
1/26/16	PPT Slide	On the Road to Summit: More Power
1/26/16	PPT Slide	OLCF Team Works to Facilitate Data Transfer
2/15/16	PPT Slide	OLCF Training Video Channel Now Open
2/15/16	PPT Slide	Staff Report Successes, Lessons from Titan-Era Early Applications
3/7/16	PPT Slide	OLCF-Developed Tool Offers a Window into Center Operations
3/7/16	PPT Slide	OLCF Storage Upgrades Paying Off for Users
3/28/16	PPT Slide	OLCF Group to Offer Spark On-Demand Data Analysis
3/28/16	PPT Slide	New Software to Enhance Performance for OLCF Users
3/28/16	PPT Slide	Intra-Lab Collaboration Automates OLCF Publication Tracking System with Cobra Tool
4/18/16	PPT Slide	Scientific Computing Group Helps Projects Prepare for Next-Generation Supercomputers
4/18/16	PPT Slide	OLCF Provides Leadership, Expertise for Dresden GPU Hackathon
4/18/16	PPT Slide	OLCF-MDF Data Collaboration Advances 3-D Printing
4/18/16	PPT Slide	New Users Tap Titan for Another Year of World-Class Supercomputing Research
5/9/16	PPT Slide	Titan Goes on Tour
5/9/16	PPT Slide	Eclipse Software Workshop Brings Together Experts, Enthusiasts
5/9/16	PPT Slide	OLCF Updates Lustre Community on New Features at Annual Workshop
5/30/16	PPT Slide	OLCF-Total Partnership Bolsters Simulation Strength for Oil and Gas Industry
5/30/16	PPT Slide	X-Stack Projects Use Titan's GPUs for Demos
5/30/16	PPT Slide	OLCF, ALCF, NERSC Co-Host HPC Software Webinar Series
6/14/16	PPT Slide	At GPU Technology Conference, OLCF Staff Teach Best Practices for Portability, OpenACC
6/14/16	PPT Slide	Accelerating Codes at University of Delaware Hackathon
6/14/16	PPT Slide	Graphic Design USA, PRSA Laud 2014–15 OLCF Annual Report
7/5/16	PPT Slide	One Billion Processor Hours Awarded to 22 Projects through ALCC

Date	Product	Title
7/5/16	PPT Slide	Wells Brings HPC Expertise to Seismology Meeting
7/5/16	PPT Slide	Big Success at OLCF's Largest User Meeting
7/5/16	PPT Slide	OLCF Expands Data Analytics Capability with Popular Programming Language
7/26/16	PPT Slide	OLCF Experts Lead and Learn at Cray User Group Meeting
7/26/16	PPT Slide	OLCF Staff Inspire Innovation at Bay Area Maker Faire
7/26/16	PPT Slide	Summer Interns' Innovations Benefit OLCF and Beyond
7/26/16	PPT Slide	Oracle's Large Tape Users Group Elects OLCF System Administrator as President
8/17/16	PPT Slide	Directorate's Daughters Run Code on Titan at WiC Event
8/17/16	PPT Slide	OLCF to Play Key Role in New International Accelerated Computing Collaboration
9/5/16	PPT Slide	Streamlining Accelerated Computing for Industry
9/5/16	PPT Slide	OLCF Dives into Deep Learning
9/5/16	PPT Slide	CSGF Annual Program Review Brings Staff and Students Together
9/5/16	PPT Slide	OLCF Staff Engage at International Supercomputing Conference
9/26/16	PPT Slide	OLCF-Fermilab Collaboration Gives ADIOS a Boost
9/26/16	PPT Slide	OLCF Team Resolves Performance Bottleneck in OpenACC Code
10/17/16	PPT Slide	ORiGAMI Folds in Open-Source Functionality, Named a Finalist for R&D 100 Awards
10/17/16	PPT Slide	OLCF User Support Specialist Receives YWCA Award Nomination
10/17/16	PPT Slide	OLCF Contributes to ORNL–UT Data Center Boot Camp
10/17/16	PPT Slide	ISSOAS Summer School
10/26/16	PPT Slide	OLCF at SC for OLCF All Hands Meeting
11/7/16	PPT Slide	October Hackathon Brings New Teams, GPU Success
11/7/16	PPT Slide	User Notifications Streamlined Over Last Decade
11/7/16	PPT Slide	OLCF Staff Members Lead HPC Sessions at 2016 Tapia Conference
11/7/16	PPT Slide	Experts Converge on ORNL for US-Japan Fusion Workshop
12/5/16	PPT Slide	ORNL Celebrates Women at Grace Hopper Conference
12/5/16	PPT Slide	SC Wrap-Up
12/5/16	PPT Slide	OLCF Data Management Course
3/1/16	PPT Slides	OA PowerPoint Slides
6/20/16	Publication	OLCF Annual Report
1/5/16	Quad Chart	Fighting Fire with FireFOAM
1/26/16	Quad Chart	Titan Targets Tumors
2/15/16	Quad Chart	Biofuel Researchers Employ Titan to Probe "Lignin Shield"
2/15/16	Quad Chart	Titan Probes Depths of Biofuel's Biggest Barrier
3/7/16	Quad Chart	Boeing Catches Caution from the Wind
3/7/16	Quad Chart	Illuminating the Universe's Ignition
3/28/16	Quad Chart	Broadening the Bilayer
4/18/16	Quad Chart	Chemistry Consortium Uses Titan to Understand Actinides
5/9/16	Quad Chart	Titan Shines Light on High-Temperature Superconductor Pathway
5/30/16	Quad Chart	Better Combustion for Power Generation
6/14/16	Quad Chart	Opening Neurotransmission's Gatekeepers
6/14/16	Quad Chart	Tiny Titan Travels Again
7/5/16	Quad Chart	Physics Researchers Question Calcium-52's Magic

Date	Product	Title
7/26/16	Quad Chart	Fundamental Fission Modeling Finds a Foothold
8/17/16	Quad Chart	Delving into the Dark Sky
8/17/16	Quad Chart	FM Global Researchers Say “ADIOS” to Bottlenecks at OpenFOAM Conference
9/26/16	Quad Chart	Physicists Quench Their Thirst for Modeling Superfluids
10/17/16	Quad Chart	Unraveling the Science Behind Biomass Breakdown
11/7/16	Quad Chart	Accelerated Cancer Research with Deep Learning
12/5/16	Quad Chart	UTRC
1/31/16	Report	ALCC Closeouts
3/1/16	Report	Contributions for Operational Assessment
6/8/16	Video	OLCF Overlook Tour in a Box Video: Overview
10/4/16	Video	OLCF Tour in a Box Video: Science
10/4/16	Video	OLCF Tour in a Box Video: Titan Deep Dive
10/4/16	Video	OLCF Tour in a Box Video: Groups
11/14/16	Website	SC16 Event Details
11/14/16	Website	SC16
11/14/16	Website	INCITE 2017

APPENDIX D. BUSINESS RESULTS FORMULAS

2016 Operational Assessment Guidance

Scheduled Availability

For HPC Facilities, scheduled availability (reference formula #1) is the percentage of time a designated level of resource is available to users, excluding scheduled downtime for maintenance and upgrades. To be considered a scheduled outage, the user community must be notified of the need for a maintenance event window no less than 24 hours in advance of the outage (emergency fixes). Users will be notified of regularly scheduled maintenance in advance, on a schedule that provides sufficient notification, and no less than 72 hours prior to the event, and preferably as much as seven calendar days prior. If that regularly scheduled maintenance is not needed, users will be informed of the cancellation of that maintenance event in a timely manner. Any interruption of service that does not meet the minimum notification window is categorized as an unscheduled outage.

A significant event that delays a return to scheduled production will be counted as an adjacent unscheduled outage. Typically, this would be for a return to service four or more hours later than the scheduled end time. The centers have not yet agreed on a specific definition for this rare scenario.

$$SA = \left(\frac{time\ in\ period - time\ unavailable\ due\ to\ outages\ in\ period}{time\ in\ period - time\ unavailable\ due\ to\ scheduled\ outages\ in\ period} \right) * 100 \qquad (1)$$

Overall Availability

Overall availability (reference formula #2) is the percentage of time a system is available to users. Outage time reflects both scheduled and unscheduled outages.

$$OA = \left(\frac{time\ in\ period - time\ unavailable\ due\ to\ outages\ in\ period}{time\ in\ period} \right) * 100 \qquad (2)$$

Mean Time to Interrupt

Time, on average, to any outage on the system, whether unscheduled or scheduled. Also known as MTBI (Mean Time between Interrupt, reference formula #3).

$$MTTI = \left(\frac{time\ in\ period - (duration\ of\ scheduled\ outages + duration\ of\ unscheduled\ outages)}{number\ of\ scheduled\ outages + number\ of\ unscheduled\ outages + 1} \right) \qquad (3)$$

Mean Time to Failure

Time, on average, to an unscheduled outage on the system (reference formula #4).

$$MTTF = \frac{time\ in\ period - (duration\ of\ unscheduled\ outages)}{number\ of\ unscheduled\ outages + 1} \qquad (4)$$

System Utilization

The percent of time that the system's computational nodes run user jobs. No adjustment is made to exclude any user group, including staff and vendors (reference formula #5).

$$SU = \left(\frac{core\ hours\ used\ in\ period}{core\ hours\ available\ in\ period} \right) * 100 \qquad (5)$$

APPENDIX E. DIRECTOR'S DISCRETIONARY PROJECTS ENABLED (AT ANY POINT) IN CY 2016

PI	Institution	Allocation	Utilization	Project Name
Vahid Ranjbar	Brookhaven National Laboratory	2,000,000	1,913,685	Spin Tracking for RHIC
Ilya Pogorelov	Tech-X Corporation	1,000,000	0	GPU-Accelerated Beam Dynamics Simulations
Remi Lehe	Lawrence Berkeley National Laboratory	500,000	10,029	Innovative Spectral Particle-In-Cell Algorithms for Laser-Wakefield Acceleration
Ji Qiang	Lawrence Berkeley National Laboratory	100,000	0	particle accelerator modeling on exascale computers
Luigi Capone	Texas A&M University - Commerce	1,000,000	372,708	HYDRA GPUs architecture migration task I
Ravichandra Srinivasan	Dresser-Rand Company	32,000,000	70,530	Compressible Flow Turbomachinery Optimization: Numerical Tools Advancement
Brian E. Mitchell	GE Global Research	200,000	174,753	TACOMA GPU Port Readiness
Kumar Srinivasan	Fiat Chrysler Automobiles	1,500,000	961,907	Advanced Shape Optimization Studies to Improve Automotive Aerodynamics
Shailendra Kaushik	General Motors	4,000,000	17,648	Multi-Disciplinary Optimization of Automobile Aerodynamic Performance, Thermal Performance, and Aero-Acoustics
Dan Williams	ANSYS Inc.	4,000,000	4,736,609	ANSYS AIM Fluid Solver Scaling
Jaime Peraire	Massachusetts Institute of Technology	4,000,000	102,361	Scalable Discontinuous Galerkin Solvers for the Implicit Large-Eddy Simulation of Transitional Flows
Peter Edward Vincent	Imperial College London	37,500,000	38,797,224	Development and Benchmarking of an In-Situ Visualization Pipeline for Next Generation CFD Tools
Mujeeb Malik	NASA Langley Research Center	3,000,000	363,569	Wall-Resolved Simulations of Complex Separated Aerodynamic Flows
Bronson Messer	Oak Ridge National Laboratory	15,000,000	0	Explosive Nucleosynthesis and Deflagration to Detonation in Type Ia Supernovae
Eric J Lentz	The University of Tennessee	35,000,000	17,445,861	Examination of Spatial and Network Resolution Core-Collapse Supernova Simulations
Kelly Holley-Bockelmann	Vanderbilt University	1,000,000	135,732	Binary Black Hole Coalescence in Galactic Nuclei
Zhihui Du	Tsinghua University	10,000,000	27,944	GPU Accelerated Gravitational Wave source modeling
Pierre Ocvirk	Universite de Strasbourg	2,000,000	12,479	Preparing Cosmic Dawn II

PI	Institution	Allocation	Utilization	Project Name
Brant Robertson	University of California - Santa Cruz	3,000,000	2,032,726	Extending the Physics of the GPU-Enabled CHOLLA Code to the Power of Titan
Joachim Gerhard Stadel	University of Zurich	6,000,000	5,250,730	Optimal Time-to-Solution Cosmology Simulation
Alexander Chekhovskoy	University of California - Berkeley	700,000	126,222	Effects of Magnetic Field on Compact Object Mergers and Nucleosynthesis
Philip Fajardo Hopkins	California Institute of Technology	4,000,000	897,092	Galaxies on FIRE: Shedding Light on Dark Matter
Charles Horowitz	Indiana University	1,000,000	14,526	Crystallization of Rapid Proton Capture Ash
Kiran Alapaty	United States Environmental Protection Agency	1,800,000	131,281	MPAS-US EPA
Daniel Jacobson	Oak Ridge National Laboratory	1,000,000	445,478	Scaling Up Parallelized Ortholog Detection Algorithms for Comparative Genomics of Bacterial Genomes
Hwee Kuan Lee	Oak Ridge National Laboratory	100,000	0	Development of Probabilistic Neural Network for Computer Aided Diagnosis of Cancer Specimens
Volkhard Helms	Universität des Saarlandes	2,000,000	2,910,418	Accessing Structural and Dynamic Features of Solute Carriers to Predict Their Substrates and Inhibitors
Peter D. Kwong	Columbia University in the City of New York	0	0	Molecular Dynamics Simulations of HIV-1 Envelope Glycoprotein Trimers
Pratul K Agarwal	Oak Ridge National Laboratory	8,000,000	0	Conformational Substates in Enzyme Catalytic Efficiency
Frank Noe	Freie Universtat Berlin	12,000,000	8,236,406	Adaptive Molecular Simulation of the Immunological Synapse
Chongle Pan	Oak Ridge National Laboratory	5,000,000	5,636,450	Large-Scale Integrated Omics Analyses of Complex Microbial Communities in Plant-Microbe Interfaces and Tropical Rainforests
Ruth Nussinov	Leidos Biomedical Research for Frederick National Lab for Cancer Research	1,000,000	0	Large Scale Simulations of Oncogenic Ras Protein Its Complexes with Effectors and Regulators
Xiaolin Cheng	Oak Ridge National Laboratory	23,500,000	7,830,374	Functional Domains in Asymmetric Membranes Probed with High-Performance Simulation and Neutron Scattering
Edward Lyman	University of Delaware	5,000,000	5,094,186	Subdiffusive Signatures of Nanoscale Lipid Bilayer Domains
Jens Meiler	Vanderbilt University	5,000,000	399,009	Structure Determination of Cellulose Synthase Using Neutron Scattering and High Performance Computing
Gerald W. Feigenson	Cornell University	2,000,000	1,844,396	Proteins Stiffen Membranes: Bending Moduli of Multicomponent Lipid Bilayers and the Effects of Peptides

PI	Institution	Allocation	Utilization	Project Name
Jun Fan	City University of Hong Kong	6,000,000	3,926,700	Multi-Scale Simulations on Membrane Remodeling by BAR Domain Proteins
D. Peter Tieleman	University of Calgary	650,000	147,517	Lipid-lipid Interactions Underlying Lateral Organizations in Cell Membranes
David N. Beratan	Duke University	1,000,000	673,332	Computer Assisted Diversity Oriented Molecular Library Design
Pak-Lee Chau	Institut Pasteur	500,000	67,896	Simulation of the Nicotinic Acetylcholine Receptor
Joseph Curtis	National Institute of Standards and Technology (NIST)	500,000	112	Atomistic Modeling of Small-Angle Scattering Experimental Data on HPC Resources via SASSIE-web
Mitchel John Doktycz	Oak Ridge National Laboratory	5,000,000	8,768,408	Multiscale Simulations in Adaptive Biosystems Imaging
Remco Hartkamp	Vanderbilt University	500,000	311,284	Permeability of Multicomponent Gel-phase Bilayers
Miguel A Fuentes-Cabrera	Oak Ridge National Laboratory	3,000,000	2,238,358	Theoretical Study of the 2D Self-Assembly of the Carboxysome
Yevgeny Moskovitz	Ben-Gurion University of the Negev	260,000	4,466,285	Validation of the Multi-Site Expansion Hypothesis of Neurological Impairments in N-methyl-D-aspartate Receptor
Robert Best	NIH National Institute of Diabetes and Digestive and Kidney Diseases	4,000,000	5,193,816	Substrate Translocation Mechanism of a Membrane Transporter Protein
Charlotte Barbier	Oak Ridge National Laboratory	2,000,000	0	Cavitation Bubbles near Solid Surfaces
Ariana Beste	The University of Tennessee	1,500,000	1,708,892	Predictive Computational Catalysis
David N. Beratan	Duke University	1,000,000	0	Toward Infra-Red Controlled Charge Transfer Reactions in Chemistry and Biology
Jacek Jakowski	The University of Tennessee	3,000,000	0	Quantum Dynamics of Nano-Scale Materials
Karl Andrew Wilkinson	University of Cape Town	4,000,000	0	Dynamic properties of porous frameworks upon the absorption of gas molecules
Vyacheslav Bryantsev	Oak Ridge National Laboratory	1,500,000	671,550	Rational Design of Flotation Agents for Rare Earth Ore Minerals
Aurora Clark	Washington State University	5,000,000	0	Real-Time Data Analytics of Molecular Dynamics Simulations
Giriprakash Palanisamy	Oak Ridge National Laboratory	0	0	ARM Archive Large Scale Data Processing System
James Joseph Hack	Oak Ridge National Laboratory	23,520,000	0	Ultra High Resolution Global Climate Simulation to Explore and Quantify Predictive Skill for Climate Means, Variability and Extremes
Katherine J. Evans	Oak Ridge National Laboratory	11,700,000	9,859,573	The Computational Climate Science Integrated Allocation

PI	Institution	Allocation	Utilization	Project Name
Francois William Primeau	University of California - Irvine	150,000	3	Implementation and Testing of a Tracer Transport Matrix for the MPAS-O Ocean Model
William I Gustafson Jr	Pacific Northwest National Laboratory	0	0	The DOE Atmospheric Radiation Measurement Programs LES ARM Symbiotic Simulation and Observation (LASSO) Initialization, Forcing and Multiscale Data Assimilation Program
Valentine Anantharaj	Oak Ridge National Laboratory	100,000	0	Workflow Automation and Advanced Analytics for Climate Simulations
Kiran Alapaty	United States Environmental Protection Agency	2,500,000	458,590	Next Generation Climate Modeling using the CESM and MPAS Modeling Systems
Valentine Anatharaj	Oak Ridge National Laboratory	10,000	20,242	Provisioning of Climate Data
Ning Ren	FM Global		0	FireFOAM Modeling of Roll Paper Fires and VOF Modeling of Sprinkler Atomization
Venkat E Tangirala	GE Global Research	5,000,000	30,658	Unsteady Combustion Processes in a Gas Turbine Combustor
Jin Yan	GE Global Research	47,000,000	3,862	Modelling of Combustion Dynamics in a Gas Turbine
Vaidyanathan Sankaran	United Technologies Research Center	1,000,000	501,109	Higher Fidelity Fire Simulations using Fire Dynamics Simulator
Eric Brown-Dymkoski	Space Exploration Technologies Corp	5,000,000	2,276,050	Scaling and Validation of Adaptive-Grid Turbulent Mixing Simulations
Josette Bellan	California Institute of Technology	5,000,000	617,382	Predictive Large-Eddy Simulation of Supercritical-Pressure Reactive Flows in the Cold Ignition Regime
Seung Hyun Kim	Ohio State University Research Foundation	3,000,000	0	Development of a Physics-Based Combustion Model for Engine Knock Prediction
Scott Michael Schnobrich	3M Company	0	0	Cobra
Siddhartha Banerjee	Pinnacle Engines Inc.	4,000,000	8	Development of Opposed-Piston Variable Compression Ratio Engine for Automotive Applications using High-Performance Computational Combustion Methodologies
Karl Virgil Meredith	FM Global	0	0	Fire Suppression Modeling
Dario Alfe	University College London	5,000,000	5,512,739	Non-Covalent Bonding in Complex Molecular Systems with Quantum Monte Carlo
Yevgeniy Puzyrev	Vanderbilt University	5,000,000	970	Flexural Phonons and Mechanical Properties of 2-D Materials

PI	Institution	Allocation	Utilization	Project Name
Shiwei Zhang	College of William & Mary	6,000,000	740,528	Quantum Many-Body Computations of Strongly Correlated Systems
Panchapakesan Ganesh	Oak Ridge National Laboratory	10,500,000	5,798,856	Data Driven Discovery by Design of Energy Materials
Timo Thonhauser	Wake Forest University	1,000,000	0	Catalysis and Diffusion in Metal Organic Frameworks
Andreas Glatz	Argonne National Laboratory	1,000,000	1,068,647	OSCon
Garnet Kin-Lic Chan	Princeton University	4,000,000	374	High Performance Quantum Chemistry and Tensor Networks
Alexander Alexeev	Georgia Institute of Technology	1,000,000	1,421,400	Mesoscale Modeling of Suspensions of Synthetic Microswimmers
Jamison Daniel, Michael Matheson	Oak Ridge National Laboratory	1,000,000	825	Cluster-based Visualization of Tera- and Peta-Scale Datasets
Neena Imam	Oak Ridge National Laboratory	5,000,000	4,326,922	Durmstrang
David Ronald Pugmire	Oak Ridge National Laboratory	5,000,000	14,994	SDAV
Sergey Panitkin	Brookhaven National Laboratory	0	73,756,883	Next Generation Workload Management System
Fernanda Schafer Foertter	Oak Ridge National Laboratory	250,000	0	Developing Scalable Heterogeneous Computing Training Code Examples
Suleyman Kocak	Dassault Systemes Simulia Corp	5,000,000	190,329	Performance, Capacity and Scalability Assessment & Improvements for Realistic Simulation Software - Abaqus/Standard and Abaqus/Explicit
John Franklin Harney	Oak Ridge National Laboratory	10,500,000	95	Data Intensive Science Incubators
Cyril Zeller	NVIDIA Corporation	13,500,000	0	CoDesign
Judith C Hill	Oak Ridge National Laboratory	40,000,000	0	Computational Partnerships
Scott A Klasky, Norbert Podhorszki	Oak Ridge National Laboratory	25,000,000	25,910,730	ADIOS
Jakub Kurzak	The University of Tennessee	200,000	0	Bench Testing Environment for Automated Software Tuning
Michael Joseph Brim	Oak Ridge National Laboratory	6,000,000	5,236,525	Lustre-Vision
Suzanne Theresa Parete-Koon	Oak Ridge National Laboratory	0	0	Data Transfer Working Group
Robert M Patton	Oak Ridge National Laboratory	3,500,000	3,008,083	Scalable Deep Learning Systems for Exascale Data Analysis
Joshua Ryan New	Oak Ridge National Laboratory	1,878,780	121,220	Big Data Mining for Building Analytics

PI	Institution	Allocation	Utilization	Project Name
Shantenu Jha	Rutgers, the State University of New Jersey	500,000	19	AIMES
Beverly Ann Sanders	University of Florida	3,000,000	1,908,193	Predicting and improving the performance of the Super Instruction Architecture
Mark Christopher Miller	Lawrence Livermore National Laboratory	100,000	181	FASTMath Installation and Portable Performance Testing
Lubomir Riha	IT4Innovations National Supercomputing Center	2,700,000	1,357,297	ESPRESO - ExaScale PaRallel FETI SOlver
Kalyan Perumalla	Oak Ridge National Laboratory	3,000,000	666,496	Cloning for Exascale
Oded Schwartz	The Hebrew University of Jerusalem	60,000	50,787	Algorithmic Linear Algebra
Manuel Arenaz	Appentra Solutions S. L.	50,000	156	Porting Parallware Tool to Large HPC Installations Including Titan
Matthew Dearing Wolf	Oak Ridge National Laboratory	2,000,000	367,080	In Situ Analytics Infrastructures
Jing Gong	Kungliga Tekniska Högskolan (KTH Royal Institute of Technology)	1,000,000	508,560	NekBone with GPU-Direct Communication and Optimized OpenACC Directives
Eric T. Phipps	Sandia National Laboratories	3,000,000	1,384,227	EQUINOX Embedded Uncertainty Quantification
Catherine Schuman	Oak Ridge National Laboratory	2,000,000	1,192,421	Scalable Evolutionary Optimization for Designing Networks
David Gutzwiller	Numeca USA, Inc.	2,000,000	966,359	Preparation of the NUMECA FINE/Open CFD Solver for Leadership Computing: Parallel I/O, Profiling, and Acceleration
Eduardo Francisco D'Azevedo	Oak Ridge National Laboratory	50,000	0	Investigation in Task-based Run-time frameworks for Geosciences Application
Vivek Sarkar	Rice University	1,000,000	635,919	Unified Portable Programming System for Integrating Task Parallelism, Accelerator Parallelism, and Message Passing on Exascale Systems
Alan Gray	University of Edinburgh	3,000,000	432,640	Performance Portability for Large-scale Scientific Applications
Ramakrishnan Kannan	Oak Ridge National Laboratory	3,000,000	44,777	Mini-Apps for Big Data
Chad Allen Steed	Oak Ridge National Laboratory	1,000,000	6,285	In Situ Visual Analytics for Extreme Scale Simulation Science
Jeffrey Young	Georgia Institute of Technology	800,000	140	Evaluation and Porting of ORNL Mini-apps to Future Directive-Based Languages and Runtimes
Cory Hauck	Oak Ridge National Laboratory	500,000	652,058	Diagnostics for Data Compression at Scale

PI	Institution	Allocation	Utilization	Project Name
Antonino Tumeo	Pacific Northwest National Laboratory	3,000,000	623,469	Scaling Graph Analytics on Heterogeneous Supercomputers
Jeff Clune	University of Wyoming	120,000	0	Improving Artificial Intelligence by Studying and Harnessing the Evolution of Structurally Organized Neural Networks
Terry Ray Jones	Oak Ridge National Laboratory	1,500,000	2,367	UNITY: Unified Memory & Storage Space
Alex Aiken	Stanford University	5,000,000	0	Scaling Legion S3D
Audris Mockus	The University of Tennessee	0	0	Software Supply Chains
Manjunath Gorentla Venkata	Oak Ridge National Laboratory	1,000,000	852,760	SharP: Shared Datastructure Programming Paradigm for Extreme Scale Systems
Abhinav Vishnu	Pacific Northwest National Laboratory	1,000,000	207	Machine Learning on Extreme Scale GPU systems
Michael Joseph Brim	Oak Ridge National Laboratory	5,000,000	0	ExCDS (Exascale Computing Data Service)
Edmon Begoli	Oak Ridge National Laboratory	0	0	Cloud Cover Study
Sherman Kisner	High Performance Imaging	240,000	199	Scalable parallelization for high performance tomography
Audris Mockus	The University of Tennessee	0	0	Fingerprinting Online Users using Doc2Vec Model model and Bayesian Networks
Shantenu Jha	Rutgers, the State University of New Jersey	2,000,000	295,382	RADICAL RESEARCH
John Cavazos	University of Delaware	1,000,000	0	Large-Scale Distributed and Deep Learning of Structured Graph Data for Real-Time Program Analysis and Characterization
Jeffrey Scott Chase	Duke University	2,500,000	774,822	Understanding File System Performance in Production Supercomputers
Misun Min	Argonne National Laboratory	1,500,000	0	CEED
Jeremy Travis Johnston	Oak Ridge National Laboratory	0	0	Surrogate Based Modeling for Deep Learning Hyper-parameter Optimization
Arvind Ramanathan	Oak Ridge National Laboratory	1,500,000	0	ECP Cancer Distributed Learning Environment
Steven Shannon	North Carolina State University	1,470,000	0	Particle-In-Cell Simulation of Radio Frequency Field Structure Near Plasma Facing Antenna Components
William M Tang	Princeton University	5,000,000	1,137,950	Extreme Scale PIC Research on Advanced Architectures
Judith C Hill	Oak Ridge National Laboratory	1,500,000	55,936	Workflow Optimization and Processing of Complex Datasets for Off-site Fusion Energy Research
Frank Jenko	University of California - Los Angeles	5,000,000	741,487	GENE

PI	Institution	Allocation	Utilization	Project Name
Zhihong Lin	University of California - Irvine	5,000,000	10,697,486	Fusion SciDAC GSEP Center: Gyrokinetic Simulation of Energetic Particle Turbulence and Transport
Jeff Candy	General Atomics	3,000,000	1,992,280	CGYRO Multicore and Accelerator Optimization
David Lindsay Green	Oak Ridge National Laboratory	1,000,000	115,820	Proto-MPEX Simulation
John Gary Shaw	University of Rochester	1,000,000	759,476	Laser Plasma Simulation Environment
Ales Necas	Tri Alpha Energy Inc.	5,000,000	67	Fusion Reactivity Enhancement: Numerical Beam-Driven FRC
Charlotte Barbier	Oak Ridge National Laboratory	1,000,000	0	Large Scale Hydraulic Fracture Simulations
Henri Francois Calandra	Total E&P	7,500,000	7,895,017	Advanced Computing for Geoscience Applications
Dilip Reddy Patlolla	Oak Ridge National Laboratory	2,000,000	1,041,580	GPU Accelerated Settlement Detection
Terry Ray Jones	Oak Ridge National Laboratory	5,000,000	726,722	Foresight
David Trebotich	Lawrence Berkeley National Laboratory	1,000,000	0	Mesoscale Simulation of Subsurface Fractured Materials
George Serban Constantinescu	University of Iowa	2,000,000	2,452,445	A Critical Assessment of the Performance of a Fully 3-D Nonhydrostatic Flow Solver to Simulate Flood Propagation in Natural Environments and Improve Efficiency of Flood Protection Measures
Thomas J. LeCompte	Argonne National Laboratory	1,000,000	0	Extending Highly Parallel HEP Codes to Titan
Gabriel Nathan Perdue	Fermi National Accelerator Laboratory	1,000,000	61,777	MACHINE Learning for MINERvA
Jason Newby	Oak Ridge National Laboratory	0	0	The COHERENT Experiment at the SNS [Spallation Neutron Source]
Constantia Alexandrou	University of Cyprus	1,000,000	1,214,520	Parton Distribution Functions From First Principles
Mark Sansom	University of Oxford	1,000,000	0	The Influenza A Frontier: Simulating the Interactions of a Full-Scale Enveloped Virus with a Complex Host Cell Membrane
Arvind Ramanathan	Oak Ridge National Laboratory	7,000,000	6,555,419	Integrating Neutron Scattering Experiments with Atomistic Simulations to Characterize Biophysical Mechanisms of Intrinsically Disordered Proteins
Liqun Zhang	Tennessee Technological University	6,000,000	7,889,140	Molecular Dynamics Simulations to Investigate the Structure, Dynamics and Functional Properties of Human Beta Defensin Type 3
Amanda Randles	Duke University	2,000,000	1,892	Massively Parallel Hemodynamic Simulation
Leonid V Zhigilei	University of Virginia	5,000,000	0	Atomistic Simulations of Laser Interactions with Metals
Byung Hoon Park	Oak Ridge National Laboratory	5,000,000	1,710,582	Accelerating Materials Modeling Loop of Leadership Computing and Spallation Neutron Source

PI	Institution	Allocation	Utilization	Project Name
Jerzy Bernholc	North Carolina State University	3,148,800	1,919,129	High Performance Simulations of Electron Spin Distribution and Dynamics in Low-Dimensional Materials
Yong Han	University of Utah	1,000,000	23	Catalysis and Diffusion in Metal Organic Frameworks
Michael J. Demkowicz	Texas A & M University - College Station	3,000,000	2,707,637	First-Principles Molecular Dynamics Modeling of Knock-on Damage in Amorphous Silicon Oxycarbide
Maarten De Jong	University of California - Berkeley	2,000,000	4,270,875	Charting the Complete Elastic Properties of All Known Inorganic Compounds: An Exploratory Study
Dongwon Shin	Oak Ridge Institute for Science and Education (ORISE)	8,000,000	6,896,692	High Performance Cast Aluminum Alloys for Next Generation Passenger Vehicle Engines
Trung Dac Nguyen	Vietnam Academy of Science and Technology	1,000,000	489,589	Implementing and Optimizing GPU-accelerated Molecular Dynamics Models
Venkat Padmanabhan	Indian Institute of Technology Kharagpur	500,000	705,286	Active Layer Morphology for Improving Efficiency of Organic Photovoltaics
Balasubramaniam Radhakrishnan	Oak Ridge National Laboratory	1,000,000	274,368	Micromagnetics Simulations for Magnetic Materials in Electric Motors
David Michael Eike	Proctor & Gamble Company	2,000,000	464,910	Large Size- and Time-Scale Investigation of Surfactant Aggregation and Structure Formation
Shanmugavelayut Kamakshi Sundaram	Alfred University	1,000,000	814,132	Structural and Mechanical Property Optimization of Photomultiplier Tube (PMT) Glasses for Neutrino Detection using MD Simulations
Efthimios Kaxiras	Harvard University	0	0	Catalyst Screening and Machine Learning for Catalytic Oxidation of Methanol on Metal Substrates
Yangyang Wang	Oak Ridge National Laboratory	3,000,000	1,103,951	Elucidating the Influence of Reversible Non-Covalent Interactions on Dynamic Properties for Rational Design of Soft Materials
Eric Jerome Lingerfelt	Oak Ridge National Laboratory	4,000,000	3,152,679	CADES/OLCF Computational Workflows for Materials Science
Hendrik Heinz	University of Colorado Boulder	4,000,000	284,006	Understanding and Designing Alloy Nanocatalysts in Atomic Resolution
Saurabh Ghosh	Vanderbilt University	2,000,000	1,066,933	Engineering Multifunctionality in Oxide Heterostructures
Jeffrey C Grossman	Robert Bosch LLC	6,000,000	2,820,969	Accurate Simulation of Li-Rich Layered Oxide Materials
Maximilian Amsler	Northwestern University	500,000	779,276	High-Throughput Calculation of Materials Properties at Finite-Temperature

PI	Institution	Allocation	Utilization	Project Name
Srikanth Allu	Oak Ridge National Laboratory	4,000,000	90,352	A Large Scale Cycling Studies for Lithium-Ion Battery Packs
Michael J. Demkowicz	Texas A & M University - College Station	0	0	Phase-field Modeling of He growth, coalescence, and stability at Cu-V interfaces
Markus Mueller	Georg-August-Universität Göttingen	1,500,000	12,578	SOMA: SOft coarse grained Monte-carlo Acceleration
Peter Thomas Cummings	Vanderbilt University	4,000,000	43	Materials Genome Screening of Soft Materials
Steven Hartman	Oak Ridge National Laboratory	0	0	ORNL Neutron Sciences Experiment Data
Adolfo Eguiluz	The University of Tennessee	2,400,000	0	GW many-body solver for ab initio multi-orbital Hamiltonians for correlated materials: Fundamental new physics on next-generation computational platforms
Weiwei Sun	Oak Ridge National Laboratory	200,000	0	Electronic structure of defects in two dimensional materials
Jeetain Mittal	Lehigh University	1,000,000	0	Biomolecular assembly processes in the design of novel functional materials
Georgia Tourassi	Oak Ridge National Laboratory	1,000,000	257,321	High Performance Infodemiology Architecture for Knowledge Extraction from Unstructured Health Text
Robert Chase Cockrell	Iowa State University	4,000,000	0	Examining the pathogenesis of pouchitis using a multiscale model of the intestinal mucosa: Spatially Explicit General-purpose Model of Enteric Tissue
Georgia Tourassi	Oak Ridge National Laboratory	4,000,000	260,897	Population Information Integration, Analysis, and Modeling for Comprehensive, Scalable, and Cost Effective Cancer Surveillance
Predrag S Krstic	State University of New York at Stony Brook	6,600,000	4,837,894	Gas-Liquid-Solid Interfaces for Energy Applications
Uwe Thumm	Kansas State University	500,000	0	Ab-initio calculations for laser-assisted XUV double ionization of helium
Klaus Bartschat	Drake University	5,000,000	0	Computational Studies of the Interaction of Time-Dependent Electromagnetic Fields and Charged Particles with Atoms and Molecules
Bamin Khomami	The University of Tennessee	1,000,000	0	Multiscale Modeling
Massoud Kaviany	University of Michigan	500,000	1,959,423	Human Thermosensation: Response of TRPV1 Domains to Temperature
Mathieu Luisier	ETH Zurich	17,000,000	0	Breaking the 10 PFLOPS Barrier with Ab-Initio Quantum Transport Simulations

PI	Institution	Allocation	Utilization	Project Name
Sefa Dag	GlobalFoundries	2,000,000	1,694,142	Computational Design and Optimization of Nano Device Architectures Based on Ab Initio Quantum Approaches
Mathieu Luisier	ETH Zurich	6,000,000	5,339,504	Contact Resistance in 2-D Semiconductors: Code Scalability and Preliminary Results
Rajan Pandey	GlobalFoundries	3,000,000	529	Ab-Initio Materials Modeling for Next Generation CMOS Devices
Christophe Calvin	CEA Saclay	1,500,000	0	NMC: New Monte Carlo
Steven Paul Hamilton	Oak Ridge National Laboratory	5,000,000	14,139,776	Radiation Transport Using Rayleigh Quotient Iteration with Multigrid in Energy Preconditioning
William Bryant Bird	Tennessee Valley Authority	1,000,000	0	TVA CASL Test Stand - Evaluation of Lower Plenum Flow Anomaly Using VERA/Hydra-TH
Guillermo Ivan Maldonaldo	The University of Tennessee	1,200,000	0	Fuel/Core Heavy Metal Design Optimization to Improve Performance/Safety of LSCR
Robert Lindsay Varner Jr	Oak Ridge National Laboratory	0	0	Majorama Demonstrator Secondary Data Archive
Jirina Rikovska Stone	The University of Tennessee	8,000,000	0	Phase Transitions in High Density Matter in Neutron Stars and Supernovae
Aurel Bulgac	University of Washington	25,000,000	0	Fission Fragment Yields within Time-Dependent Density Functional Theory
Charles Horowitz	Indiana University	2,000,000	1,341,474	Nuclear Pasta
Rajan Gupta	Los Alamos National Laboratory	3,000,000	0	Probing Novel Physics using Nucleon Matrix Elements
Andrey Vladimirovich Zakirov	Keldysh Institute of Applied Mathematics	1,000,000	0	High-Performance Wave Modeling in Nanooptics
Hanchuan Peng	Allen Institute for Brain Science	2,000,000	3,755,637	BigNeuron
Fabien Jonathan Delalondre	Ecole Polythechnique Federale de Lausanne (EPFL)	3,800,000	1	Brain Simulation Modeling
Misun Min	Argonne National Laboratory	2,500,000	0	Nek-HOM (Codes for High Order Methods)
Mark Alexander Jack	Florida Agricultural and Mechanical University	2,000,000	0	QRing: A scalable parallel software tool for quantum transport simulations in carbon nanoring devices based on NEGF formalism and a parallel C++ / MPI / PETSc algorithm
Alex Travesset	Ames Laboratory	45,000	1,065	Crystal Structure Prediction in Nanoparticle Superlattices
Amitava Bhattacharjee	Princeton University	10,000,000	0	Magnetic Reconnection and Laboratory Astrophysics with Laser-Produced Plasmas

PI	Institution	Allocation	Utilization	Project Name
Gabriel Kotliar	Rutgers, the State University of New Jersey	1,000,000	198,850	Application of MQSGW+DMFT to Delta Plutonium
Peter Zaspel	University of Basel	2,000,000	119,843	Scalable Numerical Algebra for UQ and CFD
Homer Heath Dewey	Orbital ATK	240,000	77,546	HERO Scaling Study
Helmut Katzgraber	Texas A&M University - Commerce	300,000	1	Benchmarking and Development of Quantum Annealing Architectures
Balakrishnan Naduvalath	University of Nevada - Las Vegas	500,000	0	Quantum Dynamics of Ultracold Chemical Reactions
Abhishek Kumar	Oak Ridge National Laboratory	0	0	Controlling the Movement of Leidenfrost Droplet: a Molecular Dynamics Study
Shiwei Zhang	College of William & Mary	4,000,000	671,667	Quantum Many-Body Computations of Strongly Correlated Systems
Gregory Gershom Howes	University of Iowa	0	0	Visualization and Analysis of High-Dimensional Kinetic Plasma Physics Data
Carlos Federico Lopez	Vanderbilt University	4,000,000	2,391,497	Reducing Unidentifiability in Cell-Signaling Network Models Through Multidimensional Analysis and Molecular Simulation
George Em Karniadakis	Brown University	0	0	Multiscale simulation of human pathologies
Jacopo Buongiorno	Massachusetts Institute of Technology	1,000,000	0	Development of a New, High-Fidelity, CFD-Informed Closure Relation for Taylor Bubble Velocity in Slug Flow in Pipes with Horizontal to Vertical Inclinations
Duane Lee Rosenberg	SciTec, Inc.	2,000,000	0	Small-Scale Statistics and Intermittency In Rotating Strongly-Stratified Turbulence: Verification and Connection to Bolgiano-Obokhov Phenomenology
Olivier Desjardins	Cornell University	5,000,000	1,398,347	Large-scale Computation of Particle-Laden and Multiphase Turbulent Flows
Sumanta Acharya	University of Memphis	1,707,110	2,741,279	Thermally Effective and Efficient Cooling Technologies for Advanced Gas Turbine Systems
John Schaefer	Boeing Company	0	0	Uncertainty Quantification and Sensitivity Analysis of Turbulence Model Coefficients for the Common Research Model
Said Elghobashi	University of California - Irvine	8,000,000	3,923,054	Direct Numerical Simulation of Fully-Resolved Droplets in Turbulent Flows
Rafael Omar Tinoco	University of Illinois at Urbana-Champaign	50,000	0	Vegetation-Sediment-Flow interactions: Direct numerical simulation of turbulent oscillatory flow and sediment transport on aquatic ecosystems

PI	Institution	Allocation	Utilization	Project Name
Pui Kuen Yeung	Georgia Institute of Technology	5,000,000	3,210,926	OpenACC Optimization for Direct Numerical Simulation of Turbulent Mixing Using a Hybrid Pseudo-Spectral and Compact Finite Difference Algorithm
Jesse Capecelatro	University of Michigan	500,000	8,586	Demonstration of Noise Suppression by Water Injection in High-Speed Flows with Direct Numerical Simulation
Joseph D Smith	University of Missouri	0	0	Numerical Investigation of flow instabilities inside a Westinghouse SMR
Rachelle Lea Speth	Boeing Company	0	0	Noise Control Concepts for the Main Landing Gear of a Commercial Transport Aircraft
Chad Michael Winkler	Boeing Company	3,000,000	62	Large Eddy Simulations of Jet Noise and S-duct Diffusers