ORNL/SPR-2018/XX

US Department of Energy, Office of Science High Performance Computing Facility Operational Assessment 2017 Oak Ridge Leadership Computing Facility

J. Paul Abston
Ryan M. Adamson
Ashley D. Barker
David E. Bernholdt
Katie L. Bethea
Kevin S. Bivens
Arthur S. Bland
James J. Hack
Rachel Harken
Jason J. Hill
Judith C. Hill
Jonathan Hines
Katie Jones
Bronson Messer
Stephen T. McNally
Mallikarjun Shankar
Kevin G. Thach
Suzy Tichenor
Sudharshan S. Vazhkudai
Jack C. Wells

April 2018

**Approved for public release.
Distribution is unlimited.**

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

ORNL/SPR-2018/XXX

Oak Ridge Leadership Computing Facility

HIGH PERFORMANCE COMPUTING FACILITY OPERATIONAL ASSESSMENT 2017 OAK RIDGE LEADERSHIP COMPUTING FACILITY

J. Paul Abston
Ryan M. Adamson
Ashley D. Barker
David E. Bernholdt
Katie L. Bethea
Kevin S. Bivens
Arthur S. Bland
James J. Hack
Rachel Harken
Jason J. Hill
Judith C. Hill
Jonathan Hines
Katie Jones
Bronson Messer
Stephen T. McNally
Mallikarjun Shankar
Kevin G. Thach
Suzy Tichenor
Sudharshan S. Vazhkudai
Jack C. Wells

Date Published: April 2018

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

LIST OF FIGURES

LIST OF TABLES

ACRONYMS

ACCEL	Accelerating Competitiveness through Computational Excellence
ALCC	ASCR Leadership Computing Challenge
ALCF	Argonne Leadership Computing Facility
API	Application Programming Interface
ARM	Atmospheric Radiation Monitoring
ARM	Atmospheric Radiation Measurement
ASCR	Advanced Scientific Computing Research
BER	Office of Science Biological and Environmental Research
BI	Business Intelligence
CADES	Compute and Data Environment for Science
CFD	Computational Fluid Dynamics
CRM	Customer Relationship Management
CSGF	Computational Science Graduate Fellowship
CY	Calendar Year
DART	Days Away, Restricted, or Transferred
DD	Director's Discretionary
Dev	UA Development
DFT	Density Functional Theory
DOE	US Department of Energy
DOI	Digital Object Identifier
DTN	Data Transfer Node
DVS	Data Virtualization Service
ECP	Exascale Computing Project
Esnet	Energy Sciences Network
HPC	High Performance Computing
HPE	Hewlett-Packard Enterprise
HPSS	High-Performance Storage System
INCITE	Innovative and Novel Computational Impact on Theory and Experiment
JIF	Journal Impact Factor
LASSO	LES ARM Symbiotic Simulation and Observation
LBNL	Lawrence Berkeley National Laboratory
LC	Livermore Computing
LLNL	Lawrence Livermore National Laboratory
LOSA	Battelle Laboratory Operations Supervisor Academy
LSMS	Locally Self-Consistent Multiple Scattering Code
MEP	Mechanical Energy Plant
MPI	Message Passing Interface
MTTF	Mean Time to Failure
MTTI	Mean Time to Interrupt
NAM	Not a Metric
NCCS	National Center for Computational Sciences
NCSA	National Center for Supercomputing Applications
NERSC	National Energy Research Supercomputing Center
OA	Overall Availability
OAR	Operational Assessment Report
OLCF	Oak Ridge Leadership Computing Facility
OMB	Office of Management and Budget
ORAU	Oak Ridge Associated Universities

ORISE	Oak Ridge Institute for Science Education
ORNL	Oak Ridge National Laboratory
OST	Object Storage Target
OUG	OLCF User Group
PDSW-DISC	Parallel Data Storage & Data Intensive Scalable Computing Systems
PFL	Progressive File Layout
PI	Principal Investigator
PKPass	Public Key-Based Password Distribution System
PLC	Programmable Logic Controller
PNNL	Pacific Northwest National Laboratory
QCD	Quantum Chromodynamics
RATS	Resource Allocation Tracking System
RATS	Resource and Allocation Tracking System
RFP	Request for Proposal
RSS	Research Safety Summaries
RT	Request Tracker
RUC	Resource Utilization Council
SA	Scheduled Availability
SBMS	Standards-Based Management System
SC	Office of Science
SciComp	Scientific Computing Group
SciDAC	Scientific Discovery through Advanced Computing
SDG	Scientific Data Group
SIEM	Security Information And Event Management
SNL	Sandia National Laboratories
SU	System Utilization
SVE	Scalable Vector Extensions
TPS	Thermal Protection System
UA	User Assistance
UAO	User Assistance and Outreach
WiC	Women in Computing

HIGH PERFORMANCE COMPUTING FACILITY
2017 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2018

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 computational ecosystems; 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) 2017 was filled with outstanding operational results and accomplishments: a very high rating from users on overall satisfaction for the fourth year in a row; the greatest number of core hours delivered to research projects, surpassing the previous record from CY 2016; the largest annual number of research publications since the deployment of Titan; and success in delivering on the allocation split 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 rates (overall and capability usage), represent the fulfillment of the promise of Titan: efficient facilitation of leadership-class computational applications. Table ES.1 presents a summary of the 2017 OLCF metric targets and the associated results.

The impact of OLCF's achievements is reflected in the accomplishments of OLCF users, with publications this year in such notable journals and publications as *Nature*, *Science*, *Reviews of Modern Physics*, *Nature Chemistry*, *Nature Physics*, *Nature Climate Change*, *Nano Letters*, *ACS Nano*, *Journal of the American Chemical Society*, *Nature Communications*, *Proceedings of the National Academy of Sciences*, *Journal of Physical Chemistry Letters*, *Physical Review X*, *Physical Review Letters*, and *Nanoscale*. Crucial domain-specific discoveries facilitated by resources at the OLCF are described in the *High Performance Computing Facility Operational Assessment 2017 Oak Ridge Leadership Computing Facility* (OAR). For example, Titan enabled an all-atom understanding of P-glycoprotein, which pumps foreign substances—including anticancer drugs—out of cells, helping the drug discovery community identify a pathway that will allow anticancer drugs to do their job (Section 3.2.1).

The Titan system provides the largest extant heterogeneous architecture for conducting computational science research and remains the most powerful and productive system supporting open science in the United States. Usage is high, delivering on the promise of a system well suited for accelerated capability simulations for science. This success is a result of the extraordinary work of the OLCF staff in supporting the nation's leading HPC facility for the Department of Energy (DOE). The staff are pivotal to identifying, developing, and deploying the innovative processes and technologies that have helped the OLCF, its users, and other high performance computational facilities realize success. The facility's leadership in the area of data analysis and workflows continued in 2017 with further integration of the highly successful Compute and Data Environment for Science (CADES) user facility to the National Center for Computational Sciences (NCCS). This organizational move follows the reported announcement from the 2016 OLCF OAR and shows the commitment of the OLCF to bringing user facilities together to tackle scientific grand challenges across multiple scientific domains.

Table ES.1. 2017 OLCF metric summary

2017 metric description	CY 2017 target	CY 2017 actual
Overall OLCF score on the user survey will be 3.5/5.0 based on a statistically meaningful sample.	3.5	4.6
Time between Receipt of User Query (RT Ticket) and Center Response: 80% of OLCF problems will be addressed within 3 working days (72 hours) by either resolving the problem or informing the user how the problem will be resolved.	80%	93%
Scientific and Technological Research and Innovation—Demonstrate Leadership Computing: For the calendar year following a new system/upgrade, at least 30% of the consumed node hours will be from jobs requesting 20% or more of the available node. In subsequent years, at least 35% of consumed core hours will be from jobs requiring 20% or more of cores available to the users.	35%	59.8%
Scheduled Availability, TITAN: Sustain scheduled availability to users, measured as a percentage of maximum possible scheduled.	95%	99.4%
Overall Availability, TITAN: Sustain availability to users, measured as a percentage of maximum possible.	90%	98.1%
Overall Availability, External File System: Sustain availability to users, measured as a percentage of maximum possible.	90%	Atlas 1: 98.9% Atlas 2: 98.9%
Overall Availability, Archive Storage: Sustain availability to users, measured as a percentage of maximum possible.	90%	HPSS: 98.9%

Operating a world-class HPC facility that offers the breadth of resources available at the OLCF is not without its challenges. Having celebrated its fifth year in production in 2017, the Titan continued to experience deteriorating component reliability. The OLCF staff continued their ongoing efforts to address and correct the failures of aging components while providing a highly capable leadership environment. Specific innovations are discussed in detail in this report. Through continuous improvement and daily management of its systems, the OLCF observed a significant drop in the number and frequency of failures among its leadership-class jobs.

Effective operations of the OLCF play a key role in the scientific missions and accomplishments of its users. Building on the exemplary accomplishments in 2016, 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 2017, which denotes the period from January 1, 2017, to December 31, 2017, unless otherwise specified.

ES.1 COMMUNICATIONS WITH KEY STAKEHOLDERS

ES.1.1 Communication with the Program Office

The OLCF regularly communicates with the Advanced Scientific Computing Research (ASCR) Program Office through a series of regularly occurring events. These include weekly Integrated Project Team calls with the local DOE 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.

ES.1.2 Communication with the User Community

The OLCF's communications with users 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 (Sections 1.4.5–1.4.7).

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 third year in a row. Ninety-three percent of respondents (390) rated their overall satisfaction with communications from the OLCF as "satisfied" or "very satisfied." In addition, nearly all of the 346 users felt that the OLCF kept them well informed about changes (98%), events (99%), and current issues (97%). 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 meetings
- one-on-one interactions with liaisons and analysts
- social networking
- annual face-to-face OLCF User Meeting
- targeted training events (i.e., GPU Hackathons or New User Training)

ES.2 SUMMARY OF 2017 METRICS

In consultation with the DOE program manager and as proposed in the 2016 OAR, a series of metrics and targets were identified to assess the operational performance of the OLCF in CY 2017. The 2017 metrics, target values, and actual results as of December 31, 2017, are noted throughout this report and are summarized in Section 8. The OLCF exceeded all of its agreed-upon metric targets.

ES.3 RESPONSES TO RECOMMENDATIONS FROM THE 2015 OPERATIONAL ASSESSMENT REVIEW

The OLCF did not receive any recommendations from the 2016 OAR review but did receive high praises for running a customer-focused operation that surpasses its operational targets and goals. The reviewers provided specific comments, which are included below:

"OLCF is clearly a well-run and customer-focused facility that enables users to produce outstanding science and clearly demonstrates why OLCF is a national resource."

"The Oak Ridge Leadership Computing Facility continues to provide high-end high-performance computing and data systems and services to the HPC community and surpasses its operational target goals in all areas. The center achieved high user satisfaction and met its commitments to the INCITE, ALCC, and Director's Reserve programs. OCLF had strong engagements with industrial partners."

ES.4 OPERATIONAL REALIGNMENTS TO BETTER SERVE OLCF STAKEHOLDERS

In 2017, the OLCF executed a strategic realignment of the OLCF project and program leadership. This realignment placed the OLCF-4 Project Director, Buddy Bland, into a new role as the OLCF Program Manager. The OLCF Program Manager oversees all of the ongoing OLCF projects, which were OLCF-4 and OLCF-5 in CY 2017. As the OLCF Program Manager, Buddy manages all communications-related aspects of the OLCF with the DOE ASCR office, which includes both ongoing projects and facility operations. In recognition of his accomplishments as OLCF-4 Deputy Project Director, Justin Whitt was named the OLCF-5 Project Director. This realignment will ensure dedicated oversight for each active project within the OLCF and also formalizes the OLCF program oversight role. A new Deputy Director will be hired in CY 2018 for both the OLCF-4 and OLCF-5 projects.

As reported in the 2016 OLCF OAR, ORNL's Compute and Data Environment for Science (CADES) facility was moved into the NCCS to better align ORNL high performance computing efforts. In 2017, the NCCS worked to achieve tighter integration between the OLCF and CADES. This integration is paramount for providing next-generation systems that are capable of bridging modeling and simulation with experimental science. The CADES facility operates three distinct security enclaves for open, moderate, and HIPAA-required scientific discoveries. Staff from the OLCF and CADES worked diligently together to provide a closely coupled computing environment for research teams who would like to take their experimental data from another facility to the OLCF moderate production enclave for further simulation on Titan and then back into the CADES open enclave for analysis, workflow, and/or presentation through scientific portals. Having a flexible organization that operates within an open research environment, such as CADES, has proved beneficial to the OLCF and ORNL for a variety of reasons. See Section 4.1.4 for more information regarding the accomplishments of the 2017 integration.

In July of 2017, the OLCF hired Installation Manager Paul Abston to oversee the construction and system installation in the new Summit data center and to ensure the safety of all the center's activities. Given that the Summit data center was still under very active construction when the first system components were being delivered and installed, additional oversight on this data center seemed most appropriate. The hiring of an installation manager was first put in place at the Argonne Leadership Computing Facility (ALCF) and was recognized as a best practice by the OLCF. Given the number of people working in the same space and the complexity of the activities involved, the OLCF felt it best to select a professional with a strong background in safety. Paul successfully guided work within the new data center among a diverse group of subcontractors, ORNL craft, and OLCF staff to finish the year strong. This accomplishment is highlighted further in the Safety section.

HIGH PERFORMANCE COMPUTING FACILITY
2017 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2018

1. USER RESULTS

CHARGE QUESTION 1: Are the processes for supporting the users, resolving users' problems, and conducting outreach to the user population effective?

OLCF RESPONSE: Yes. In 2017, the Oak Ridge Leadership Computing Facility (OLCF) supported 1,336 users and 328 user 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 2017 was 4.6.

Overall ratings for the OLCF were positive; 96% 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 2017 with 93% 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 2017 survey was launched on October 04, 2017, and remained open for participation through November 12, 2017. The survey was sent to 1,115 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, 2017, and September 30, 2017. OLCF staff members were excluded from participation. A total of 448 users completed the survey, for an overall response rate of 40.2%. The results of the 2017 survey can be found on the OLCF website.

The effectiveness of the processes for supporting customers, resolving problems, and conducting outreach[1] are defined by the metrics in Table 1.1 and are assessed by the user survey.

[1] https://www.olcf.ornl.gov/olcf-media/center-reports/2017-outreach-survey/

Table 1.1. 2017 user result metrics summary

Metric description	2016 target	2016 actual	2017 target	2017 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	Results will show improvement in at least half 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.	Results will show improvement in at least half 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%	93%
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 2017. 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 2016		CY 2017	
	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 were 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 96% of users responding that they were 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 that 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	4.7/5.0
Overall satisfaction with support services	4.5/5.0	4.5/5.0	4.5/5.0	4.5/5.0
Overall satisfaction with compute resources	4.6/5.0	4.5/5.0	4.4/5.0	4.6/5.0
Overall satisfaction with data resources	4.5/5.0	4.4/5.0	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 2017 survey was 4.6/5.0, indicating that the OLCF exceeded the 2017 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?"

OLCF provides top-of-the-line compute resources to carry on our scientific mission. These resources are backed up by a highly capable and talented team providing training and technical support to the users to allow us to utilize OLCF resources efficiently and to place us in a good position for writing successful INCITE and ALCC proposals. The team is sharply focused, and is proactive, in proposing solutions for us to meet and exceed our scientific and computational milestones with the tools at their disposal and discretion, such as for example tweaking queue priorities etc. The OLCF truly values our success.

Cutting edge hardware, excellent software environment, excellent communication and documentation.

The support staff is always very responsive and are able to resolve issues quickly and effectively. The computing resources are world class. The training offered by OLCF has changed the trajectory of my career.

A comprehensive system with the right balance of hardware, software resources, very proficient and helpful staff. I've been a user for 7 years, and its standards have never dipped.

Access to OLCF's vast HPC resources has allowed me to perform research that would otherwise be impossible at my company. All of my interactions with the staff have been very positive, and I get the sense that everyone there is really personally invested in making sure users have a positive experience. The website and user guides are also very helpful.

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 2016 and 2017 surveys. However, based on feedback received from the 2016 User Survey in conjunction with other feedback channels, the OLCF took the following actions in 2017 to enhance the user experience at the OLCF.

- The OLCF upgraded the Data Transfer Node (DTN) cluster to help increase data transfer rates and improve throughput.

    - The available nodes were increased to 4 login, 8 batch, and 12 Globus transfer nodes.
    - Network connectivity was upgraded from 10GbE to 40GbE on each node.
    - The uplink to the Internet was upgraded to 100GbE.
- The OLCF now offers 24/7 password assistance via the User Support phone line.
- The OLCF completed the Constellation Digital Object Identifier (DOI) project, making it possible for researchers to obtain a 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 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 OLCF offered more frequent Getting Started webinars to better assist new users.
- The OLCF worked 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.
- The OLCF is in the process of deploying Singularity containers for Titan.

1.2.4 Assessing the Effectiveness of the OLCF User Survey

The survey was created by Oak Ridge Associated Universities' (ORAU'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 (ORISE) 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 4, 2017, and subsequent follow-up reminders were sent by Jack Wells, the National Center for Computational Sciences (NCCS) Director of Science, the OLCF Executive Board, and again by Ashley Barker on November 6, 2017.

The survey was advertised on the OLCF website and in the weekly communications via email 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 and users who have been at the center for 1–2 years. The OLCF saw growth in 2017 of users who have been using OLCF resources more than 2 years (Table 1.4).

Table 1.4. User survey participation

	2016 survey	**2017 survey**
Total number of respondents (Total percentage responding to survey)	369 (39%)	448 (40%)
New users (OLCF user <1 year)	32%	26%
OLCF user 1–2 years	25%	27%
OLCF user >2 years	43%	47%

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. ORAU 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, ranging from “very dissatisfied” (1) to “very satisfied” (5). These responses were close 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 how responses, needs, and preferences varied by types of OLCF users, close-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 very comparable across all three allocation programs, and the variations are statistically insignificant.

Table 1.5. Statistical analysis of survey results

Survey area	INCITE			ALCC			DD			ALL		
	Mean	Variance	Std. Deviation	Mean	Variance	Std. Deviation	Me an	Variance	Std. Deviation	Mean	Variance	Std. Deviation
Overall satisfaction with the OLCF	4.6	0.8	0.7	4.6	0.8	0.8	4.7	0.8	0.7	4.6	0.8	0.7
Overall satisfaction with support services	4.5	0.8	0.7	4.5	0.8	0.7	4.5	0.9	0.7	4.5	0.9	0.7
Overall satisfaction with user assistance team	4.6	0.8	0.7	4.6	0.7	0.6	4.6	0.8	0.6	4.6	0.8	0.6
Overall satisfaction with accounts services team	4.6	0.9	0.8	4.7	0.7	0.5	4.6	0.8	0.7	4.6	0.9	0.7
Overall satisfaction with compute resources	4.5	0.9	0.7	4.4	0.8	0.7	4.6	0.8	0.7	4.6	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 in which 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 the 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 2,100 tickets in response to user queries for CY 2017. The center exceeded the problem-resolution metric and responded to 93% of the queries within 3 business days (Table 1.6).

Table 1.6. Problem resolution metric summary

Survey Area	CY 2016		CY 2016	
	Target	Actual	Target	Actual
Percentage of problems addressed in 3 business days	80%	92%	80%	93%
Average of problem resolution ratings	3.5/5.0	4.6/5.0	3.5/5.0	4.6/5.0

Tickets are categorized by the most common types. The top three reported categories in 2017 were running jobs, account access, and job compilations (Figure 1.1).

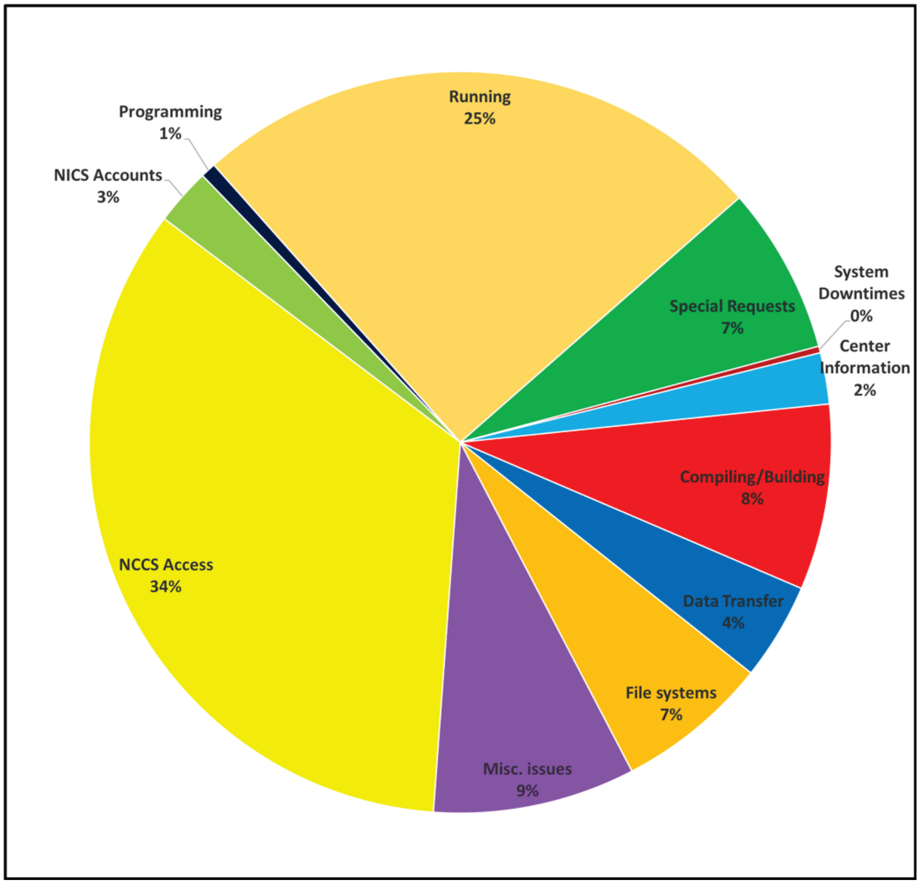

Figure 1.1. Categorization of help desk tickets.

1.4 USER SUPPORT AND OUTREACH

The Operational Assessment Report (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 a wide range of 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 2017 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 Containers

The OLCF deployed Singularity software containers on Titan, Summitdev, and Summit systems in CY 2017. This effort involved not only installing the container runtime environment but also developing methods for ensuring the performant use of vendor software, such as MPI, within the container. OLCF staff worked with Singularity developers to ensure that Singularity was compatible with Titan's unique software stack. A command line utility, container-builder, was created to allow users to build containers fully within the OLCF. A set of system-specific base container images is maintained, which come preloaded with system software required for the performant use of resources. These containers are made available through the container-builder utility.

1.4.2.2 OLCF Website

The OLCF's website serves not only as the home for news on scientific discovery occurring at the center but also as a gateway for OLCF users to access information critical to the success of their projects. The OLCF undertook a project in 2017 to redesign the website with an eye toward improving navigation, making the site more mobile friendly, updating the aesthetics, and adding a host of new features, many of which were specifically requested by OLCF users through previous years' surveys and through one-on-one interaction with users.

A beta version of the new OLCF website was launched in January 2018 and remained in place for 4 weeks. Making the beta site available for 4 weeks allowed users, media, and OLCF stakeholders to familiarize themselves with the new site before the current site was retired on February 15, 2018. Users were also invited to participate in a monthly conference call where OLCF staff walked through the new website and pointed out key new features.

Some of the new features include the following:

- Each staff member at the OLCF now has a dedicated page in the new Staff Directory on the OLCF website that provides details such as their biography, publication history, R&D projects, contact information, and more.
- New OLCF R&D project pages to give projects to which OLCF staff contribute more visibility.
- New OLCF INCITE and ALCC User project pages allow the OLCF to share more information about its flagship projects with the public.
- Enhanced mobile capability provides enhanced interaction from tablets and cell phones.
- Enhanced access controls allow OLCF staff to have access to and update their staff profile pages and individual project pages.

- Navigation within the OLCF system user guides has been improved.
- Information on current and future downtimes is now automated.

1.4.2.3 Resource and Allocation Tracking System (RATS)

The User Assistance (UA) team continued to make feature improvements to the OLCF's in-house customer relationship management (CRM) software, known as RATS CRM, in CY 2017. A total of 12 new major and minor versions (3.28.0–3.36.0) were deployed to production with minimal downtime. New feature highlights include business intelligence–based scheduling, future downtime tracking, multi-enclave support, staff profiles, and user/project application tools.

Business Intelligence-Based Scheduling**:** In a first for the center, UA Development (Dev) staff redesigned RATS CRM to use business intelligence data gathered from business intelligence reporting tools to automate continuous changes to batch scheduling configurations on Titan. This allows Titan's scheduler configuration to change automatically in real-time in response to user behavior. For example, when business intelligence (BI) reporting tools notice that a project begins running large, GPU-intensive jobs, that project's scheduler configuration records can be automatically updated so that the project's subsequent job submissions use queues or "quality of services" that are most suited to those types of jobs.

Historical and Future Downtime Tracking**:** In order to support improvements to the public website as well as decommission old infrastructure, UA Dev staff added integrated downtime tracking into RATS. Staff can now easily use the RATS CRM Application Programming Interface (API) to provide both historical and future downtimes, and downtime information can be provided in real-time to the public OLCF website.

Multi-Enclave Support**:** UA Dev staff redesigned RATS to support multiple coexisting security enclaves. Now staff can use the same project-based access controls used for the OLCF's Moderate enclave in other operational security enclaves, improving divisional efficiency and reducing operational complexity.

Staff Profiles**:** UA Dev staff extended RATS to include staff profiles (including images and PDFs), staff group profiles, and staff R&D profiles. This feature also includes fine-grained permissions that enable individual staff members, group leads, and administrative staff to create and manage their own profile data. This data can be easily exported to any number of other web services, including the public OLCF website.

User/Project Application Tools**:** Historically, account and project application submissions through the public OLCF website were entered by hand into RATS for project processing. UA Dev staff redesigned the application forms on the website and added functionality to pull submissions into RATS automatically where they can be viewed and processed by accounts staff using a number of time-saving, in-app tools.

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, and 1.4.3.3.

1.4.3 Scientific Liaison Collaborations

The following sections highlight specific collaborative areas where OLCF staff scientists partnered with INCITE research teams to maximize their productivity on the provided leadership class resources.

1.4.3.1 Turbulent Mixing at High Schmidt Number

The objective of the INCITE project "Turbulence mixing at high Schmidt number," led by P. K. Yeung from the Georgia Institute of Technology, is to study the turbulent mixing of scalar fields with very low diffusivity (i.e., high Schmidt number) in regimes that have been difficult to analyze in the past. One specific target is to reach a Schmidt number of 512, which is comparable to that for salinity transport in the ocean, while retaining the flow properties that characterize high Reynolds number turbulence. A high Schmidt number is challenging for both experiment and computation because fluctuations arise at scales much smaller than those of the velocity field.

Like many leadership-class applications, performance portability is critical for PSDNS, as the code is run on multiple platforms and must perform efficiently on each system. However, the current diversity of architectures and the developing state of relevant programming models have made it difficult for codes to achieve performance portability.

PSDNS achieves GPU acceleration on Titan in a performance portable way using OpenMP 4.5 compiler directives. The team recently produced the first production implementation they were aware of that combines the GPU offload capabilities introduced in OpenMP 4.5 with the specification's asynchronous tasking capabilities, making PSDNS the first production-level application run at the OLCF to combine the latest GPU offload features of OpenMP with its tasking capabilities. This constitutes a major step forward beyond using proprietary programming models for accelerated hardware to instead use a performance portable solution applicable across many platforms.

OLCF computational scientist Wayne Joubert, along with Oscar Hernandez of ORNL's Computer Science Research Group and Jeff Larkin of NVIDIA, assisted Yeung's team in porting PSDNS to the OpenMP 4.5 performance portable programming model. This involved helping the PSDNS developers conform their code to the OpenMP 4.5 standard, working through compiler issues encountered, measuring and optimizing code performance, and subsequently taking the lessons learned from this work back to the standards bodies and other users. This distillation and dissemination of their experience have already benefited the broader community. The resulting work has contributed to the forthcoming OpenMP 5.0 specification with additional features and capabilities to better support this "asynchronous offload" programming style. With encouragement and assistance from Hernandez, the approach has also been featured in talks[2,3] and tutorials.[4]

1.4.3.2 Determining the Physics of Galactic Winds

The INCITE project "Revealing the Physics of Galactic Winds with Petascale GPU Simulations" aims to understand the role of galactic-scale winds in the formation and evolution of galaxies.

These winds are streams of high-speed particles driven by the energetic feedback processes associated with supernovae and the stellar winds of massive stars. Galactic winds regulate baryonic content, star formation rates, and stellar masses of galaxies. This project, led by PI Brant Robertson of the University of California Santa Cruz and co-PI Evan Schneider of Princeton University, tackles this prime theoretical challenge, which was highlighted in the report *New Worlds, New Horizons in Astronomy and Astrophysics*, by the National Research Council and released in 2010.

---

[2] M. P. Clay, D. Buaria, and P. K. Yeung, "Improving Scalability and Accelerating Petascale Turbulence Simulations Using OpenMP," OpenMP Developers Conference 2017, Stony Brook University. Slides:https://openmpcon.org/wp-content/uploads/openmpcon2017/Day1-Session1-Clay.pdf

[3] M. P. Clay, D. Buaria, and P. K. Yeung, "OpenMP 4.5 Acceleration for Turbulence Simulations on GPUs." Scientific Computing in Times of MPI+X: Looking at Multiple "X" with Regards to Performance and Portability Symposia at Platform for Advanced Scientific Computing (PASC) 2018, July 2–4, 2018. (Accepted talk).

[4] Barbara Chapman, Oscar Hernandez, Yun (Helen) He, Martin Kong, and Geoffroy Vallee, "MPI+OpenMP tutorial," February 9th, 2018. 2nd ECP meeting, Knoxville, TN. Slides:https://docs.google.com/presentation/d/1LzhXlVFlr_WdMRJ1UnZgquTSIJkXyhWqlLl_5X-3eaE/edit?usp=sharing

The project uses the code Cholla to perform numerical simulations of galactic winds by leveraging the massively parallel GPUs on Titan to provide computational power unmatched by earlier simulation codes. The Cholla code is unique among similar astrophysical codes in that it performs all of its calculations on GPUs at scale. This feature allows it to perform simulations with numerical fidelity that can resolve the multiphase structure of the galactic outflows to reveal how mass and energy are entrained into a hot wind.

Timely progress in large simulations requires the code to be able to handle the challenges that become more common in large simulations, such as inadvertent node failures. To this end, OLCF computational scientist Reuben Budiardja developed a checkpoint handling and automated restart capability to enable continued running in the event of a node failure. He also identified optimization opportunities in the I/O portion of the code and is working with the team to implement these improvements. Because large simulations generate prodigious amounts of data, better I/O management is not only beneficial during the execution of the simulations but also ensures that data analysis can be done efficiently to provide scientific insight. Budiardja also directly supported the team in their data analysis efforts by providing software and tools on the OLCF data analysis and visualization cluster Rhea.

1.4.3.3 Statistical Physics of Materials

Many materials properties studies require first-principles calculations of many configurations. These simulations are sometimes individually not appropriate for high performance computing, but combined they require substantial resources, placing them in the leadership computing category. The INCITE project “First principles based statistical physics of alloys and functional materials” investigates the effects of chemical disorder in materials using first-principles calculations. The calculations solve the Schrödinger equation for the electrons inside these solids using Density Functional Theory (DFT). Disorder plays an important role in understanding the properties of real materials, and the computational treatment of disorder requires both large simulation cells and many representative samples.

The team has calculated the magnetic properties of a FePt nanoparticle based on measured atom positions and chemical composition determined experimentally with atomic resolution. They calculate the spatially resolved spin and orbital magnetic moments, as well as the magneto-crystalline anisotropy of this particle, taking into account the deviation from the ideal FePt crystal structure. The calculation of the local anisotropies required the bundling of *thousands* of DFT calculations, each only requiring a few Titan nodes. To be able to perform these calculations efficiently on OLCF resources, OLCF computational scientist M. Eisenbach bundled these VASP calculations using the wraprun capability developed at the OLCF, and presented in previous OLCF OARs, for this purpose. In conjunction with the large-scale Locally Self-consistent Multiple Scattering code (LSMS) that Eisenbach codeveloped and optimized for OLCF resources, the team could efficiently analyze the magnetic properties of a real nanoparticle using OLCF’s Titan.

1.4.4 OLCF User Group and Executive Board

The OLCF User Group (OUG) is open to all PIs and users on approved OLCF user projects and will remain so for 3 years following the conclusion of their OLCF project. The OUG meets once a month via BlueJeans webinar to discuss OLCF news, resources, policies, and timely HPC tutorials and techniques. The OUG executive board represents the OUG and is made up of 10 users who give feedback to OLCF staff on its services and represent the OLCF user community. Members of the executive board meet 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. Elections are held yearly during the annual OLCF Users Meeting to select three new board members to replace those members who have completed their 3 year appointments. A total of 119 users voted in the 2017 election and selected James McClure, Abhi Singharoy, and Mike Zingale to serve on the OUG executive board for 3 year terms. The current board chair is Balint Joo. The board elected Joe Oefelein to serve as the vice chair this year, and he will automatically become the chair for the 2018–2019 cycle.

The OLCF hosted a total of 10 monthly conference calls in 2017, with a total of 345 attendees, and the annual OLCF Users Meeting, which was attended by 151 participants. More information about the OUG, including a list of the executive board members, can be found at https://www.olcf.ornl.gov/about-olcf/oug/.

1.4.5 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 or jointly collaborated on 10 user conference calls, four week-long hackathons, three INCITE proposal writing webinars, eight Interoperable Design of Extreme-Scale Application Software–ECP (IDEAS-ECP) webinars, a Summitdev workshop, and many other training and seminar activities.

See Appendix B for a complete summary of these events. A few of the notable 2017 events are highlighted in Sections 1.4.5.1–1.4.5.4.

1.4.5.1 The Interoperable Design of Extreme-Scale Application Software

The OLCF, the Argonne Leadership Computing Facility (ALCF), the National Energy Research Scientific Computing Center (NERSC), and the IDEAS-ECP continued to partner to produce a series of webinars in 2017—"Best Practices for HPC Software Developers"—to help users of HPC systems carry out their software development more productively. A total of eight webinars were presented in 2017. The IDEAS-ECP 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 (Argonne), Lawrence Berkeley National Laboratory (LBNL), Los Alamos National Laboratory (LANL), ORNL, Pacific Northwest National Laboratory (PNNL), and Sandia National Laboratories (SNL)—with the Colorado School of Mines.

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 both the recordings and the webinar materials were posted online. Approximately 600 people attended these events. The OLCF helped coordinate and promote each of these events, and a few OLCF staff members contributed material to the webinars.

1.4.5.2 2017 Hackathons

In 2017, the OLCF and their partners continued their efforts to teach new GPU programmers how to leverage accelerated computing in their own applications and to help existing GPU programmers to further optimize their codes. Hackathons were held at the Jülich Supercomputing Centre in Germany, Brookhaven National Laboratory in New York, NASA Langley Research Center in Virginia, the CSCS Swiss National Supercomputing Center in Switzerland, and the OLCF in Tennessee. The events at these five institutions hosted 47 teams (compared to 26 teams last year) totaling more than 300 attendees from universities, national laboratories, supercomputing centers, government institutions, and industry. A special Hackathon was held alongside the annual Computational Science Graduate Fellowship (CSGF) annual program review in Arlington, Virginia. During this Hackathon, participants gained experiencing developing GPU-enabled applications on early access Summit architecture. This Hackathon is discussed in more detail in section 1.4.6.1. Building on the theme of interagency collaboration, the NASA hackathon brought together the expertise of teams from NASA Langley Research Center, NASA Ames, and NASA Glen. In a letter to the organizers, Dana Hammond (NASA Langley) wrote

> *"This five day accelerated Hackathon demonstrates how impressive performance gains can be obtained by leveraging interagency and multi-center expertise and hardware resulting in near-term impact and technological excellence. The LaRC HARA code increased its GPU speedup by 10x; a speedup that will immediately be used by NASA flight programs such as Orion and Mars 2020, which apply HARA for radiative heating simulations for thermal protection system (TPS) sizing. The FUN3D code used the next generation GPU to demonstrate significant performance gains that could influence the Agency's next purchase of GPUs. Your daily orchestration and coordination yielded an effective, yet creative unstructured work environment providing the teams not only the motivation to achieve success, but also to share ideas with other teams."*

During the 2017 events, teams developed applications covering a wide range of scientific domains, including astrophysics, CFD, combustion, molecular dynamics, plasma physics, quantum mechanics, climate modeling, and machine learning. The teams consisted of developers with intimate knowledge of their own applications along with two mentors with extensive GPU programming experience. Together, they developed these applications on world-class HPC systems (e.g., Titan and Piz Daint), local clusters, and development systems (SummitDev) for upcoming architectures, such as Summit.

In addition to improving the participants' GPU programming knowledge, these events help to expand the use of accelerated computing in the larger scientific community by highlighting the benefits (e.g., speedups) of accelerated applications within many scientific domains. As an example, the atmospheric CFD application, PALM, was initially ported to GPUs during the TU Dresden/Forschungszentrum event in 2016. Afterward, the team continued porting the PALM code and published a conference paper describing their efforts,[5] crediting the hackathon as the catalyst for their success. In fact, many previous hackathon teams report continued GPU development of their applications well after the event. Some teams even carry out this development under new OLCF projects; examples from the 2017 events include a NASA DD proposal for GPU development of the FUN3D CFD code and a Defense Threat Reduction Agency DD proposal intended to scale up their GPU-enabled HIGRAD/FIRETEC code for modeling wildland and urban fires. A team from one of our first hackathons even went on to win a 2016 R&D100 award,[6] as well as an ECP award.[7] Papers specifically resulting from the 2017 hackathons have now been submitted as well and are currently awaiting acceptance. Examples include "Meeting the Challenges of Modeling Astrophysical Thermonuclear Explosions: Castro, Maestro, and the AMReX Astrophysical Suite,GALARIO: a GPU Accelerated Library for Analysing Radio Interferometer Observations,"[8] "Pushing Memory Bandwidth Limitations through Efficient Implementations of Block-Krylov Space Solvers on GPUs,"[9] and "GALARIO: a GPU Accelerated Library for Analysing Radio Interferometer Observations."[10]

These events also serve as an outreach opportunity for the OLCF; participants are introduced to the facility and its resources, teams are encouraged to submit DD proposals for new projects, and

---

[5] H. Knoop, T. Gronemeier, C. Knigge, and P. Steinbach, "Porting the MPI Parallelized LES Model PALM to Multi-GPU Systems – An Experience Report," in M. Taufer, B. Mohr, J. Kunkel (eds), High Performance Computing, ISC High Performance 2016, Lecture Notes in Computer Science, vol 9945. Springer, Cham.

[6] https://www.rd100conference.com/awards/winners-finalists/6546/nekcemnek5000-scalable-high-order-simulation-codes/

[7] https://www.exascaleproject.org/project/ceed-center-efficient-exascale-discretizations/

[8] Meeting the Challenges of Modeling Astrophysical Thermonuclear Explosions: Castro, Maestro, and the AMReX Astrophysical Suite, arXiv:1711.06203 [astro-ph.IM] (submitted to Proceedings of AstroNum 2017 / *Journal of Physics Conference Series*).

[9] Pushing Memory Bandwidth Limitations Through Efficient Implementations of Block-Krylov Space Solvers on GPUs, arXiv:1710.09745 [hep-lat] (submitted to Elsevier).

[10] GALARIO: a GPU Accelerated Library for Analysing Radio Interferometer Observations, arXiv:1709.06999 [astro-ph.IM] (Submitted to MNRAS)

relationships vendor partners and leaders of the host institutions are established. An example of the benefits of such relationships, a new team member was recently added to the OLCF's User Assistance and Outreach Group thanks to a previous hackathon attendee reaching out to Fernanda Foertter after meeting her at a 2016 event.

Participants from previous hackathons show their appreciation/recognition of these events through acknowledgement sections in papers and simple thank you letters and emails. Michael Zingale (Stony Brook University; OLCF User Group Executive Board 2014–present) has attended many hackathons along with his students.

> *"Dear Meifeng and Fernanda, thank you again for hosting the hackathon last week at BNL. It was a great experience for our team and we learned and accomplished a lot. I can't imagine how much work it is to put on this type of show. The format was excellent and our entire team appreciates the work. We are much further ahead on GPUs after that intensive week of hacking then we were when we started."*

As previously noted, Helge Knoop (Leibniz Universitat Hannover) credited his work at a previous hackathon as the catalyst for a conference paper he went on to publish.[11] In the acknowledgements, he writes

> *"We would like to thank the Oak Ridge National Laboratory (US), Nvidia Corporation Inc. (US), the Portland Group Inc. (US), the standards OpenACC committee as well as the Center for Information Services and High Performance Computing (ZIH) at Technische Universität Dresden and the Forschungszentrum Jülich for organizing the OpenACC Hackathon in March 2016. We would like to thank personally Fernanda Foertter, Guido Juckeland and Dirk Pleiter for organizing the Hackathon in Dresden. Further, we express our deep gratitude to Dave Norton (Portland Group) and Alexander Grund (HZDR; Rossendorf) for their instrumental contribution as members of the mentoring team during the Hackathon. The author team consists of three PALM developers (Knoop, Gronemeier, and Knigge) and one mentor of the Hackathon (Steinbach)."*

Jose Monsalve (University of Delaware) gives a graduate student's perspective in an email he sent after attending his second hackathon.

> *"This was the second GPU Hackathon I have attended and I cannot be any more satisfied with these events. Not only because of how much you learn, which is invaluable, but also because the hackathon is the right environment to get things done, to really move forward in a project. I feel like one week of this event represents the effort of a month for a whole team.*
>
> *Having an expert like Randy next to us during the whole week, having constant discussions with him and your teammates, constantly finding ways to explain your progress and showstoppers, and being surrounded by great enthusiastic people who are joining you thought the pain of seeing your screen full of compiler errors and debugging them, are just some of the key factors that makes this event a completely unique one."*

It should be noted that mature, optimized applications require more development than can be accomplished at a single week-long hackathon event, but the testimonials received from previous attendees as well as their publication records continue to show the impact of these events.

[11] H. Knoop, T. Gronemeier, C. Knigge, and P. Steinbach, "Porting the MPI Parallelized LES Model PALM to Multi-GPU Systems – An Experience Report," in M. Taufer, B. Mohr, J. Kunkel (eds), High Performance Computing, ISC High Performance 2016, Lecture Notes in Computer Science, vol. 9945. Springer, Cham.

1.4.5.3 2017 OLCF User Meeting

One hundred twenty-three OLCF users and staff members attended the annual OLCF Users Meeting held May 23–25 to share achievements on Titan, discuss the next big Summit supercomputer, and delve into deep learning concepts. The event took place as OLCF staff announced a milestone in the history of the center: its 25th anniversary. On the first day of the meeting, talks revolved around Titan—the OLCF's current leadership-class supercomputer, capable of 27 petaflops—and the best ways to manage workflows and analyze data during and after simulations. Virginia Tech's James McClure and OLCF computational scientist Mark Berrill gave invited talks entitled "GPU-Accelerated Digital Rock Physics Workflows for Titan and Summit" and "Task-Parallel In Situ Analysis for Two-Fluid Flow on Heterogenous Supercomputers," respectively. Arizona State University's Abhishek Singharoy also gave an invited talk, during which he described the molecular dynamics NAMD code for large biomolecular systems. His presentation was entitled "Simulating Biological Energy Transfer with NAMD: The Source, the Sink, and Everything that Happens in Between." Singharoy said the meeting provided an opportunity for scientists to share their research and staff members to relay valuable information about Titan and Summit. The first day also included a poster slam, giving users opportunities to present brief overviews of their topics before a poster session. The night concluded with the poster session chaired by OLCF user support specialist Suzanne Parete-Koon, showcasing 26 posters on topics ranging from combustion to genomics to nanomaterials.

The second day focused on the arrival of Summit. The National Center for Computational Sciences (NCCS) Director of Science Jack Wells kicked off the day with a talk entitled "Summit Transition to Operations." Wells explained the milestones of transitioning to Summit, the goals for developing best practices on the new system, and the progress of the Center for Accelerated Application Readiness (CAAR), whose teams optimize the application codes that will show a boost in performance on the forthcoming supercomputer. Summit will feature NVIDIA's Volta GPUs and IBM's POWER9 line of processors.

The last day was dedicated to machine learning. The OLCF's Arjun Shankar, group leader for the OLCF's Advanced Data and Workflow Group and the director of the Compute and Data Environment for Science (CADES) at ORNL, gave the first talk, "Big Data and Analytics at the OLCF." Shankar guided users through the OLCF's deep learning tools, including services for observational data, workflow systems, data analysis, and data sharing. Deep learning, a branch of machine learning that uses algorithms to mimic the ways that humans extract and understand information, was a special area of interest for many users, especially those using the OLCF's latest NVIDIA DGX-1 deep learning system.

Users who were unable to attend the meeting could participate remotely via BlueJeans webcasts. Presentations from the meeting are accessible through the 2017 user meeting webpage.

During the meeting, the OLCF OUG also elected three of its executive board members.

1.4.5.4 Summitdev Workshop

The OLCF hosted a 3-day SUMMITDEV workshop January 10–13, 2017, at ORNL. The purpose of the workshop was to introduce OLCF staff and OLCF CAAR teams to the new Summit Early Access development platform. Both hardware and software aspects of the new system were covered. The event consisted of lectures in the mornings and hands-on activities in the afternoons. Attendees brought their own code and partnered with members of the IBM/NVIDIA Center of Excellence during the hands-on sessions. A total of 87 people participated in the workshop.

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

Since 1946, ORAU and ORNL have partnered to provide internships in subject areas ranging from climate research to nuclear nonproliferation policy. Over these 71 years, more than 9,000 participants from 885 US universities and more than 300 international universities have been provided the opportunity to gain hands-on experience in their respective fields of interest. Each summer, the OLCF hosts student interns of various interests and background, ranging from high school students exploring their interest in programming and HPC to PhD candidates honing their research skills. The OLCF hosted 22 students during summer 2017, 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. Three such interns were Mark Mudrick, Joanna Reed, and Daniel Barry.

Mark Mudrick received his bachelor's degree in physics from Villanova University and is currently pursuing a PhD in computational physics at the University of Georgia. Working with his mentor, computational scientist Markus Eisenbach, Mark focused on combined molecular and spin dynamics simulations.

Joanna Reed is a recent graduate of L&N STEM Academy in Knoxville, where she focused her high school studies on computer science. For her internship, Joanna worked as part of a team charged with programming one of the most crucial day-to-day aspects of operating a supercomputer—keeping it cool. In the Summit Temperature Project, the group was charged with developing software that will take into account factors such as water flow and weather in order to keep Summit's nodes at an acceptable temperature. Joanna's work is part of the big picture plan to increase operational efficiency and to reduce the cost of cooling a machine as powerful as Summit. When the supercomputer goes live in 2018, it will feature a number of cost saving innovations, including the use of warm-water cooling, a method that is projected to lower the OLCF's cooling costs by half.

Daniel Barry is a senior at the University of Tennessee–Knoxville (UT) with a double major in computer engineering and mathematics. He first came to ORNL as a part of the High Performance Computing Student Cluster Competition, an event designed to develop students' skills and confidence in HPC as well as make connections with professionals in the field. This summer Daniel worked with his mentor, Stephen McNally, to inventory operational data flow from the OLCF systems in an effort to understand how to make the best use of the information the systems contain.

The full list of OLCF summer interns is as follow: Aaron Barlow, Daniel Barry, Cade Brown, Ahana Roy Choudhury, Gregory Croisdale, Joshua Cunningham, Swapnil Desai, Evan Fann, Alfred Farris, Rachel Harken, Kratika Jain, Harsh Khetawat, Ben Klein, Hannah Klion, Justin Lietz, Seth Maxwell, Jacob McDaniel, Mark Mudrick, Christopher Muzyn, Joanna Reed, Jeremy Rogers, and Jake Wynne.

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.

1.4.6.1 CSGF Annual Program Review Brings Staff and Students Together

For the eighth 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 July 24–27, 2017, in Arlington, Virginia, provided CSGF fellows with an outlet for research discussion and one-on-one meetings with DOE laboratory staff. The event also offered HPC workshops facilitated by staff members from the DOE laboratories that provided incoming through fourth-year fellows with hands-on training and access to DOE supercomputers. This year the training was organized into two tracks—Computation on an Intel Xeon Phi (aka Knights Landing or KNL) Platform, facilitated by staff from NERSC and the ALCF, and a Mini-GPU Hackathon, facilitated by staff from the OLCF. As part of the mini-GPU hackathon, preassigned fellows gained access to OLCF's Summit Early Access Development Platform, Summitdev, and had the opportunity to work in small teams with an assigned mentor to build, profile, accelerate, and visualize an HPC application using OpenACC or CUDA. To facilitate the mini-hackathon, OLCF staff worked with attendees prior to the event to set up accounts and provided them with a set of basic exercises to prepare them for the day-long event.

1.4.6.2 Introduce Your Daughter to Code

For the second year in a row, the OLCF partnered with ORNL's Women in Computing (WiC) Group to host the event entitled "Introduce Your Daughter to Code." The event introduced ORNL staff members' daughters in middle and high school to the computational sciences with programming activities that center around running code on the OLCF's flagship supercomputer, Titan. This year, 25 girls ages 10 to 18 participated in the labwide event. OLCF's Suzanne Parete-Koon kicked off the event with an introduction to parallel computing and Titan, explaining how the massive machine helps scientists "blow up" stars, study cancer cells and blood flow, and build better engines. Dasha Herrmannova, an intern in ORNL's Computational Data Analytics Group, and Anne Berres, ORNL postdoctoral research associate, walked the girls through the basics of coding in Python, a widely used open-source programming language. Using Python, the girls learned how to type in commands, assign variables, create lists of numbers, and use loops—tools that cycle through sets of values to "find" specific values. Katie Schuman, a Liane Russell Distinguished Early Career Fellow and WiC cochair, helped the girls use fractalName, a program designed by former OLCF intern Susheela Singh, that takes input values and generates colored fractals—repeating patterns that form shapes. The girls used their names and ages to make the colorful pictures, which were projected onto the visualization wall in the Exploratory Visualization Environment for Research in Science and Technology, or EVEREST, at the end of the day. The girls also used a program called Birthday Pi, which Schuman designed, to find their birthday sequences in the first 100,000 digits of the number pi. This exercise demonstrated how loops mine through data and find relevant information using specified customized parameters. After they coded on Titan, the girls explored the interactive Tiny Titan, which provided another visual representation of how processors work together to generate what appears on the screen. The coding experience gave the girls a glimpse into computing, and it gave their parents the opportunity to share insights into their careers with them.

1.4.7 Outreach

The OLCF Outreach team works to engage new and next-generation users and showcases OLCF research through strategic communication activities such as highlights, fact sheets, posters, snapshots, the OLCF website, and center publications (Appendix C). In 2017, the Outreach team was responsible for the creation of 66 highlights—including science and technology highlights and features about OLCF staff members and users—and more than 198 total outreach products. In addition, throughout the year, the OLCF provides tours to groups of visitors who range from middle-school students through senior-level government officials. The center conducted tours for 261 groups in 2017, which included more than 3,600 individuals. The team hosted at least five media-specific tours, including the *New York Times*, *Washington Post*, *Bloomberg News*, WBIR, and the *Knoxville News-Sentinel*, all of which resulted in news stories that featured elements of the OLCF and its resources.

In all, the Outreach team produced a total of 19 science highlights in 2017. Those highlights touched on a variety of science domains, from biophysics ("Assembling Life's Molecular Motor"), to fusion ("Decades-Long Physics Mystery Elucidated with Titan"), to materials ("Putting the Pedal to the Metal in the Hunt for Alloys"). In addition, the team completed 16 technology stories and 25 people features.

In 2017, the Outreach team focused much of their efforts on two themes: the 25th anniversary of the OLCF, and the 2018 launch of Summit. These themes allowed the team to create campaigns, promoting activities and milestones within each topic. The 25th anniversary of the OLCF was the theme for the 2016–17 OLCF annual report and a "breakthroughs" document entitled *Twenty-five Years of Leadership Science at the Oak Ridge Leadership Computing Facility*. The team repurposed the 12 features in the Breakthroughs document for social media and as stories on the OLCF homepage. These items, along with "flashbacks" on social media, gained a lot of attention from followers of the OLCF. The 25th anniversary theme culminated in a large event on November 7, 2017, where ORNL hosted leaders from government, academia, and industry for a half-day of presentations to recognize the technical and scientific achievements of OLCF staff over the last quarter century and discuss the future of supercomputing and the leadership computing facility.

The team used the theme of the 2018 launch of Summit to promote Summit activities and develop a few new products in anticipation of a big Summit rollout in 2018. They created a profile series called "Faces of Summit," writing three profiles in 2017. They also created a branding scheme for Summit, which was used extensively at the Supercomputing Conference, SC17, by the Summit vendor partners as they discussed their progress in the installation of Summit components.

The team also continued to be more intentional and proactive in amplification of their features and news stories. These efforts included using social media, leveraging partner institutions for amplification, pitching stories to specific outlets, and pushing science highlights to a media list through the ping. The net result was that Outreach science highlights were picked up by media outlets (science journals, trade publications and, in a few cases, the regular press) a total of 230 times in 2017. OLCF Outreach science highlights were published by *Science Daily*, *Phys.Org*, *HPCWire*, and other outlets. Outreach highlights were also published on the DOE SC home page and other DOE outlets more than 17 times this year.

In June of 2017, the team made a conscious effort to enhance their social media presence on their two most followed platforms: Twitter and Facebook. The plan included a commitment to making more weekly social media posts and to diversifying the type of posts emanating from the OLCF's accounts. By employing new types of social media posts on the OLCF feeds, which included animated .gifs, reoccurring social media posts such as #ThrowbackThursday, and posting live tweets during OLCF events, the team garnered more social media interactions and views than ever before achieved through their social media efforts. For example, November 2017 was the facility's most trafficked month on Twitter with the @OLCFGOV page receiving 93,000 tweet impressions and 2,390 profile visits, compared to the approximately 30,000 tweet impressions and 510 profile visits for the same month in 2016. These numbers can be attributed to the new types of posts made on this platform, specifically the live tweeting of the OLCF 25th anniversary event.

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 from the OLCF together with Katie Antypas from NERSC and Timothy Williams from ALCF the worked on an initiative to coordinate application readiness activities and develop a strategy to provide guidance and tools encouraging application development that is portable across different architectures. This effort culminated in the recently published book, "Exascale Scientific Applications: Scalability and Performance Portability" that highlighted the application readiness efforts at the ASCR computing facilities.

Additionally, the OLCF collaborated with the three facilities on a project to assess the impact of portable programming approaches on performance in applications. Led at the OLCF by computational scientist Arnold Tharrington and postdoctoral associate Ada Sedova, they compared the performance of various implementations of the calculation of short-range non-bonded forces (SNF), the computational bottleneck in most molecular dynamics simulations, in OpenACC. Their conclusion was that an SNF algorithm that implemented the calculation of the pairwise distances in the SNF computation using matrix-matrix multiplication, though creating more floating-point operations, performs equally well as an OpenACC approach because of the platform-optimized BLAS implementations available. Details of their work, as well as the general conclusions and recommendations for best practices for portability resulting from this joint ASCR facilities collaboration, are publicly available at http://performanceportability.org.

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. The applications that are part of the CAAR program are summarized in Table 1.7.

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

With an initial partition of Summit delivered and accepted in CY 2017, and with the remainder of Summit expected to be fully delivered in CY 2018 and available to the user programs beginning in CY 2019, the OLCF released its Early Science call for proposals in late CY 2017. More than 60 teams, including all 13 CAAR teams, submitted letters of intent expressing interest in submitting a full proposal to the OLCF Early Science program. These letters were due December 31, 2017. The selected teams will be given access to a partition of Summit in early CY 2018 for benchmarking purposes so they may prepare a full Early Science proposal, which is due in mid-2018. The Early Science projects will have access to Summit from the time it has finished the acceptance testing period until the system is fully transitioned to the INCITE and ALCC user programs.

1.5.3 Computational Scientists for Energy, the Environment, and National Security 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 continuing and increasing 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 2017, the CSEEN Postdoctoral program at the OLCF supported eight 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 assisted in the development of and conducted research with particle-based plasma models, particularly the XGC gyrokinetic particle-in-cell simulation code that is being prepared for Summit under the CAAR. He left the OLCF CSEEN program in August 2017 to join NVIDIA as a Solutions Architect.

Yangkang Chen joined the SciComp group in April 2016 after receiving 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 is 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. Amelia became an Assistant Professor of Chemistry at Lakeland University in August 2017.

Kalyana Gottiparthi joined SciComp in September 2015. He earned a bachelor of technology degree in aerospace engineering from the Indian Institute of Technology–Kharagpur in 2007 and received a PhD in aerospace engineering from the Georgia Institute of Technology in 2015. During his postdoctoral appointment, he performed high-fidelity simulations using the Raptor large eddy simulation (LES) code, which is one of the CAAR applications. The code's scalability on Titan has been improved, and his simulations showcase the physics and the performance developments in the code. Kalyana jointed United Technologies Research Center as a Senior Research Scientist in July 2017.

Austin Harris joined the SciComp group in April 2017 after working as a postdoctoral researcher at LBNL. He received a PhD in physics from the University of Tennessee. He is working on extending the number of simulated elements in a core collapse supernova from 13 to 150 in the FLASH code, which will greatly enhance the detail and accuracy of simulations and help uncover the origins of heavy elements in nature.

Anikesh Pal joined SciComp in June 2017 after working as a postdoctoral researcher at the University of California, Los Angeles. He received his PhD in mechanical engineering from the University of California, San Diego, and his master's degree in mechanical engineering from the Indian Institute of Technology Kanpur, India. Anikesh is working on the computational and physical aspects of climate science. He will be focusing on the climate simulations and the so-called super parameterization approaches to resolving multi-scale atmospheric processes in a cost-effective manner with the ACME CAAR project.

Micah Schuster joined SciComp in August 2015 and worked 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. Micah became an Assistant Professor of Computer Science at Wentworth College in July 2017.

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.

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 2017 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 successfully managed policies and job-scheduling priorities that maximized access to these production systems, even in the face of some daunting resiliency challenges on Titan. In 2017, the OLCF once again delivered all of the compute hours committed to the three major allocation programs: INCITE, ALCC, and DD. The business results demonstrate that the OLCF delivered a record year of reliable and technically sufficient resources to the scientific research community.

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 AND WOLF) 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.

In March 2017, the OLCF procured, installed, and deployed the Wolf GPFS file system, which serves as the center-wide file system for the computational resources in the Open Production enclave. Wolf provides a total storage capacity of 8 PB and up to 120 GB/s performance.

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 two 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, Los Alamos National Laboratory (LANL), Sandia National Laboratories (SNL), Lawrence Livermore National Laboratory (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 2017 (Table 2.1). Supporting systems such as EVEREST, Rhea, and data transfer nodes were also offered. Metrics for these supporting systems are not provided.

Table 2.1. OLCF production computer systems, 2017

System	Access	Type	CPU	GPU	Computational description			Interconnect
					Nodes	Node configuration	Memory configuration	
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

2.8.1 OLCF HPC Resource Production Schedule

The OLCF 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 2.2. 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

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

System	Type	Production date[a]	Performance end date[b]	Notes
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	2016 target	2016 actual	2017 target	2017 actual
Cray XK7 (Titan)	Scheduled availability	95%	99.64%	95%	99.39%
	Overall availability	90%	97.03%	90%	98.09%
	MTTI[a] (h)	NAM[c]	473.52	NAM	660.95
	MTTF[b] (h)	NAM	1,750.73	NAM	1,741.49
	Total usage	NAM	90%	NAM	91%
	Core-hours used[d]	NAM	4,323,608,405	NAM	4,389,163,123
	Core-hours available	NAM	4,778,502,912	NAM	4,817,215,104
	Capability usage				
	INCITE projects[e]	NAM	62%	NAM	68.22%
	All projects	35%	59.01%	35%	59.81%

[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.
[d] Does not include usage recorded during an outage.
[e] Does not include INCITE 2016 13th month usage in January 2017.

Table 2.4. OLCF business results summary for Eos

	Measurement	2016 target	2016 actual	2017 target	2017 actual
Cray XC30 (Eos)	Scheduled availability	NAM[c]	99.89%	NAM	99.61%
	Overall availability	NAM	97.97%	NAM	98.39%
	MTTI[a] (h)	NAM	661.97	NAM	783.58
	MTTF[b] (h)	NAM	2,924.91	NAM	2,908.61

[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	2016 target	2016 actual	2017 target	2017 actual
HPSS	Scheduled availability	95%	99.89%	95%	99.46%
	Overall availability	90%	98.54%	90%	98.87%
	MTTI[a] (h)	NAM[c]	376.33	NAM	541.3
	MTTF[b] (h)	NAM	974.93	NAM	1,244.73

[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	2016 target	2016 actual	2017 target	2017 actual
/atlas1	Scheduled availability	95%	99.92%	95%	99.59%
	Overall availability	90%	98.2%	90%	98.87%
	MTTI[a] (h)	NAM[c]	616.12	NAM	509.48
	MTTF[b] (h)	NAM	1,462.79	NAM	872.39
/atlas2	Scheduled availability	95%	99.92%	95%	99.59%
	Overall availability	90%	98.2%	90%	98.88%
	MTTI (h)	NAM	718.81	NAM	433.08
	MTTF (h)	NAM	2,194.18	NAM	671.11

[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 4 year production anniversaries in 2017. The reported results for each system measure are for CY 2017 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 and beyond targets for the accompanying metrics.

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. While this did not occur in CY 2017, the data reflected here artificially assume calculation boundaries of 00:00 on January 1, 2017, and January 1, 2018.

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 Eq. (1). The OLCF has exceeded the SA targets for the facility's computational resources for 2016 and 2017 (Table 2.7).

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

Table 2.7. OLCF business results summary: Scheduled availability

	System	2016 target	2016 actual	2017 target	2017 actual
Scheduled availability	Cray XK7	95%	99.64%	95%	99.39%
	Cray XC30	NAM[a]	99.89%	NAM	99.61%
	HPSS	95%	99.89%	95%	99.46%
	/atlas1	95%	99.92%	95%	99.59%
	/atlas2	95%	99.92%	95%	99.59%

[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, the OLCF HPC Operations group, and the NCCS Operations Council. Typical preventative maintenance activities include software updates, application of field notices, and hardware maintenance to replace failed components. Without concurrence, the systems remain in their respective normal operating conditions.

In 2017, OLCF staff executed scheduled maintenance on the Cray XK7 a total of eight times, which included parallel file system software testing, scheduled facility work, firmware patching, security patching, hardware maintenance, and system software updates. Four unscheduled outages were experienced in 2017, which included a generator failure during a previously scheduled outage, a utility power sag caused by a strong thunderstorm, and two separate high-speed network issues. Similarly, OLCF performed scheduled maintenance on Eos eight times in 2017, with two unscheduled outages.

2.9.2 Overall Availability

The overall availability of OLCF resources is derived using Eq. (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 2016 and 2017.

Table 2.8. OLCF business results summary: Overall availability

	System	2016 target	2016 actual	2017 target	2017 actual
Overall Availability	Cray XK7	90%	97.03%	90%	98.09%
	Cray XC30	NAM[a]	97.97%	NAM	98.39%
	HPSS	90%	98.54%	90%	98.87%
	/atlas1	90%	98.2%	90%	98.87%
	/atlas2	90%	98.2%	90%	98.88%

[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 Eq. (3), and a summary is shown in Table 2.9.

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

The MTTI summary is shown in Table 2.9.

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

	System	2016 actual	2017 actual
MTTI (h)	Cray XK7	473.52	660.95
	Cray XC30	661.97	783.58
	HPSS	376.33	541.3
	/atlas1	616.12	509.48
	/atlas2	718.81	433.08

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 Eq. (4), and a summary is provided in Table 2.10.

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

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

	System	2016 actual	2017 actual
MTTF (h)	Cray XK7	1,750.73	1,741.49
	Cray XC30	2,924.91	2,908.61
	HPSS	974.93	1,244.73
	/atlas1	1,462.79	872.39
	/atlas2	2,194.18	671.11

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

2.10 RESOURCE UTILIZATION 2017

Operational Assessment Guidance

The facility reports Total System Utilization for each HPC computational system as agreed upon with the program manager.

The numbers that are reported for the Cray XK7 resource are Titan core-hours, which are composed 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, 2017, 4,389,163,123 Titan core-hours were used outside of outage periods from an available 4,817,215,104 Titan core-hours. The total system utilization for the Cray XK7 was 91%.

2.10.2 Total System Utilization

2.10.2.1 2017 Operational Assessment Guidance

System utilization (SU) is the percentage of time that the system's computational nodes run user jobs. No adjustment is made to exclude any user group, including staff and vendors.

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

The measurement period is for 2017, 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 2017 was 91%, which marks the fifth 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 2017. Figure 2.1 represents the three major OLCF user programs and does not include consumed core-hours from staff or vendor projects.

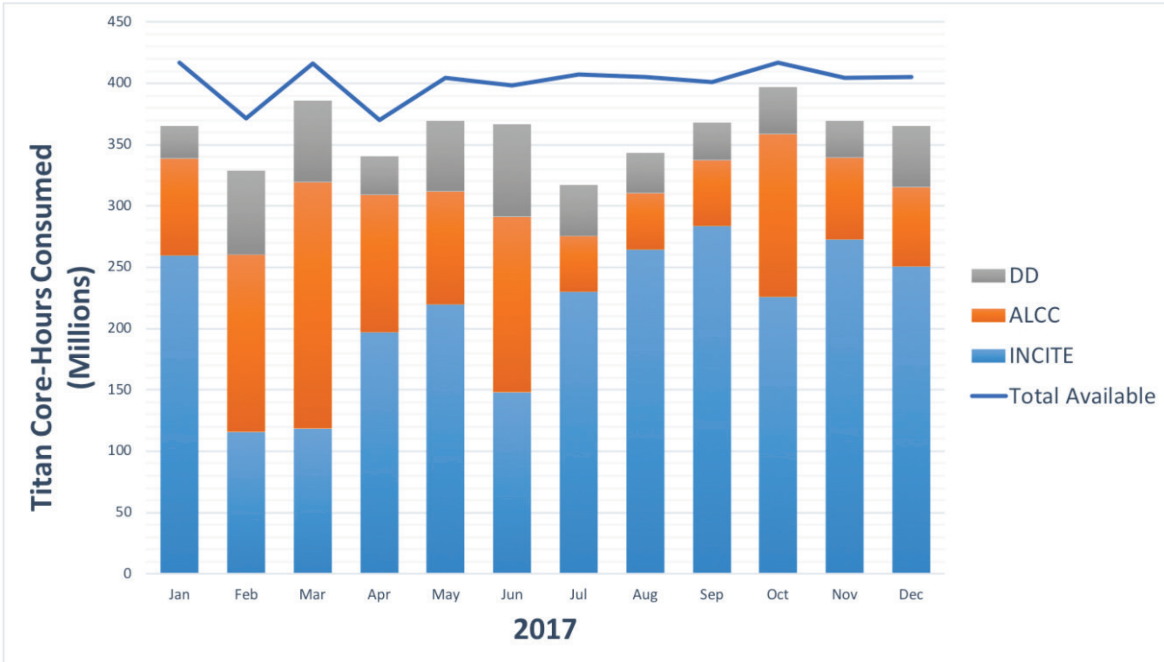

Figure 2.1. 2017 XK7 resource utilization—Titan core-hours by program.

2.10.2.2 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 2017 INCITE allocation program was the largest program in 2017, 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 is in place to permit an additional, final month for completion and 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 2017 allocation period (ending June 30, 2017) were also granted extensions where appropriate.

Table 2.11. The 2017 allocated program performance on Titan

Program	Allocation	Hours consumed	Percent of total
INCITE[a]	2,250,000,000	2,587,030,630	59.86%
ALCC[b]	Allocation spans multiple CY	1,181,882,126	27.35%
DD	—	552,809,167	12.79%
Total		4,321,721,922	100%

[a] Includes all INCITE program usage for CY 2017
[b] Includes all ALCC program usage for CY 2017

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 2017 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 hour wall clock times in the queue.

The OLCF continues to exceed expectations for capability usage of its HPC resources (Table 2.12). Keys to successful demonstration 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.

Table 2.12. OLCF capability usage on the Cray XK7 system

Leadership usage	CY 2016 target	CY 2016 actual	CY 2017 target	CY 2017 actual
INCITE	NAM[a]	62.00%	NAM	68.2%
ALCC	NAM	61.32%	NAM	56.7%
All projects	35%	59.01%	35%	59.8%

[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, 59.8%, was once again well above the 2017 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.

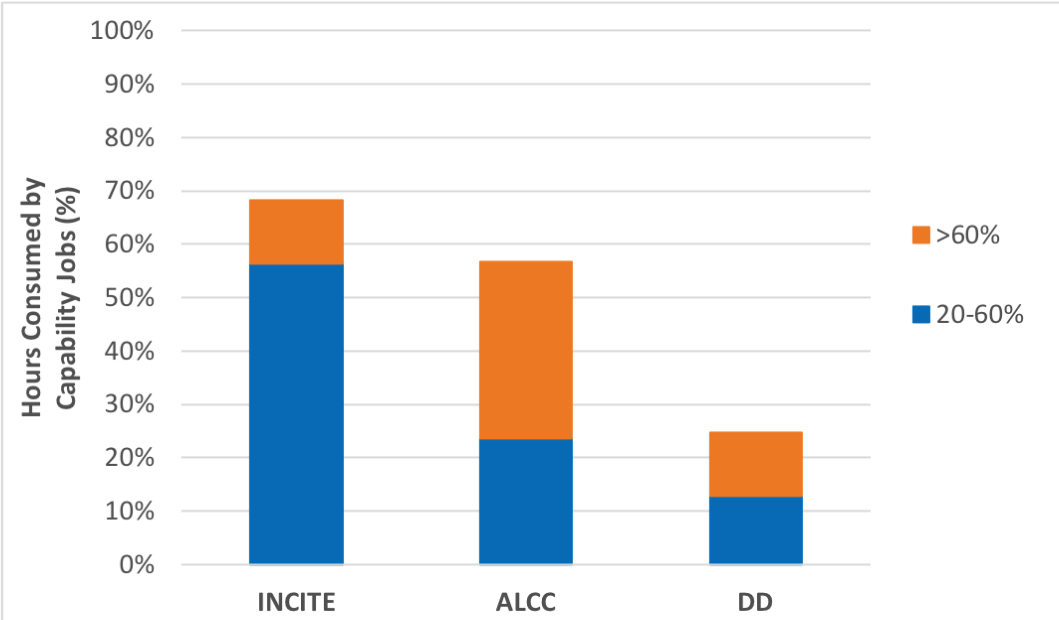

Figure 2.2. Capability usage by job size bins and project type.

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 2017, 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 program reported just under two-thirds usage for GPU-enabled applications, and the ALCC program reported almost three-fourths GPU-enabled usage. The DD program totaled roughly 30% usage for GPU-enabled compute hours. When compared with CY 2016, the INCITE program consumed roughly 90 million less core-hours on Titan but just over 200 million more GPU-enabled core-hours. The ALCC program consumed roughly 225 million more GPU-enabled core-hours in CY 2017 compared to CY 2016. The DD program showed the least preference for GPU-enabled computing. The DD GPU-enabled usage showed a slight decline over the usage that was reported in the 2016 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, and the DD program supports projects that may be in the beginning phases of porting code to GPUs. While the ALCC program does not require computational readiness reviews specifically for GPU usage, the ALCC research teams excelled in the use of GPU-enabled time on Titan, reporting over 70% of the consumed time within the program for the first time ever. For most months in CY 2017, GPU-enabled INCITE applications were consistently responsible for more than half of the delivered hours to those projects and averaged 60.46% for all user programs in 2017, which is an overall increase of 3.5% and the highest overall GPU-enabled consumption for Titan to date.

Table 2.13. 2017 GPU-enabled and CPU-only usage by program

Program	Percentage	Hours
INCITE (GPU-enabled)	61.79	1,598,412,929
INCITE (CPU-only)	38.21	988,617,701
ALCC (GPU-enabled)	72.06	851,632,081
ALCC (CPU-only)	27.94	330,232,044
DD (GPU-enabled)	29.49	163,017,980
DD (CPU-only)	70.51	389,791,188

Approximately 60% of all delivered compute time on Titan was by GPU-enabled applications in CY 2017. Figure 2.3 shows the percentage of GPU-enabled compute time by month.

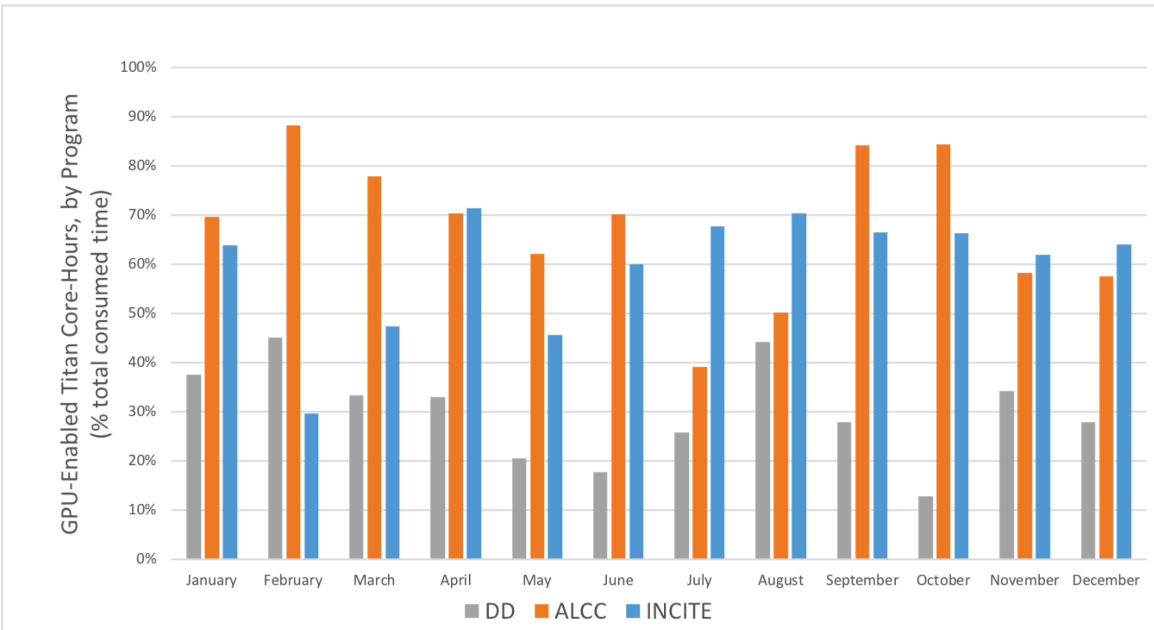

Figure 2.3. GPU-enabled percentage of compute time for the DD, ALCC, and INCITE user programs.

2.13 GPU SELECT SCHEDULING

Beginning in July of 2015, the GPU SXM modules on Titan exhibited a higher-than-expected failure rate. A multiyear investigation was conducted with Cray and Nvidia. A defect in manufacturing was identified, and as a result, newer GPU SXM modules, without this defect, were manufactured and installed into the system. Unfortunately, there were not enough K20X products to entirely replace all of Titan's 18,688 GPU modules. By definition, leadership-class jobs use 20% or more of the nodes in the system. At this scale, applications are more frequently susceptible and are often impacted by these failures. The impacts of failures are typically magnified due to longer waits in queues and greater runtimes.

In 2017, the OLCF studied a method to better utilize the Titan scheduling queue to preferentially schedule leadership-class jobs on nodes that contain the newer GPU modules. Using knowledge from previous work in dual-ended scheduling (that was deployed on Titan in 2015), it was known that the MOAB scheduler uses a first fit scheduling algorithm in Titan's configuration. The application of this scheduling algorithm led to increased utilization of nodes occurring first within the scheduling list, which is created by ALPS, in a network-aware enumeration of Titan's 3D Gemini interconnect.

Two complementary strategies were developed and tested in CY 2017 that showed great promise at improving outcomes for leadership jobs. The first strategy was a reordering of the ALPS enumeration, to maintain network ordering but moving the new GPU modules to the top of the scheduling list. The second strategy employed a technique built upon dual-ended scheduling, which moves smaller jobs to the opposite end of the queue but also now includes jobs from projects which were known to be CPU-only projects. Both techniques were studied in a simulated environment against historic workloads from Titan. These results showed that by using these strategies, leadership-class, GPU-enabled jobs would get 33% more of the new GPU modules in their job allocations, decreasing their likelihood of failure.

The first of these strategies was put into production on Titan in July of 2017. The second (CPU-only preference) is currently under development. Ongoing analysis is showing month to month that failures in leadership jobs have decreased. Figure 2.4 shows that failures would occur in leadership jobs 55–65% of the time prior to July 2017. This frequency has dropped to 47% for the last 6 months of 2017 (July–December) after this GPU Select Scheduling method was put into practice.

As of March of 2018, 11,000 GPU SXM modules were replaced on Titan and 500 new components were withheld as spares. The scheduling changes, coupled with the installation of new components, greatly reduced overall job failures and node failures on Titan, allowing the OLCF to realize the most productive year of record to date.

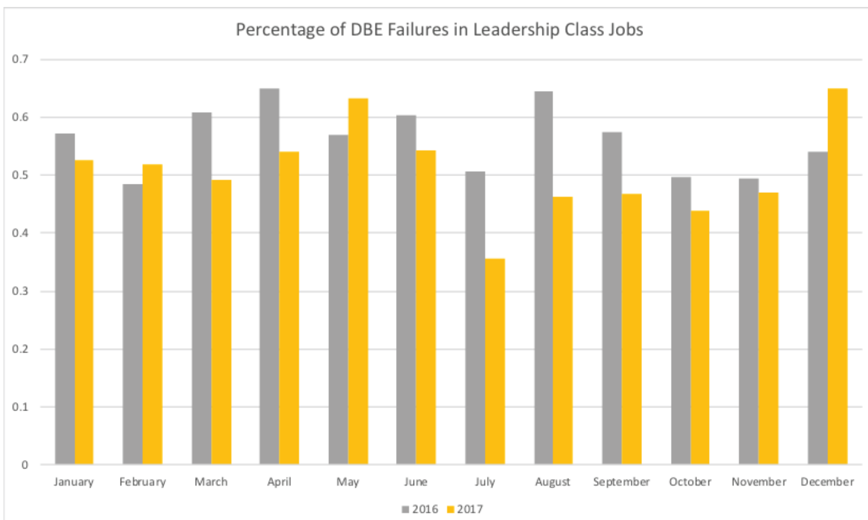

Figure 2.4. Percentage of GPU failures in leadership-class jobs.

2.14 CONSTELLATION DIGITAL OBJECT IDENTIFIER (DOI) SERVICE

The OLCF first announced the Constellation DOI service in the 2016 OLCF OAR. This service had entered the pilot phase in CY 2016. In 2017, the Constellation DOI service was made available to all OLCF users. The DOI service was presented to the OLCF user group executive board, the OLCF user meeting, and also during a monthly OLCF user call. During these presentations, OLCF staff discussed the capabilities and sought feedback on how best to serve the data needs of the OLCF scientific user community. These interactions have helped develop greater understanding of the data needs of the scientific community and the importance of a DOI service that is integrated with scientific workflows. Four DOIs were created in the year 2017; two of the data sets "3-D VPIC simulation of an vortex-induced reconnection event observed by MMS" (DOI: 10.13139/OLCF/1395321) and "Ultrafast current imaging by Bayesian Inversion" (DOI: 10.13139/OLCF/1410993) were associated with articles published in the *Nature Communications*.

One of the major developments with the Constellation DOI portal in CY 2017 was the integration with Globus. With this development, OLCF users can now use Globus for all data transfers. Published DOI data sets that are open to the wider scientific community (non OLCF users) can use Globus to download the data sets. Also, the initially deployed Constellation custom upload tool was replaced with Globus. To date, the development efforts have aimed at making the software more robust and resilient to support diverse scientific needs of the OLCF users.

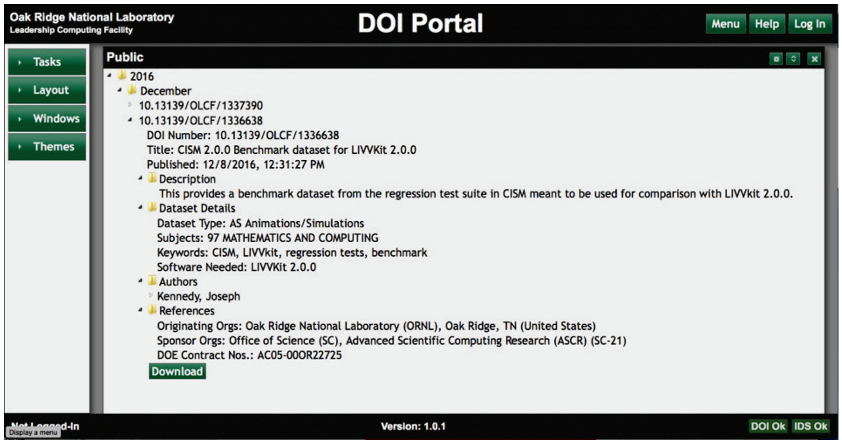

Figure 2.5. Constellation DOI Portal screenshot.

2.15 GRAND UNIFIED INFORMATION DIRECTORY ENVIRONMENT (GUIDE)

The GUIDE tool that was introduced in the 2015 OLCF OAR (referenced as Grand Unified Information Directory for OLCF Operations) presents a window into OLCF operations based on the collection, federation, and analysis of system logs. This service has been in operation for over 2 years at the OLCF. For the most part, the system is in maintenance mode with only a few minor operational tweaks. Members of the Technology Integration group published a paper in the Proceedings of Supercomputing 2017 based on 2 years of operations of the GUIDE infrastructure in the "State of the Practice" track. The paper citation is as follows: Sudharshan S. Vazhkudai, Ross Miller, Devesh Tiwari, Christopher Zimmer, Feiyi Wang, Sarp Oral, Raghul Gunasekaran, Deryl Steinert, "GUIDE: A Scalable Information Directory Service to Collect, Federate, and Analyze Logs for Operational Insights into a Leadership HPC Facility," *Proceedings of Supercomputing 2017 (SC17): 30th International Conference on High Performance Computing, Networking, Storage and Analysis*, Denver, Colorado, November 2017.

3. STRATEGIC RESULTS

CHARGE QUESTION 3: *Are the processes for engaging with programmatic and strategic stakeholders (i.e., beyond the user population) effective and do these processes enable scientific outputs and accomplishments consistent with the DOE strategic goals? Is the allocation of resources reasonable and effective?*

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 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

2017 Operational Assessment Guidance

The Facility tracks and reports annually the number of refereed publications resulting (at least in part) from use of 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 2017, 452 publications resulting from the use of OLCF resources were published, based on the data collection completed on April 6, 2018, representing an 10.4% increase in publications over the previous year.[12,13] A list of 2014–2017 publications is available on the OLCF website.[14] In 2017, users and OLCF staff jointly authored 54 publications, almost matching the number of jointly authored

[12] In this document, "year" refers to the calendar year unless it carries the prefix "FY", indicating the fiscal year.
[13] 405 publications were reported in the 2016 OLCF OAR.
[14] https://www.olcf.ornl.gov/leadership-science/publications/

publications reported for 2016 (53). OLCF users published 391 without OLCF staff coauthorship, and OLCF staff published 7 without user coauthorship.

Sponsor guidance allows accepted and in press publications to be reported, but the OLCF only reports publications appearing in print in the year under review. However, the OLCF continues to search for publications after the OAR is submitted to DOE each year, and the number of publications shown in previous OARs is updated in the current report. Table 3.1 provides the updated, verified, and validated publications count for the 2012–2017 period, showing continued growth in both the total publications count and the number of publications in journals with high impact factors.

Table 3.1 Summary of unique OLCF publications for 2012–2017

Year	Unique, confirmed OLCF publications	High-impact publications with JIF* >10
2017	452	25
2016	453	27
2015	349	18
2014	294	13
2013	359	8
2012	333	17

*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 the OLCF DD program outreach. The following subsections provide brief summaries of select scientific and engineering accomplishments, as well as resources for obtaining additional information. While they cannot capture the full scope and scale of achievements enabled by the OLCF in 2017, 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 these achievements, OLCF users published many breakthrough publications in high-impact journals in 2017, as shown in Table 3.2.

Altogether in 2017, OLCF users published 58 papers in journals with a journal impact factor (JIF) of greater than 7 and 25 papers with a JIF greater than 10. Also noteworthy are the more than 190 OLCF publications that appeared in a wide range of journal types and conferences in 2017.

Table 3.2. Publications in high-impact journals in 2017

Journal	Number of publications
Nature	3
Science	4
Reviews of Modern Physics	1
Nature Chemistry	1
Nature Physics	1
Nature Climate Change	2
Nano Letters	1
ACS Nano	2
Journal of the American Chemical Society	1
Nature Communications	6
Proceedings of the National Academy of Sciences	1
Journal of Physical Chemistry Letters	4
Physical Review X	1
Physical Review Letters	16
Nanoscale	4

3.2.1 Titan Helps Researchers Suck the Mystery Out of Cell's "Vacuum Cleaners"

PI: Emad Tajkhorshid, University of Illinois at Urbana–Champaign (UIUC)
Allocation Program: INCITE

The Science: In the fight against cancer, cancer cells often find ways to fight back. One means is by stocking the cell membrane with a membrane transport protein called P-glycoprotein, or Pgp, that pumps foreign substances—including anticancer drugs—out of the cell. Improved understanding of Pgp's molecular machinery could help the medical research community combat multidrug resistance in cancer cells.

The Impact: Pulling the plug on these molecular pumps requires targeted drugs based on a thorough understanding of Pgp's makeup and mechanics. An all-atom understanding of this protein will give the drug discovery community more information to develop Pgp inhibitors, drugs that block the transporter function and make it easier for anticancer drugs to do their job.

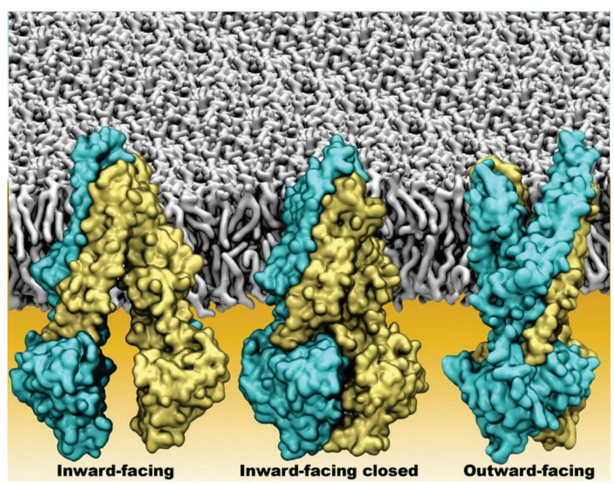

Figure 3.1. UIUC simulation aids understanding of multidrug resistance in cancer cells. This image depicts P-glycoproteins' transition from its inward-facing state to its outward-facing state. For the first time, UIUC researchers successfully captured the protein's outward-facing dynamics. *Credit: University of Illinois at Urbana-Champaign*

Summary: Using Titan, Tajkhorshid's team produced the first experimentally verified all-atom simulation of a mammalian Pgp in its outward-facing state (open to the outside of the cell) within a biological membrane. The UIUC team relied on new experimental data captured via a spectroscopy method called double electron–electron resonance from Vanderbilt University and the molecular dynamics code NAMD to create five distinct simulation sets of Pgp components. Run on Titan, the setup allowed the team to scrutinize Pgp's parts in isolation before bringing them all together.

The substantial sum of simulation data captured on Titan unlocked a trove of new insights, notably how the protein transitions from its inward-facing to its outward-facing state. Additional new findings include a comprehensive look at the ATP-powered mechanism Pgp uses to alter substrate access from inside the cell to outside the cell and a "salt bridge"—an ionic coupling between positively and negatively charged amino acids—that contributes to the protein's structural integrity.

Access to Titan allowed Tajkhorshid's team to compile 300 nanoseconds of simulation time for each protein subunit. In total, the team aggregated 1.5 microseconds. Seventy-two percent of the team's simulation runs required 20 to 60% of Titan, while 28% of the project used less than 20% of the machine.

Contact: Emad Tajkhorshid, University of Illinois at Urbana Champaign, emad@life.illinois.edu

Funding: This research was funded by the National Institutes of Health. Computational resources were provided by the OLCF through the INCITE program supported by DOE's Office of Science (SC).

Publication: Brandy Verhalen, Reza Dastvan, Sundarapandian Thangapandian, Yelena Peskova, Hanane A. Koteiche, Robert K. Nakamoto, Emad Tajkhorshid, and Hassane S. Mchaourab, "Energy transduction and alternating access of the mammalian ABC transporter P-glycoprotein," *Nature* 543, no. 7647 (2017): 738–741. [DOI:10.1038/nature21414]

Related Links: Oak Ridge Leadership Computing Facility highlight: https://www.olcf.ornl.gov/2017/09/19/titan-helps-researchers-suck-mystery-out-of-cells-vacuum-cleaners/

3.2.2 Researchers Uncover Crucial Fusion Phenomena

PI: C. S. Chang, Princeton Plasma Physics Laboratory
Allocation Program: ALCC

The Science: In a fusion reactor, the divertor—which removes exhaust heat from the vacuum vessel—will withstand the highest surface heat load. The divertor must be engineered to withstand this heat so that it is not damaged in the process. A team performed simulations and determined that in ITER, the experimental fusion reactor being built in France, "blobby"-shaped edge turbulence will spread heat across a larger area of the divertor, reducing the risk of rapid damage. The team also simulated the edge of the plasma and found the mechanism by which plasma transitions from low- to high-confinement mode, a requirement for fusion.

The Impact: To avoid damage or costly replacements to the interior material surface of a tokamak, scientists need to understand how to control and remove exhaust heat from a fusion reactor. They also must understand the high-confinement mode (H-mode) because it will enable them to plan for the necessary plasma heating power required to successfully achieve ITER's goal of tenfold energy production. Plasma edge simulations can help guide planning for ITER experiments, ultimately reducing cost through more accurate predictions.

Summary: Using Titan, the team used XGC to simulate where the plasma edge meets the tokamak's divertor. The team found that the blobby turbulence at the plasma edge widens the heat load and that turbulence may spread heat across a larger area of the divertor surface, significantly increasing the heat-flux width to more than 5 cm, which is wider than current smaller-scale fusion devices and less likely to cause damage to the divertor. They also performed secondary simulations to reveal the physics behind the low- to high-confinement transition in the plasma of MIT's Alcator C-Mod tokamak. The team found that two separate theories together explain the transition: turbulence generates a turbulent stress-driven flow, and in the background, some ions are lost to orbits crossing the closed magnetic field lines, generating another kind of flow: an orbit-loss-driven flow. This finding is crucial because if the orbit-loss-driven flow is weaker in ITER than it is in current tokamaks, scientists may need to plan for a higher plasma heating power to successfully achieve ITER's goal of tenfold energy production. Therefore, future simulations will study the transition into H-mode at ITER scale.

The XGC code places tremendous demands on Titan's compute and data-handling capabilities. Ninety-four percent of the team's simulation runs required 60–100% of Titan, and 6% of the project used less than 20% of the machine.

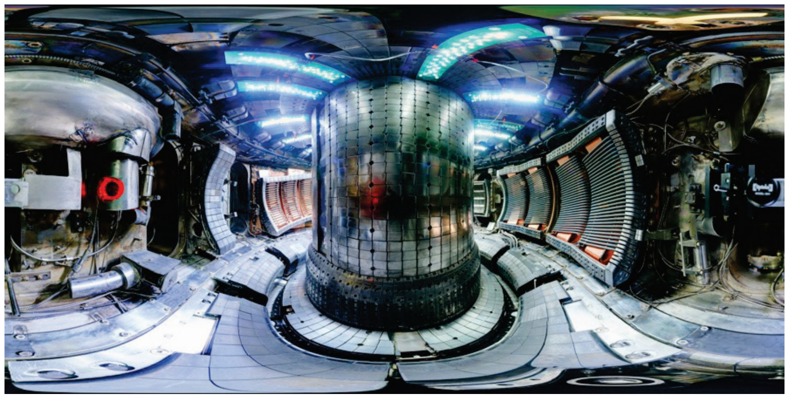

Figure 3.2. The interior of the Massachusetts Institute of Technology's alcator c-mod tokamak. A team led by Princeton Plasma Physics Laboratory's C. S. Chang recently used the Titan supercomputer to simulate c-mod in the first computational demonstration of a spontaneous transition of a tokamak plasma into high-confinement mode, a mode that will be necessary for fusion in ITER. *Image credit: Robert Mumgaard, MIT*

Funding: This work is supported mostly by the DOE SC Fusion Energy Sciences and ASCR through the SciDAC (Scientific Discovery through Advanced Computing) program. Computing resources were provided by the OLCF under the ALCC (ASCR Leadership Computing Challenge) program.

Publications:

C. S. Chang, C. S., S. Ku, G. R. Tynan, R. Hager, R. M. Churchill, I. Cziegler, M. Greenwald, A. E. Hubbard, and J. W. Hughes, "Fast Low-to-High Confinement Mode Bifurcation Dynamics in a Tokamak Edge Plasma Gyrokinetic Simulation," *Physical Review Letters* 118, no. 17 (2017). doi:10.1103/physrevlett.118.175001.

C. S. Chang, S. Ku, A. Loarte, V. Parail, F. Köchl, M. Romanelli, R. Maingi, J.-W. Ahn, T. Gray, J. Hughes, B. LaBombard, T. Leonard, M. Makowski, and J. Terry, "Gyrokinetic Projection of the Divertor Heat-Flux Width from Present Tokamaks to ITER," *Nuclear Fusion* 57, no. 11 (2017). doi:10.1088/1741-4326/aa7efb.

R. M. Churchill, C. S. Chang, and S. Ku, "Pedestal and Edge Turbulence Characteristics from an XGC1 Gyrokinetic Simulation," submitted to *Plasma Phys. Cont. Fusion* 59 (2017). doi:10.1088/1361-6587/aa7c03.

Related Links: Oak Ridge Leadership Computing Facility highlights:
https://www.olcf.ornl.gov/2017/06/27/how-hot-is-too-hot-in-fusion/
https://www.olcf.ornl.gov/2017/10/17/decades-long-physics-mystery-elucidated-with-titan/

3.2.3 Galactic Winds Push Researchers to Probe Galaxies at Unprecedented Scale

PI**:** Brant Robertson, University of California, Santa Cruz
Allocation Program: INCITE

The Science: Galactic wind affects star formation in galaxies by propelling outflows of gas into interstellar space. Simulating this phenomenon at multiple scales, with the goal of creating a trillion-cell simulation of an entire galaxy, could lead to improved understanding of the forces that regulate the life cycle of many galaxies.

The Impact: Low-mass galaxies constitute the most common type of galaxy observed in the universe, but according to astrophysicists' most advanced models, these galaxies should consist of many more stars than they appear to contain. A leading theory for this discrepancy hinges on the fountain-like outflows of gas—galactic wind—observed exiting some galaxies. These outflows are driven by the life and death of stars, specifically stellar winds and supernova explosions. Detailed modeling of galactic wind will help researchers bridge the gap between theory and observation and clarify scientific understanding of galaxy evolution.

Summary: Brant Robertson and University of Arizona graduate student Evan Schneider (now a Hubble Fellow at Princeton University) scaled up their Cholla hydrodynamics code on the Titan supercomputer to create highly detailed simulations of galactic wind. High-fidelity simulation allowed them to study phases and properties of galactic wind in isolation and helped them rule out a theory that cold clouds close to a galaxy's center could be pushed out by fast-moving, hot wind from supernovas.

Robertson and Schneider are moving ahead with a full-galaxy simulation about 10 to 20 times larger than their initial effort. Expanding the size of the simulation to hundreds of billions of cells (representing more than 30,000 light years of space) will allow the team to test an alternate theory for the emergence of galactic wind in disk galaxies. The team is relying on its detailed gas cloud simulations to bridge scales and inform unresolved physics within the larger simulation.

After testing their code on a small GPU cluster, Robertson and Schneider benchmarked Cholla under two OLCF DD awards before letting the code loose under INCITE. In test runs, the code has maintained scaling across more than 16,000 GPUs. Eighty-seven percent of the team's simulation runs required 20 to 60% of Titan, and 13% of the project used less than 20% of the machine.

Funding: This research was supported in part by the National Science Foundation. Computational resources were provided by the OLCF through its DD program and the INCITE program supported by DOE's SC.

Publication: Evan E. Schneider and Brant E. Robertson, "Hydrodynamical Coupling of Mass and Momentum in Multiphase Galactic Winds." *The Astrophysical Journal* 834, no. 2 (2017): 144.

Related Links: Oak Ridge Leadership Computing Facility highlight: https://www.olcf.ornl.gov/2017/08/09/galactic-winds-push-researchers-to-probe-galaxies-at-unprecedented-scale/

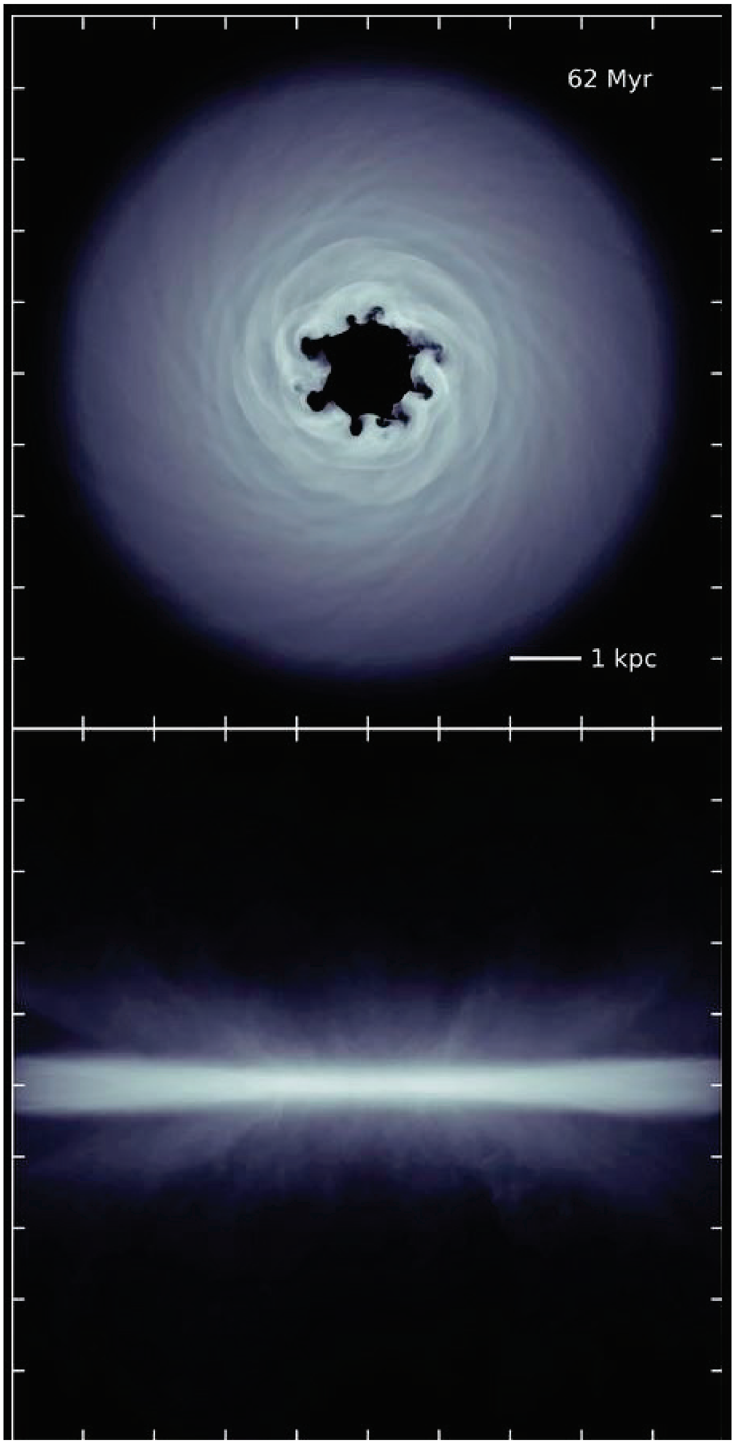

Figure 3.3. A density projection derived from a preliminary full-galaxy simulation of galactic winds. The top panel looks "down-the-barrel" of the outflow, while the bottom panel shows the disk galaxy at the edge. *Image courtesy of Evan Schneider, Princeton University*

3.2.4 Scaling Deep Learning for Sciences

PI: Robert Patton, ORNL
Allocation Program: ALCC

The Science: Deep neural networks—a form of artificial intelligence—have demonstrated mastery of tasks once thought to be uniquely human. Automated tools capable of creating high-performing neural networks without specialized knowledge could expand the benefits of deep learning for the science community by eliminating much of the time-intensive, trial-and-error tuning traditionally required.

The Impact: Because scientific data often look much different from other kinds of data, developing the right artificial neural network can seem like an impossible guessing game for nonexperts. Developing an evolutionary algorithm capable of generating custom neural networks that match or exceed the performance of handcrafted artificial intelligence systems could reduce the time it takes to produce high-performing neural networks from a matter of months to a matter of days.

Summary: Using Titan, Patton's team deployed an evolutionary algorithm called MENNDL (Multinode Evolutionary Neural Networks for Deep Learning) to evaluate, evolve, and optimize neural networks for unique datasets at scale. Central to the process is identifying the best hyperparameters—initial network variables that are typically adjusted on a trial-and-error basis—for varying types of datasets.

To demonstrate MENNDL's versatility, the team applied the algorithm to several datasets, training networks to identify sub-cellular structures for medical research, classify satellite images with clouds, and categorize high-energy physics data. The results matched or exceeded the performance of networks designed by experts. Currently, researchers at the Fermi National Accelerator Laboratory are applying MENDDL to physics data from neutrino experiments to generate optimized algorithms. In addition to improved physics measurements, the results could provide insight into how and why machines learn.

MENNDL leveraged Titan's GPUs to test and train thousands of potential networks for a science problem simultaneously, eliminating poor performers and averaging high performers until an optimal network emerges. Ninety-eight percent of the project used between 20 and 60% of the supercomputer, while 2% of the work required less than 20% of the system.

Contact: Robert Patton, ORNL, pattonrm@ornl.gov

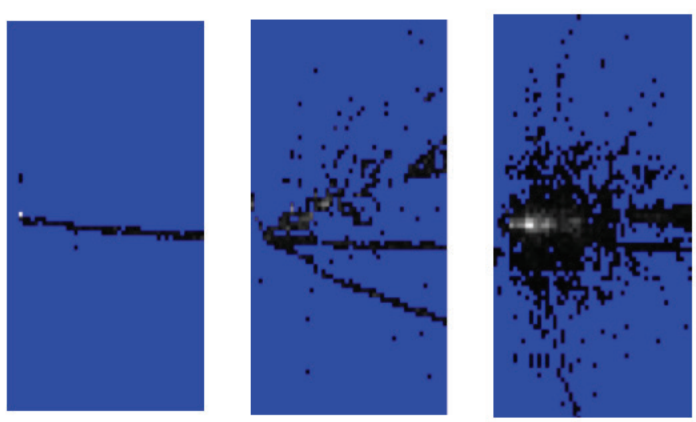

Figure 3.4. An example of the type of scientific data being analyzed by MENNDL to develop high-performing neural networks. This image comes from neutrino scattering data captured by the MINERvA detector at the Fermi National Accelerator Laboratory. *Image courtesy of Fermi National Accelerator Laboratory*

Funding: Computational resources provided by the OLCF through the ASCR Leadership Computing Challenge (ALCC).

Publication: Adam M. Terwilliger, Gabriel N. Perdue, David Isele, Robert M. Patton, and Steven R. Young, "Vertex Reconstruction of Neutrino Interactions Using Deep Learning," in *2017 International Joint Conference on Neural Networks (IJCNN)*, IEEE (2017): 2275–2281.
[**DOI:** 10.1109/IJCNN.2017.7966131]

Related Link: Oak Ridge Leadership Computing Facility highlight:
https://www.olcf.ornl.gov/2017/11/28/scaling-deep-learning-for-science/

3.2.5 Computing the Axial Charge of the Nucleon Using Lattice QCD

PI: André Walker-Loud, LBNL
Allocation Program: INCITE

The Science: On the Titan supercomputer, researchers used lattice quantum chromodynamics (QCD) to calculate the strength with which the weak axial current couples to a nucleon (a proton or neutron)—a fundamental property of protons and neutrons known as the nucleon axial coupling. Ultimately, these calculations will aid in the search for dark matter, help scientists understand the observed abundance of matter over antimatter, and address other probing questions related to the nature of the Universe.

The Impact: As computing power increases, nuclear physicists are pushing the frontier of applying lattice QCD, which has predominantly been used for high-energy physics problems, to low-energy nuclear physics problems. Historically, nuclear physics has been a data-driven field without a direct connection to the underlying theory of the nuclear strong interactions described by QCD. In lattice QCD, the fundamental constituents of the nucleon—quarks and gluons—are formulated on a finite space-time lattice. More accurate low-energy nuclear physics computer simulations, enabled by lattice QCD, could reveal minor discrepancies in the current Standard Model of particle physics that would indicate new discoveries.

Summary: Using Titan's accelerated architecture and an unconventional computational strategy, researchers completed the first lattice QCD calculation of a nucleon axial coupling with a percent-level precision—at least 99% certainty—the highest ever attained for this calculation. Before this project, the research community predicted that this accomplishment could only be achieved on next-generation high-performance computing systems in the 2020 time frame. One of the major challenges of calculating the nucleon axial coupling, and one of the major challenges in applying lattice QCD to nucleon systems, is an exponentially bad signal-to-noise problem. Considerable statistical samples are required for high-precision lattice QCD calculations. Just as large error bars clutter a graph, the tiny mass of the pion particles involved in nucleon interactions generates a lot of unwanted noise, making it difficult to understand the signal.

Using publicly available QCD configurations from the MILC Collaboration, the lattice QCD code Chroma, and the lattice QCD library for NVIDIA GPU-accelerated compute nodes Quda, the team's new computational strategy was able to utilize exponentially more precise numerical data and reduce the amount of statistics needed to reach a precise answer by a factor of 10. This compute and data reduction made the computationally intensive lattice QCD approach feasible. Further, the team was able to enhance the signal of axial coupling by averaging the interaction of the weak axial current over time during the coupling of a neutron and proton, as opposed to selecting one interaction time, as in previous approaches. Their calculation of the nucleon axial coupling at an unprecedented uncertainty of less than 1% is also in line with experimental data.

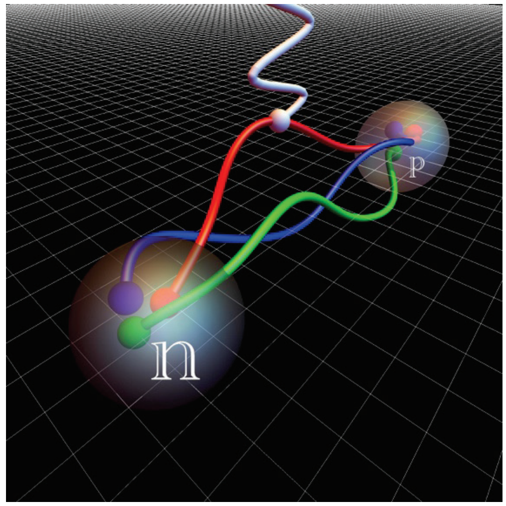

Figure 3.5. A conceptual illustration of a neutron converting to a proton through an interaction with the weak current of the Standard Model. *Image courtesy of André Walker-Loud, LBNL*

Contact: André Walker-Loud, LBNL, Awalker-loud@lbl.gov

Funding: This research was supported by the DOE SC Advanced Scientific Computing and Nuclear Physics offices and the National Science Foundation.

Publication: C. C. Chang, A. N. Nicholson, E. Rinaldi, E. Berkowitz, N. Garron, D.A. Brantley, H. Monge-Camacho, C. Monahan, C. Bouchard, M.A. Clark, B. Joó, T. Kurth, K. Orginos, P. Vranas, and A. Walker-Loud, "A percent-level determination of the nucleon axial coupling from Quantum Chromodynamics," submitted to *Nature*.

3.2.6 Researchers Seek Sigma Meson on the Path to Heavier Hadrons

PI: Robert Edwards, Thomas Jefferson National Accelerator Facility
Allocation Program: ALCC

The Science: Researchers at Jefferson Lab are working to help understand subatomic particles, specifically the principles governing the theory of quantum chromodynamics (QCD), through the GlueX experiment. They use QCD to better understand the fundamental interactions between quarks—the building blocks of protons and neutrons—and the gluons that bind them in a way similar to how atoms are bound together into molecules by the electromagnetic force. GlueX is designed to study "hadrons," composite particles built out of quarks and gluons, one class of which are known as mesons, made of one quark and one antiquark.

The Impact: This work provides resolution of the nature of the sigma meson, a long-outstanding issue in particle physics. Understanding exotic mesons and the fundamental particles that comprise atomic building blocks will help researchers gain a better understanding of the standard model of particle physics and the laws that govern the atomic world.

Summary: Robert Edwards at Jefferson Lab led a team searching for exotic mesons using the Titan supercomputer. When produced, exotic mesons are in existence for only a tiny fraction of a second before they decay. And because of mesons' short-lived, ultra-small nature, Edwards and his team used supercomputer simulations in tandem with experiment to gain a deeper understanding of quark, gluon, and meson behavior. This two-pronged approach helped the team confirm the existence of the sigma meson—a mystery in QCD for over 50 years and an essential step toward observing the heavier, exotic particles that the GlueX is seeking. The team also made the first calculation of the radiative decay of an unstable particle. This calculation is necessary for computing the photo-coupling of exotic mesons and for the GlueX experiment.

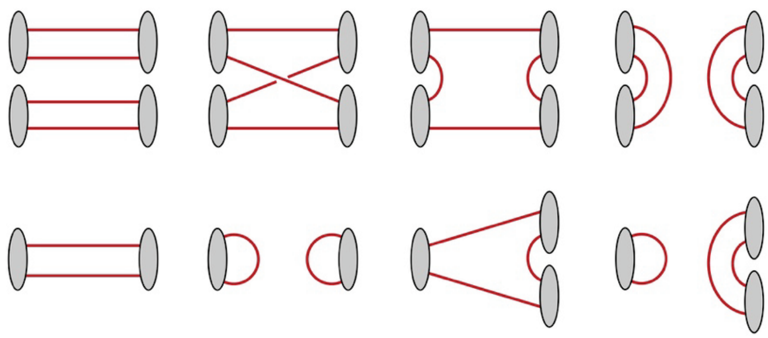

Figure 3.6. Schematic diagrams showing possible quark propagation channels in pion-pion scattering relevant for the sigma meson. *Image courtesy of Thomas Jefferson National Accelerator Facility*

Contact: Robert Edwards, Thomas Jefferson National Accelerator Facility, edwards@jlab.org

Funding: Computational resources provided by the OLCF through ASCR Leadership Computing Challenge (ALCC). The contractions were performed on clusters at Jefferson Lab under the USQCD Collaboration and the Scientific Discovery through Advanced Computing (ScIDAC) program.

Publication: Raul A. Briceño, Jozef J. Dudek, Robert G. Edwards, and David J. Wilson, "Isoscalar ππ Scattering and the σ Meson Resonance from QCD," *Physical Review Letters* 118 (2017): 022002, doi:10.1103/PhysRevLett.118.022002.

3.2.7 Researchers Flip Script for Li-Ion Electrolytes to Simulate Better Batteries

PI: Thomas Miller, Caltech
Allocation Program: INCITE

The Science: Lithium-ion batteries—which are lighter, longer lasting, and more functional than standard batteries under a wider range of temperatures—power everything from cell phones to electric cars. One of the main challenges researchers face in dealing with battery components is finding novel, nonflammable materials for the electrolyte, or the liquid or solid that shuttles ions during charging and discharging, transferring the energy that enables a battery's use.

The Impact: The ubiquitous use of lithium-ion batteries makes their stability, efficiency, and safety important for businesses and consumers alike. Typically lithium-ion batteries feature liquid electrolytes, but new research is focusing on solid polymeric electrolytes, which are known to be more stable, less flammable, and less volatile. Promising electrolyte candidates must be not only stable but also conductive to lithium ions. This characteristic allows batteries to maintain efficiency during charge cycles.

Summary: A team led by the California Institute of Technology's (Caltech's) Thomas Miller used Titan to identify potential electrolyte materials and predict which ones could enhance the performance of lithium-ion batteries. The team first created a coarse-grained simulation protocol called the chemically specific dynamic bond percolation model to screen electrolyte materials based on short molecular dynamics trajectories. Running its simulations on LAMMPS, a classic molecular dynamics code, Miller's team analyzed several dozen polymer–salt combinations under different salt concentrations. About 400 simulations at a time were run in parallel, each consisting of around 3,000 atoms periodically replicated in 3-D space.

The team ultimately found that a class of polymers called Lewis-acidic polymers not only conducted anions more slowly than previous solid electrolytes did but also conducted the positive lithium ions more quickly. Because the positive regions of Lewis-acidic polymers are contained in a small amount of space and their negative regions are spread out over a large amount of space, they give lithium ions more opportunities to dissolve. The simulations showed that these polymers may be capable of producing an eightfold increase in desired lithium conduction and a marked decrease in the unwanted anion conduction. This would be—given the historically slow pace of discovering new polymer materials—a very large jump.

Sixty percent of the team's simulation runs required 20 to 60% of Titan, and 17% of the project used less than 20% of the machine.

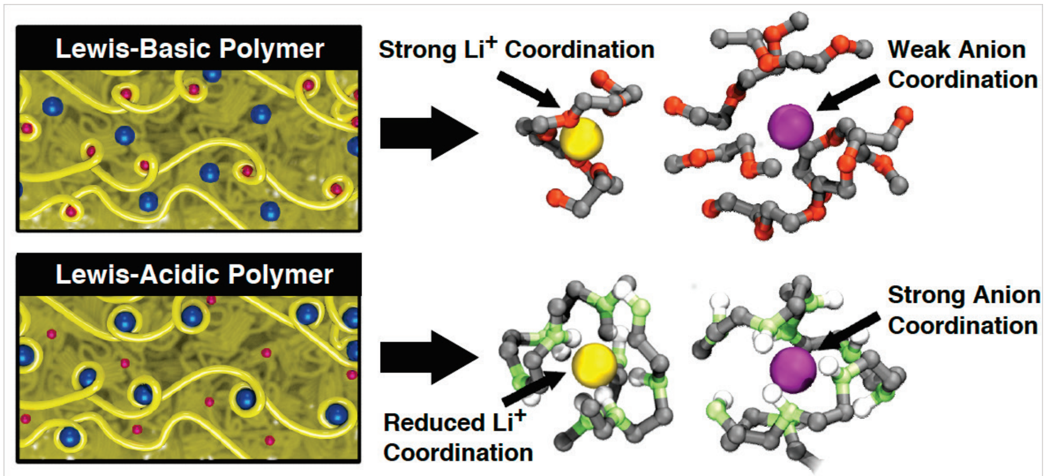

Figure 3.7. Using Titan, researchers are creating models to screen electrolyte materials that could make lighter, longer lasting batteries. Shown here are contrasting solvation strategies in conventional electrolytes and a new class of Lewis-acidic polymer electrolytes. *Image courtesy of Caltech*

Contact: Thomas Miller, Caltech, tfm@caltech.edu

Funding: The research was supported by the National Science Foundation. Researchers also made use of the National Energy Research Scientific Computing Center, a DOE SC User Facility at LBNL.

Publication: Brett M. Savoie, Michael A. Webb, and Thomas F. Miller III, "Enhancing Cation Diffusion and Suppressing Anion Diffusion via Lewis-Acidic Polymer Electrolytes," *Journal of Physical Chemistry Letters* (2016), doi: 10.1021/acs.jpclett.6b02662.

Related Links: Oak Ridge Leadership Computing Facility highlight, https://www.olcf.ornl.gov/2017/01/31/researchers-flip-script-for-li-ion-electrolytes-to-simulate-better-batteries/

3.3 ALLOCATION OF RESOURCES: FACILITY DIRECTOR'S DISCRETIONARY RESERVE TIME

3.3.1 2017 Operational Assessment Guidance

This section should provide insight into the strategic rationale behind use of the Director's Discretionary reserve. The Facility should describe how the Director's Discretionary reserve is allocated and list the awarded projects, showing the PI name, sponsor organization(s), hours awarded, and project title.

3.3.2 The OLCF Director's Discretionary Program

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 to which awardees can effectively use leadership resources. Further, through the ALCC program, the ASCR office allocates up to 20% 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 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. Examples of projects are those associated with the ORNL Laboratory Directed Research and Development program; 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 Exascale Computing Project; the Critical Materials Institute hub led by Ames National Laboratory; the Energy Exascale Earth System Model (E3SM) project; the Center for Nanophase Materials Sciences; and large experimental facilities such as the Spallation Neutron Source and the ATLAS (DOE Office of High Energy Physics) experiment at CERN demonstrating at scale the PanDA workflow management system to achieve the integration of Titan into the Worldwide Large Hadron Collider Computing Grid. In March 2017, the PanDA demonstrator project at OLCF implemented a new PanDA WMS (Workload Management System) instance within OLCF operating under Red Hat OpenShift—a powerful container cluster management and orchestration system in order to serve a broader set of use cases, especially experimental and observational project, at the Titan supercomputer. By running on the premise Red Hat OpenShift built on Kubernetes, the OLCF provides a container orchestration service that allows users to schedule and run their HPC middleware service containers while maintaining a high level of support for many diverse service workloads. The containers run in the context of the project, which allows for direct-container access to all OLCF shared resources such as parallel file systems and batch schedulers. With this new PanDA instance, new demonstrations serving diverse scientific workflows are ongoing including the IceCube experiment, the Neutron Electric Dipole Moment Experiment, LQCD and IceCube, biology studies of the genes and human brain, molecular dynamics, and lattice QCD studies.

Also included in this broad category are projects that come to the OLCF through the Accelerating Competitiveness through Computational ExceLlence (ACCEL) Industrial HPC Partnerships outreach, which encourages opportunities for industrial researchers to access the leadership systems through the usual leadership-computing user programs to carry out research that would not otherwise be possible. A science achievement highlight from Pinnacle Engines is described in Section 3.3.3. 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, 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 decisions 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 2017, the OLCF DD program participants used 12.6% of total user resources consumed for these DD program goals, consuming 553 million Titan core hours.

3.3.3 At the Pinnacle of Cleaner Engines

With an eye toward reducing greenhouse gas emissions and petroleum usage for light-duty vehicles, Pinnacle Engines (www.pinnacle-engines.com) is developing a new, efficient, clean engine technology. The OLCF has provided the Pinnacle Engines team the unique ability to extensively explore a promising class of engines that could provide a diesel-like fuel economy benefit at reduced cost and emissions.

Using the CONVERGE CFD solver, the team performed combustion simulations on Titan and Eos that provided them with a breakthrough combustion recipe, revealing a complex in-cylinder flow development in a combination of swirl and tumble flow structures. The simulations showed that the flow field generates high turbulence, which could enable an engine to operate at lower combustion temperatures. The team validated their results and concluded that a pre-mixed, highly dilute gasoline engine with opposed piston architecture will meet or exceed the team's fuel economy goal and federal emission regulation targets, paving the way for market acceptance of new engine technology. For the Pinnacle Engines problem set, CONVERGE solves the stiff ODEs required for chemical kinetics solvers twice as fast on Titan's GPU-accelerated nodes as compared to Eos' homogeneous architecture. Eos' larger memory per node (64 GB) offers this project some advantages over Titan's smaller memory per node (32 GB + 6 GB). This project is making effective use of both platforms.

Publications:

Siddhartha Banerjee, "High Performance Computing and Analysis-Led Development of High Efficiency Dilute Opposed Piston Gasoline Engine," Internal Combustion Engine Fall Technical Conference, The American Society of Mechanical Engineers, Seattle, Washington, October 15–18, 2017.

Siddhartha Banerjee, "Multi-Dimensional Computational Combustion of Highly Dilute, Pre-Mixed Spark-Ignited Opposed-Piston Gasoline Engine Using Direct Chemistry with a New Primary Reference Fuel Mechanism," Internal Combustion Engine Fall Technical Conference, The American Society of Mechanical Engineers, Seattle, Washington, October 15–18, 2017.

3.4 STAKEHOLDER ENGAGEMENT AND OUTREACH

3.4.1 Programmatic Engagement

In 2018, the aggregate leadership computing capability will reach approximately 220 petaflops based on the delivery of the Summit supercomputer. This falls far short of application requirements that will need a 50–100× increase in application performance in the 2021–2023 timeframe compared to what can be achieved on today's DOE leadership deployments, based on recent and ongoing analyses. This capability gap and the impact of the potential accomplishments have been identified by the US science community and described in a series of recent DOE exascale requirements workshop in 2015, 2016, and 2017 (http://exascaleage.org/). The OLCF partnered with the other DOE ASCR Facility Division programs (ALCF, ESNet, and NERSC) to execute these programmatic requirements activities. Meetings were held for each of the DOE's six SC program offices, plus a cross-cut workshop, as follows:

- High-Energy Physics (HEP)—June 2015
- Basic Energy Sciences (BES)—November 3–5, 2015
- Fusion Energy Sciences (FES)—January 2016
- Biological and Environmental Research (BER)—March 2016
- Nuclear Physics (NP)—June 2016
- Advanced Scientific Computing Research (ASCR)—September 2016
- Cross-cut report, March 2017

In addition, DOE's workshop on management, analysis, and visualization of experimental and observational data (https://science.energy.gov/~/media/ascr/pdf/programdocuments/docs/ascr-eod-workshop-2015-report_160524.pdf) outlines the data science and analytics challenges that next-generation leadership capabilities can overcome. In these reports, world-leading scientists from DOE experimental and observational programs and prominent researchers in mathematics and computer science articulated the scientific needs that closer-coupling of advanced computational and data capabilities can help resolve. Future leadership computing facility systems will enable solutions to these challenging scientific questions.

3.4.2 Community Engagement

A primary and natural place for community engagement has been DOE's Exascale Computing Project (ECP), whose goal is to develop software and applications and influence the development of hardware technology, all to facilitate the successful deployment and operation of capable exascale systems. The ECP has funded proposals from national labs, academia, and industry with the expressed goal of producing usable software and applications and influencing the development of hardware technology for the exascale systems in the 2021–2022 time frame. These investments are very timely and will significantly aid the OLCF in delivering capable exascale systems with robust system software and application software that can address the science gaps immediately upon delivery and acceptance of the systems. In order to plan for the mutual interdependence and capture opportunity, the OLCF and sister DOE SC facilities are partnering with ECP to produce the DOE SC Facilities Engagement Plan with the ECP. This document outlines how ECP, OLCF, and the other DOE computing facilities will work together to ensure mutual success and will accelerate progress toward delivering on their respective missions. The document outlines the rules of engagement, addressing a strategy for communication and conflict resolution. The framework enables the identification of actionable activities that can come from a formal engagement between the DOE facilities and ECP to the benefit of both efforts

Additional notable community engagements include the rapidly developing activities in data analytics, machine learning, and neuromorphic computing and quantum information sciences. OLCF staff worked closely with science collaborators in mapping science problems to artificial intelligence (AI) frameworks and scaling them to run across significant fractions of Titan nodes. These methods required incorporating best-in-class software with cutting-edge libraries and deployment mechanisms.

A DOE laboratory collaborative effort called the Future Laboratory Computing Working Group was co-chaired by Arjun Shankar. This working group reviewed state-of-the-art technologies and DOE facility capabilities to identify methods for establishing a Distributed Computing and Data Ecosystem (DCDE). The working group engaged key program, community, and industry stakeholders, and tools and DOE facilities subject matter experts to identify future requirements and state-of-the-art capabilities available to enable a DCDE.

3.4.3 Industry Engagement

Our industrial partnerships program, Accelerating Competitiveness through Computational ExceLlence (ACCEL), continues to attract new industrial users as well as firms that are returning to solve new and larger problems and grow their internal computational expertise. Project results affirm that ACCEL is helping OLCF deliver high-impact science and engineering results. It also is also enabling DOE to meet its goals to contribute to national competitiveness and grow the community of researchers able to use next-generation leadership computing resources, thereby strengthening the nation's innovation infrastructure.

Thirty-three industrial projects were under way during 2017, which represented 10% of the total number of projects provided to external user programs (INCITE, ALCC, and DD). These projects used 149,836,821 hours, representing approximately 3.5% of the total hours that Titan delivered in 2017.

- In 2017, 58% of the industrial project hours were allocated through INCITE, 31% via ALCC, and 11% through the OLCF DD program.
- Of 33 projects, 13 were new. These firms received awards via INCITE (one project), ALCC (two projects), and DD (10 projects).
- The new ALCC projects were awardees through DOE's Advanced Manufacturing Office HPC4Manufacturing (HPC4Mfg) program. Under this program, projects selected through a competitive call for proposals link companies to national laboratory computational science experts and leadership computing resources in order to apply modeling, simulation, and data analysis to advance innovation in energy-efficient manufacturing and clean energy technologies.
- Three firms that received DD awards were new to ACCEL: Westinghouse, MicrosurgeonBot, and Silicon Therapeutics.
- Microsurgeonbot and Silicon Therapeutics are both small businesses, demonstrating again that small companies are also pursuing innovative, large-scale problems and can benefit from access to OLCF's leadership computing capabilities. MicrosurgeonBot is developing intelligent "middleware" designed to enable someone without specialized knowledge to make use of high performance computing for specialist capabilities such as computer-aided engineering and achieve useful outcomes. Silicon Therapeutics is exploring new approaches to more efficiently and reliably screen ("dock") the millions of compounds that must be investigated in the search for next generation, breakthrough drugs.

3.4.3.1 GM Revs up Diesel Combustion Modeling on Titan Supercomputer

GM researchers used Titan to improve combustion models for diesel passenger car engines with an ultimate goal of accelerating innovative engine designs while meeting strict emissions standards. Diesel engines are about 10% more fuel efficient than gasoline engines, but they produce more emissions, like soot and nitrous oxide, than gasoline engines because of how they combust fuel and air.

Studying the intricacies of combustion—including the thousands of chemical species (types of molecules) created in the process—is difficult in a laboratory and requires significant computational resources to simulate in a virtual environment. At the beginning of this project, GM researchers routinely computed about 50 chemical species in-house. The team wanted to increase the number of species to better understand the chemical reactions taking place during combustion, but in-house computational resources could not compute such complex chemical changes with high accuracy within a reasonable time frame. To improve their combustion model, the team pursued time on Titan through the ASCR Leadership Computing Challenge program. For Titan simulations, GM researchers brought experimental data for about 600 operating conditions—points measuring the balance of engine load (a measure of work output from the engine) and engine speed (revolutions per minute) that mimic realistic driving conditions in which a driver is braking, accelerating, driving up or downhill, idling in traffic, and more. Working with a team at ORNL's National Transportation Research Center and utilizing a GPU-enabled chemistry kinetics solver called Zero-RK, developed at LLNL, the team optimized their combustion model for Titan's accelerators.

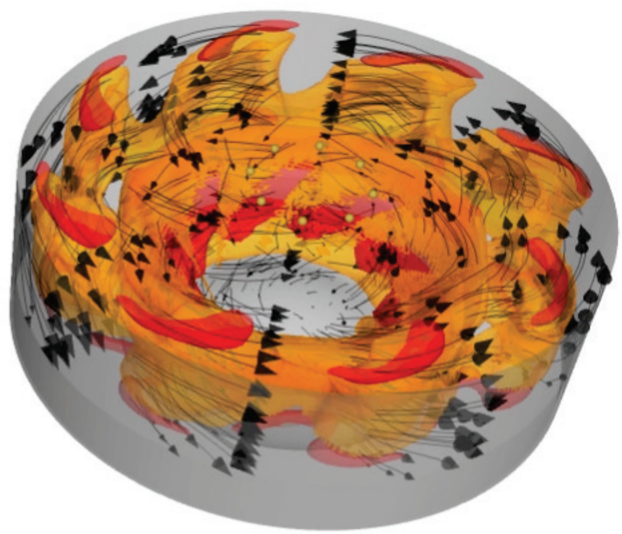

Figure 3.8. A model of a 1.6 liter engine cylinder; liquid fuel (shown in red and orange) is converted to fuel vapor under high temperatures during ignition. *Image courtesy of Ronald Grover.*

Researchers simulated a baseline model of 50 chemical species that matched what GM routinely computed in-house and then increased the number of species for a total of 144. On Titan using Zero-RK, the chemistry computations ran 33% faster. Following these promising results, the team increased the number of chemical species to 766—a number they had previously attempted to simulate in-house at 280 crank angle degrees (a measure of engine cycle progress, with a full engine cycle being 720 crank angle degrees) but could not complete within 15 days.

What had taken the team over 2 weeks to do in-house—modeling 766 species across 150 crank angle degrees—they completed in 5 days on Titan. In addition, they were able to complete the calculations over the desired 280 crank angle degrees, something they had not been able to do using in-house resources.

Compared to the baseline Titan simulation, the refined Titan simulation with 766 chemical species improved nitrous oxide predictions by 10 to 20%. The simulation results are now guiding new research goals to study the effect of heat transfer and combustion chamber wall temperatures on the formation and oxidation of emissions species.

Related Publication: J. Gao, R. Grover, V. Gopalakrishnan, R. Diwakar, W. Elwasif, K. Edwards, C. Finney, and R. Whitesides, "Steady-State Calibration of a Diesel Engine in CFD Using a GPU-based Chemistry Solver," paper presented at the AMSE 2017 Internal Combustion Fall Technical Conference, Seattle, Washington, October 15–18, 2017, doi:10.1115/ICEF2017-3631.

4. INNOVATION

CHARGE QUESTION 4: Have innovations been implemented that have improved the facility's operations? Is the facility advancing research, either intramurally or through external collaborations, that will impact next generation high performance computing platforms, consistent with the DOE strategic goals? Is the facility tracking research results from the broader community, evaluating strategies to incorporate them into next generation high performance computing platforms?

OLCF RESPONSE: Yes. The OLCF actively pursues innovations that can enhance facility operations. Through collaborations with users, other facilities, vendors, and the broader digital infrastructure community, 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 2017, the OLCF pursued strategic organizational and technological innovations and external collaborations to remain the state-of-the-art HPC facility in the United States. 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 2017: key operational innovations, cross-facility integration, strategic external collaborations, and early infrastructure pilot projects for next generation systems.

4.1 OPERATIONAL INNOVATION

4.1.1 Progressive File Layout

A multiyear research and development collaboration effort between Oak Ridge National Laboratory (ORNL) and Intel Corporation completed its final milestone in 2017, significantly improving the usability, performance, and capacity management of the Lustre parallel file system. Lustre is a technology that is widely used on the world's fastest supercomputers and large-scale clusters and is designed mostly to cater to the needs of large-scale scientific simulations and datasets. However, additional work was required for Lustre to better accommodate data analytic and experimental facility workloads, which are beginning to play an increasingly important role in HPC-enabled scientific discovery.

Progressive File Layout (PFL) is a novel storage scheme that absolves Lustre users of the responsibility of file striping, allowing files to be striped dynamically depending on their size, at which point another striping scheme is implemented as the file size grows and so on as the files surpass various size thresholds. PFL allows users to benefit from Lustre's performance and scalability without intimate parallel file system knowledge. It also provides an opportunity for more cost-efficient parallel file system solutions to be implemented in the future by allowing the reliability requirements to be moved into the parallel file system software layer from the hardware.

Large-scale testing of PFL was performed on Titan in June, in partnership with the US Department of Defense and the Extreme Scale Systems Center at ORNL. PFL provided significant improvement in file I/O performance and flexibility and is now available in the Lustre 2.10 community release.

PFL is the result of a 3 year effort in which ORNL codefined the architecture with Intel, oversaw the development efforts, and performed extensive testing at scale.

4.1.2 Fprof Tech Transfer and Results

As first introduced in the 2015 OLCF OAR, fprof is a lightweight and scalable profiling solution that can efficiently walk, analyze, and profile very large-scale parallel file systems. It addresses one of the most pressing challenges in operating increasingly complex and large-scale file systems in production: to gain a deeper understanding of file system characteristics at grand scale without negatively impacting the system performance. This tool has been deployed and is in regular use on very large-scale production parallel file systems at both the OLCF and LLNL Livermore Computing (LC) facilities.

Members of the OLCF's Technology Integration group, in collaboration with LC, conducted an analysis on fprof data obtained from these two institutions. The results and operational insights were published in the 2nd Joint International Workshop on Parallel Data Storage & Data Intensive Scalable Computing Systems (PDSW-DISC). Major findings of this effort were as follows: Stripe pattern analysis demonstrated the disconnect between file system design and actual user/application behavior. This suggested the importance of and urgent need for features such as the PFL and simulated block analysis. Furthermore, the simulated block analysis not only quantified the space overhead for a complete system migration but also revealed optimization points for better performance and space trade-offs by employing more efficient native file system block sizes.

4.1.3 Spider Metadata Snapshot Analysis

In order to understand the behavioral trends of the OLCF user base (over 1000 users) and scientific projects (over 380 projects from 35 science domains), members of the Technology Integration group conducted a detailed analysis of the Spider parallel file system's daily metadata snapshots, collected over 500 days. Specifically, the team analyzed both the individual and collective behavior of users and projects, highlighting needs from individual communities and the overall requirements to operate the file system.

The metadata was analyzed across three dimensions: (i) the projects' file generation and usage trends, using quantitative file system-centric metrics, (ii) scientific user behavior on the file system, and (iii) the data-sharing trends of users and projects. The results indicate the following trends.

1. Over the measurement period, more than 30% of the science domains (11 out of 35) generated over 100 million files. Moreover, many domains had a tendency to create a large number of files in a small number of directories.
2. A project typically contains a 10-fold increase in the number of files when compared to a single user; i.e., a median user account contains 2,000 files, while a median project contains 20,000 files. Most scientific users organize files using a shallow subdirectory hierarchy (less than a depth of 10).
3. Scientific formats such as .c (NetCDF) and .mat (matlab) are within the top 20 popular extensions. However, many scientific applications adopt domain-specific data formats.
4. Storage system performance tuning is actively explored by many projects. For instance, scientists from 20 out of 35 science domains manually configure object storage target (OST) counts for achieving a higher I/O bandwidth.
5. The number of files has grown to over 1 billion.
6. A large portion of the files are not accessed until they are at least 1 week old, but many files are repeatedly accessed beyond the 90 day purge window.

7. Collaboration at the data level is not very common in cross-domain or intra-domain projects. However, projects in climate science typically show active collaboration among users within that domain.

The results of this metadata analysis are published in Seung-Hwan Lim, Hyogi Sim, Raghul Gunasekaran, Sudharshan S. Vazhkudai, "Scientific User Behavior and Data-Sharing Trends in a Petascale File System," Proceedings of Supercomputing 2017 (SC17): 30th International Conference for High Performance Computing, Networking, Storage and Analysis, Denver, Colorado, November 2017.

To the best of our knowledge, this work is the first of its kind, providing comprehensive insights on user behavior from multiple science domains through metadata analysis of a large-scale shared file system. The OLCF envisions that this case study will provide valuable insights for the design, operation, and management of storage systems at scale while also encouraging other HPC facilities to undertake similar such efforts.

4.1.4 CADES Integration into NCCS

During the reporting period, the OLCF engaged in activities that were targeted at user productivity and increasing the use of HPC within nontraditional scientific domains. The first of these activities was the final restructuring of personnel resources that provide services under the ORNL Compute and Data Environment for Science (CADES) umbrella into the NCCS so that their efforts could be delivered in a more streamlined way that mimics the highly efficient OLCF HPC Operations service deployments. The appointment of Arjun Shankar as the group leader of the Advanced Data and Workflows Group and head of the existing CADES staff further integrated the CADES staff into the data, workflow, and development efforts of the NCCS. This section describes some of the innovative and collaborative efforts that CADES and the OLCF accomplished in CY 2017. These new capabilities are a direct result of the recognized leadership and expertise of the NCCS within both the laboratory as well as the HPC community.

The CADES processes and user support structures were still in their infancy when the CADES staff were incorporated into the NCCS. This presented a great opportunity to incorporate the best processes from both facilities (CADES and the OLCF) into a more multi-programmatic NCCS. The existing user support services of NCCS was a natural start. The user-facing Help Desk function of CADES was deployed in the Request Tracker (RT) system alongside the support queues that the OLCF uses to support the users of Titan. This alignment of processes allowed the CADES system administration staff to focus on other system and software deployments. Additionally, the CADES team has deployed their own Resource Utilization Council (RUC), modeling it after the OLCF's RUC that is used to manage allocations within the user programs. Another positive change to CADES operations as part of NCCS is that the CADES team members are now active participants in the weekly NCCS Operations meeting where issues are discussed; outages for resources and services are planned, coordinated, and finalized; and current operational metrics are covered.

NCCS staff engaged in activities that supported an Office of Science Biological and Environmental Research (BER) Atmospheric Radiation Monitoring (ARM) Data Center workflow, which provided HPC capability to data processing and scalable simulations. NCCS has been the archival location for ARM observational data since the early 1990s. The delivered workflow now has data flowing from the ARM data center to CADES through multiple protection zones for processing and storage in the NCCS HPSS. In addition, this workflow enables data to be transferred from the moderate production security enclave instance of the HPSS archive into the CADES scalable simulation platform for the BER ARM program to perform simulation runs such as large eddy simulations and stores the product of that analysis back in the HPSS archive. The significant innovation here is that the NCCS has enabled a cross-facility workflow that allows observational scientists to take advantage of leadership-scale computing and infrastructure.

In support of the cross-facility workflows such as the example above, the OLCF adapted its user access and authorization system (RATS—the Resource Allocation Tracking System) to closely link with ORNL's access control mechanisms (UCAMS/XCAMS), which will provide a flexible pathway for collaborators of ORNL science to apply for access to OLCF resources in the Open Production or Development security enclaves. This mechanism will provide easier access and a more natural pathway through the CADES open environment, further enabling CADES to serve as a gateway user facility for future OLCF researchers and collaborators.

For the community that runs the Accelerated Climate Modeling for Energy project, now the Energy Exascale Earth System Model (ACME/E3SM) in the OLCF, outward facing code and data parameter repositories are being provided via CADES cloud resources that were updated in 2017. In addition, the ACME/E3SM team demonstrated and deployed an analytics platform using the CADES cloud-based Apache Zeppelin in-memory analytics capability on ACME/E3SM data. An older version of this portal had previously been operated on a VM within the NCCS.

The OLCF piloted the use of a deep learning appliance—the NVIDIA DGX1—which was used by several research groups and was demonstrated as an important tool for research. External and internal teams used the resource more flexibly as it was deployed in the CADES ORNL Open Research Protection Zone. These teams are now progressively moving their codes to the Summit architecture. Having a quick-start open environment without the formal procedures of user credentialing and export control was perfect for prototyping open research exploration—the CADES environment was a natural fit for this innovative pilot effort.

As the OLCF continued to research methodologies for using containers in production for users of OLCF systems, the expertise of the CADES team became very beneficial. The CADES team has significant expertise in developing, deploying, and managing containers, and the OLCF team was able to leverage much of that work to create an almost fully automated process for creating a prototype container and deploying it for use on Titan. As the process was being developed and tested, transferring the container image between the CADES resource and Titan was found to be a significant bottleneck. The NCCS felt it imperative that a high bandwidth network connection be established to eliminate this bottleneck if this workflow was to be deployed in production for users. Removing this bottleneck is discussed below.

Both CADES and the OLCF are engaged in ongoing efforts at CERN's Large Hadron Collider (LHC), in Switzerland. The OLCF takes part by supporting an opportunistic workflow called PANDA where jobs are scheduled on Titan nodes that would have otherwise sat idle as the system works to drain nodes for capability jobs. During these opportunistic jobs, data is rapidly transferred from the LHC for analysis. The CADES team participates in the ALICE project by providing compute analysis capability for data generated on the LHC. The CADES team noticed that their incoming network link was being pushed to the limit and that additional traffic from the OLCF's use of CADES resources to support users with data portals and container creation could potentially slow the ingest of data from ALICE and other cross-facility workflows as they progress toward running on leadership resources. The network architecture previously required the CADES traffic destined for the OLCF to cross the outbound border of the Science DMZ of the CADES network to reach the OLCF. The teams developed a plan to deploy a direct link between the open science DMZ networks in each facility (CADES and the OLCF) to support these data movement activities and limit the amount of data that would be required to cross the outbound Science DMZ of the CADES network. This link was deployed and operational in October of 2017.

4.1.5 Routed Spectrum Scale Using IP over InfiniBand and Native Linux Kernel Tools

In 2017, the NCCS, in collaboration with the DOE Atmospheric Radiation Measurement Program's LES ARM Symbiotic Simulation and Observation (LASSO) Initialization, deployed a Cray Rhine/Redwood XC40 supercomputer, Cumulus, and a Spectrum Scale file system, Wolf, using IP over InfiniBand (IPoIB) and native Linux kernel tools.

Spectrum Scale lacks a routing facility similar to Lustre's networking (LNET) capability. In order to facilitate communication between the storage and compute platforms, Cumulus initially projected the Wolf file system through Cray's Data Virtualization Service (DVS). The original DVS mount of the Wolf file system on Cumulus created issues for LASSO's users. While the raw bandwidth of the DVS mount was deemed satisfactory, the lack of full POSIX compatibility left users unable to adapt their workflows to the new environment without making significant changes. Additionally, certain software packages and libraries were unavailable to them while they awaited vendor-provided patches.

To address these needs and support similar use cases in the future, the DVS projection method was replaced by a native Spectrum Scale cluster on Cumulus that routes traffic to and from Wolf at comparable performance. This configuration also provides better support for nontraditional HPC workloads. For example, the IPoIB routing method shows increased performance for workloads that are not I/O optimized. Initial benchmarking has revealed better performance for two common bottlenecks: small-file I/O and metadata operations. While some traditional HPC applications benefit from highly specialized I/O routines performing large streaming read and writes to the filesystem, this is not always the case. The routed IPoIB method supports non-optimized and interactive I/O to the file system better than the DVS equivalent.

Additionally, the Cumulus compute nodes now run the Spectrum Scale client natively, allowing the use of an operating system call included in several software packages. Prior to running the clients natively, Cumulus applications required recompilation without this operating system call because DVS did not support it. In practice, the OLCF has also measured performance gains with the native Spectrum Scale client when compared to the DVS method in metadata-intensive workloads—coupled with the added benefits of byte-range locking and detailed client I/O statistics.

This operational comparison has demonstrated tangible positive results for current applications and next generation and nontraditional HPC applications.

4.1.6 Workforce Development through Undergraduate Data Center Course and Minor Degree

As reported in the 2015 and 2016 OLCF OARs, the OLCF continued its partnership with local ORNL facility and data center experts and the University of Tennessee to offer a first-of-its-kind undergraduate data center design and management course and minor degree. The program completed its third year in CY 2017.

The OLCF extended its partnership to private sector data center and digital infrastructure organizations such as the Infrastructure Masons (www.imasons.org). The iMasons organization was established to provide infrastructure executives and technical professionals an independent forum to connect, grow, and give back. When the OLCF discovered the iMasons, it was clear that there was a very natural opportunity to collaborate. The OLCF became the first federally funded organization to officially join with the iMasons. Shortly after joining, the iMasons asked the OLCF Operations Manager, Stephen McNally, to chair the very active Education Committee, which he accepted in early 2018 (https://imasons.org/imasons-blog/stephen-mcnally-of-oak-ridge-national-lab-appointed-chair-of-imasons-education-committee/ and https://www.olcf.ornl.gov/2018/02/27/mcnally-appointed-education-chair-for-imasons/). The OLCF values partnerships such as this one, as they seek to ensure that there is a highly skilled and available pool of talented individuals for the future of digital infrastructure, which includes high performance computing.

The ability to directly interact with digital infrastructure industry pioneers such as Google, Facebook, Microsoft, Apple, and many others is an opportunity that few universities are able to take advantage of. The OLCF will directly benefit from such partnerships and will also play an important role in shaping the future of the digital infrastructure workforce both locally and nationally.

4.2 RESEARCH ACTIVITIES FOR NEXT GENERATION SYSTEMS

4.2.1 Next Generation Workflows Enabled through OpenShift Infrastructure

OpenShift is a container orchestration platform developed by Red Hat that makes it possible for users to run services in containers and is a certified distribution of the Kubernetes container orchestration software. The strong security sandboxing inherent to Kubernetes allows user applications running in these containers to be accessible from the Internet. Additionally, OLCF resources such as file systems and batch schedulers are also accessible to these applications, depending on user requirements. The OLCF has provisioned two OpenShift clusters to prototype this new capability—Marble and Granite. Marble is a cluster specifically targeted to fulfill users' container requests, and Granite is a cluster used for OLCF-related core infrastructure services.

Users schedule their containers to run on Marble in a manner similar to how users schedule their jobs to run on compute resources such as Summit, Titan, or Rhea. Because the allocable units of an OpenShift cluster are based on CPU cores and memory, the user can request these resources for their container based on the needs of their application. Quotas are enabled, and quota exceptions are handled by the OLCF Resource Utilization Council (RUC). Additional container requirements such as network ports or hostnames for a website can be requested by the end user. These capabilities are defined in part by both workflow and data portal requirements; however, a container is flexible and can be used to run many diverse workloads such as webservers, databases, and custom service daemons.

In 2017, the OLCF demonstrated that these OpenShift clusters can run containers that mount Lustre; make use of the center-wide file system Atlas; and schedule jobs that run on Titan, Rhea, and the Data Transfer Node cluster. These tenants are the basic building blocks of any workflow application. The Marble cluster was deployed in the moderate production security enclave and provides access to experienced OLCF users to test their deployment of web dashboards and workflow systems. The BigPanDA project is hosting a PanDA workflow server, and the OLCF applications team is using OpenShift to run a dashboard that shows test results for application performance testing. To preserve a moderate level of security, when users expose their services to the Internet, they are automatically protected by a two-factor authentication, and by default only project members have access to the service. This is configurable, but it ensures that only project members have access to project data by default.

The OLCF has also been working on making this service available to other staff and groups so that they can take advantage of this novel approach to deploy internal and user-facing websites. The OLCF deployed the Granite cluster for core infrastructure services and has started to migrate some OLCF applications and operational services. This infrastructure provides OLCF staff more visibility into their applications and lets the application owner rapidly deploy new versions. This is especially helpful for ongoing internal software and analytic development frameworks. Two internal pilots were launched on this framework. They are the Summit Temperature Control project and an instance of Kafka that the OLCF security team is maintaining.

4.2.2 Notable Data and Visualization Software Enhancements

To better serve the data and visualization communities, the OLCF funded and provided software enhancements to two key software applications in CY 2017. The VisIt received targets releases focused on tool stability, as well as added new features for users. In CY 2017, the Scientific Data Group (SDG) within the OLCF was an active contributor to the VisIt project and provided installation and maintenance support for VisIt software on the Rhea analysis cluster.

The SDG also provided bug fixes and enhancements to the ADIOS data readers for the XGC and SPECFEM3D_GLOBE codes. The enhancements for XGC include better support for 2D views of the data, alternative coordinate systems, and better support for reading particle files. An additional XGC reader was added to facilitate post-processing requirements of XGC collaborators. Several bugs were fixed in the SPECFEM3D_GLOBE ADIOS reader, and some enhancements were made that made alternative coordinate systems more flexible. To support Radiation Transport codes, an ADIOS reader was added to VisIt to read new data formats for SCALE/Denovo codes.

The pbdR collection of R packages for HPC data analysis released several updates and issued a few new releases in CY 2017. The most notable addition was *pbdADIOS*, a parallel data reader based on ADIOS software, bringing in situ capabilities to the R/pbdR data analysis codes. Another notable addition was the *kazaam* package for parallel data analysis kernels that involve skinny dense matrices, where speed is gained by exploiting the smaller dimension of a data matrix.

4.2.3 Forward-Looking Engagements

OLCF staff actively participate on many technical committees, advisory boards, and collaborative groups. While it is impractical to list each and every community interaction, the following examples are representative of key areas in which the OLCF staff engaged with the broader science and technology community in CY 2017.

4.2.3.1 The Petascale DTN Project

As mentioned in the 2016 OLCF OAR, the OLCF deployed a new data transfer node (DTN) cluster with updated hardware and software specifically designed for high performance data transfers between HPC user facilities. This project was titled the "Petascale DTN Project" and was completed in 2017, along with other partners such as the Energy Sciences Network (ESnet), the National Energy Research Supercomputing Center (NERSC), the Argonne Leadership Computing Facility (ALCF), and the National Center for Supercomputing Applications (NCSA). The goal of the project was to demonstrate sustained data transfers between the facilities of at least the 1 petabyte per week. The project officially concluded in November of 2017 with sustained data transfers between the facilities at better than the 1-petabyte-per-week measure originally instituted. Figure 4.1 shows data transfer rates that were demonstrated between the HPC facilities in November 2017. Based on this data, the OLCF can, on average, transfer 1 petabyte in 2–3 days to and from the other HPC user facilities of this project. This is a significant improvement over the multi-week process that preceded this project. ESnet provided the network infrastructure used to facilitate the transfers across the wide-area network, and researchers provided the data set used to test the connectivity and transfer rates. More information about the project can be found at https://esnetupdates.wordpress.com/2017/12/11/esnets-petascale-dtn-project-speeds-up-data-transfers-between-leading-hpc-centers/.

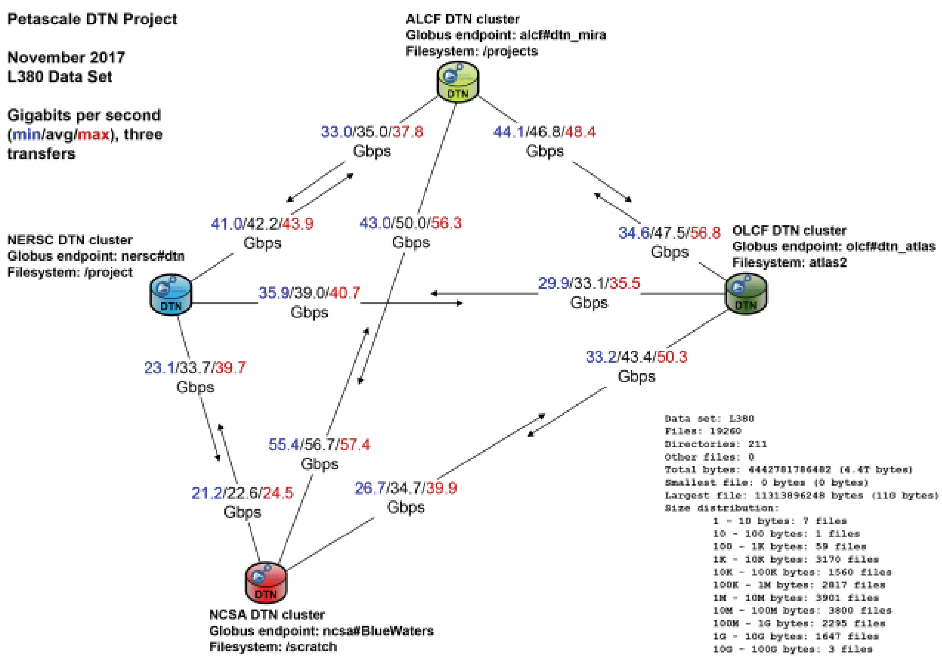

Figure 4.1. Performance measurements from November 2017 at the end of the Petascale DTN project. All sites met or exceeded project goals. *Image credit: Eli Dart, ESnet*

4.2.3.2 ARM Testbed Deployment, Operations, Outreach, and Findings

The OLCF deployed a Cray Envoy system in December 2016. Approximately two dozen users have access to this test and development system and are mostly working on application portability to the ARM architecture. Jack Dongarra's group included this system in a survey of the performance of High Performance Linpack benchmark on different architectures.

In March 2017, the OLCF began working with Hewlett-Packard Enterprise (HPE) as a member of their ARM-based Comanche project. The system the OLCF plans to acquire as part of this project will provide researchers and HPC practitioners with a newer generation ARM processer than that of the Cray Envoy system and will also contain some interesting technologies that the Envoy lacks, such as GPUs in some of the nodes. The first two nodes arrived in CY 2017, and the rest of the system is scheduled to arrive before the end of the CY 2018. This system will allow the OLCF to continue exploring the use of ARM architecture for HPC workloads. The OLCF is specifically interested in

- exploring the use of GPU's with ARM CPUs;
- exploring the use of ARM's "Scalable Vector Extensions" (SVE) instruction set, which will have to be done via emulation, since hardware implementing SVE does not exist yet; and
- cleaning up the Lustre codebase for ARM in an effort to make Lustre on ARM architecture "production class" software.

The team was particularly successful at outreach in CY 2017 as there were presentations at ISC'17 (as part of the ARM BoF session), HP-CAST, and two separate sessions at SC'17 (ARM BoF session and the ARM-HPC User Group meeting). Through their evaluations, the team has provided direct feedback to Cray about their experience with the Envoy system. This included specific issues such as "The compiler fails with an internal error building Kokkos" and more generic ones like "Having to cross-compile on the Envoy system adds a significant burden and makes it harder to get work done." Cray has listened to this feedback: the compiler bug was fixed quickly, and their next-generation ARM systems will leverage a native compiling environment. User experiences around the performance of the architecture have been limited at present due to the sensitive nature of this disruptive technology. This work was enabled through the innovative security enclave developments as discussed in the cyber security section and the cross-facility accounting work carried out through the CADES and OLCF integration activities.

4.2.3.3 OLCF Technical Community Participation

The OLCF staff serve on a variety of technical review and programmatic advisory committees through organizations, conferences, and workshops. Through these engagement activities, the OLCF is able to stay abreast of the notable technological advancements and potentially disruptive or innovative technologies or approaches. OLCF staff are encouraged to participate in these capacities to educate the community on the novel and innovative activities that the OLCF is conducting and to learn from other state-of-the-art user facilities, industry partners, and researchers. Through this culture of service through these organizations, the OLCF maintains an active and forward-looking pulse on the greater technology community. A representative set of engagements is included below.

ECP PathForward Engagement**:** OLCF's Scott Atchley serves as the Technical Representative for AMD's PF contract (and AMD's previous FastForward2 contract). In that role, Scott acts as DOE's point-of-contact with the labs. Scott's responsibilities include distributing deliverables and soliciting feedback from all involved as well as contract management and monitoring of progress and risks. Scott and his LLNL counter-part, Pavlos Vranas, then determine if a deliverable has met its requirements and authorize payment. Within OLCF, members from the Technology Integration, Scientific Computing, and Computer Science Research groups review and provide feedback on ECP PF deliverables from all six of the PF vendors as well as attend deep-dives by the different vendors on their projects.

Open Fabrics Association: Scott Atchley also serves on the board of the Open Fabrics Association (OFA) as OLCF's representative. The OFA serves as the focal point for HPC interconnect vendors, system operators, and users to promote common programming interfaces and to exchange best practices. Every high performance interconnect vendor participates in OFA. Scott also served on the OFA's Technical Program Committee for the 2017 OFA Workshop, an annual workshop that highlights the latest in research, development, and deployment of HPC interconnects.

UnifyCR: OLCF staff participate in a collaborative ECP project to develop a production-grade user-level file system, UnifyCR, that is highly-specialized for shared file access on HPC systems with distributed burst buffers (e.g., Summit at ORNL and Sierra at LLNL systems, as well as similar future systems). UnifyCR is a collaborative design and implementation project between LLNL and ORNL. Once completed, it will allow users to integrate UnifyCR either through resource managers or via I/O and checkpoint/restart libraries for transparent use by applications. UnifyCR is in its second year of development. Technology Integration group members Sarp Oral, Feiyi Wang, Hyogi Sim and Joseph Moore participate in this effort.

OpenSFS**:** OLCF participates in the Open Scalable File Systems (OpenSFS) community organization, promoting open-source parallel file systems, such as Lustre. OpenSFS's goal is to unite the Lustre end-users and vendors under the same organization and collect and orchestrate Lustre end-users' needs and requests and communicate that to vendors. It is also responsible for organizing the Lustre User Group (LUG) meeting annually. LUG is the largest gathering of Lustre community to discuss current and future issues and development plans for Lustre. OpenSFS also hosts the Lustre Working Group (LWG), which acts as the technical clearing house and decision-making body on all technical issues related to Lustre. OLCF contributes significantly to the Lustre community and development. It has shaped the development of several efforts, such as Progressive File Layouts (PFL), Imperative Recovery (IR), Distributed Name Space (DNE). Currently ORNL is represented at OpenSFS by Sarp Oral as the President, and Dustin Leverman as the LWG co-chair.

Scientific Data Management**:** Sudharshan Vazhkudai, OLCF's Technology Integration group lead, is a Co-PI on a DOE award that delves into the development novel storage systems concepts for future HPC systems. In this project, he has explored the construction and integration of advanced data management services such as analysis-awareness, tagging, searching, metadata extraction from datasets, indexing and active operations into the HPC file systems themselves instead of being decoupled external services. Such an approach helps bridge the gap between file systems and such higher-level data management services and enables future file systems to be more adept at serving scientific queries. This work has resulted in the following publication:

Hyogi Sim, Youngjae Kim, Sudharshan S. Vazhkudai, Geoffroy R. Vallee, Seung-Hwan Lim, Ali R. Butt, ``TagIt: An Integrated Indexing and Search Service for File Systems," *Proceedings of Supercomputing 2017 (SC17): 30th Int'l Conference on High Performance Computing, Networking, Storage and Analysis*, Denver, CO, November 2017.

5. RISK MANAGEMENT

CHARGE QUESTION 5: Is the facility effectively managing operational risks?

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 2017 are listed and described in this section. Planned mitigations and implementations are included in the subsequent descriptions. As of this writing, the OLCF has one high-priority operational risk.

- Decreasing availability of spare parts as Titan ages. The OLCF addressed this risk during this reporting period by continuing its partnership with Cray and NVIDIA to analyze historical failures and monitoring the proactively replaced GPU components to maintain Titan's high level of productivity.
- More information on this risk can be found in Sections 5.2, 5.3, and 5.6 of this report.

5.1 RISK MANAGEMENT SUMMARY

The OLCF's Risk Management Plan describes a regular, rigorous, proactive, and highly successful review process that is reviewed at least annually and 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 for the projects and are attended by the federal project director, facility management, OLCF group leaders, and others subject matter experts and risk owners. Operational risks are discussed in the weekly NCCS Operations Meeting attended by the risk owners, facility management team, OLCF group leaders, and other stakeholders. 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, 27 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 issues as safety, funding and expenditures, and staffing. The specific, more focused risks are concerned with reliability, availability, and use of the system or its components (e.g., the computing platforms, power and cooling infrastructure, storage, networks, software, and user support). In addition to operational risks, at the time of this report, there are 38 tracked risks for the OLCF-4 project and 54 risks tracked for the OLCF-5 project.

The costs of 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 2017

Table 5.1 contains the major risks tracked for OLCF operations in 2017. The full OLCF operations risk register is available on request. The selected risks are all rated medium or high in impact.

Table 5.1. 2017 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 an increase in the number of applications to the INCITE 2018 program. Additionally, the 62 letters of intent for the Summit Early Science Program signify the continued move by application scientists to 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) level 4 requirement for privileged access to OLCF resources.
906: Programming environment tools may be insufficient	Low/medium	Mitigating	The OLCF leverages the expertise of the NCCS Software Tools team 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 team.
917: Robust support will not be available to ensure portability of restructured applications	Medium/medium	Mitigating	The OLCF deployed multiple compilers which maximize 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
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 2017. 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 2017, the data center customer base remained static, and slightly fewer systems were deployed. As 2018 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 2017. 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 continued to impact user jobs. Components were proactively replaced with new parts that are not susceptible to the discovered issue. The OLCF continued the successful work with vendor partners from CY 2016 to execute the 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	The OLCF 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.		

The OLCF continues to have a limited supply of GPU parts to replace failures. This supply is bolstered by a supply of predicted failure parts that were proactively replaced due to the component failure discovered and resolved in CY 2016.

5.4 RISKS RETIRED DURING THE CURRENT YEAR

There were no risks retired from OLCF operations during this reporting period.

5.5 NEW OR RECHARACTERIZED RISKS SINCE LAST REVIEW

5.5.1 Recharacterized Risks

The status or impact of the following risks changed during the reporting period.

Risk No. 721	Lustre metadata performance continues to impact applications		
Risk owner	Sudharshan S. Vazhkudai		
Status	Mitigating**—Current**		
Probability	Medium → Low		
Impact	***Cost:*** Low	***Schedule:*** Low	***Scope/Tech:*** Medium
Mitigations	The OLCF deployed Lustre version 2.8 in CY 2017, which had metadata improvements in its release. The measured improvement was 2× that of the previously deployed Lustre version 2.7. Due to this software change, the OLCF has decided against deploying the Distributed Namespace (DNE) Phase 2 as it would require significant changes and testing time that would reduce the availability of Titan.		
Triggers	Direct observations of application performance decreases by users or staff.		

Risk No. 722	**Safety – personal injury**		
Risk owner	Stephen McNally		
Status	Accept**—Current**		
Probability	Low → Medium		
Impact	***Cost:*** Low	***Schedule:*** Medium	***Scope/Tech:*** Low

Mitigations	In CY 2017, the OLCF introduced two new roles to provide additional support in reducing the potential of encountering this risk and the impact of this risk should it occur. The OLCF, in conjunction with ORNL, introduced a Division Electrical Safety Officer (DESO) and, based on observations of another DOE facility project, an installation manager. The installation manager is responsible for managing the activities by the various subcontractors and ORNL craft support in the datacenter where OLCF-4 is being installed. The OLCF monitors worker compliance with existing safety requirements, conducts safety meetings daily when work is being conducted, performs periodic surveillances using independent safety professionals, and encourages stop-work authority for all staff. Any indication of a negative trend receives prompt intervention from management.
Triggers	Any safety incident or observation is taken seriously and will trigger additional mitigation activities that will be determined by the analysis of the situation.

Risk No. 948	Lack of adequate facility for exascale system
Risk owner	James H. Rogers
Status	Accept**—Current**
Probability	Low → Medium
Impact	***Cost:*** High ***Schedule:*** High ***Scope/Tech:*** Medium
Mitigations	The OLCF has a plan to house the OLCF-5 system in Building 5600 by moving other systems to alternate locations. The preferred approach would be to construct a new building meeting the requirements of an exascale system, but the Office of Management and Budget (OMB) has rejected third-party financing as a method of constructing such a facility. The CD-1 review committee approved a cost range that would support performing the work to house the system in Building 5600.
Triggers	Based on CORAL-2 Request for Proposal (RFP) responses that the space in 5600 would not meet the requirements of the chosen system.

5.5.2 New Risks in This Reporting Period

The following risks were created and tracked during CY 2017. They are included with their risk creation date, mitigations, and triggers.

Risk No. 1240	Failure to handle Export Controlled Information (ECI) properly
Risk creation date	2017-12-01
Risk owner	Ryan Adamson
Status	Mitigating**—Current**
Probability	Medium
Impact	***Cost:*** Low ***Schedule:*** Low ***Scope/Tech:*** Low ***Other:*** Medium
Mitigations	Staff with elevated privileges on system where ECI can be accessed go through annual training and refreshers on how to handle ECI. Project Principal Investigators (PIs) are subject to an initial project briefing where an ORNL Export Control Analyst describes the categorization of the project based on the responses in the application for resources at the OLCF. OLCF Cyber Security staff participate in this briefing. The OLCF has automated permission enforcement on storage resources/areas where ECI is can be stored.
Triggers	Information/indications that ECI information has been released.

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	Low
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 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.

Risk 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. The 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.

Risk No. 1154	Decreasing availability of spare parts as Titan ages
Risk owner	Don E. Maxwell
Status	Mitigate
Probability	High
Impact	***Cost:*** Medium ***Schedule:*** Low ***Scope/Tech:*** High
Mitigations	The OLCF 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 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 the acceptance milestone is achieved in the OLCF-4 project in Q3 CY 2018. They are listed with their mitigations and triggers.

Risk No. 1092	Water damage to equipment from overhead water piping
Risk owner	Bart A. Hammontree
Status	Mitigating
Probability	Low
Impact	***Cost:*** Medium ***Schedule:*** Low ***Scope/Tech:*** Low
Mitigations	Catch pans are installed where overhead routing is needed to redirect water away from sensitive equipment. Leak detection systems are being installed to alert operators when a leak has occurred.
Triggers	Information on potential for or actual water damage.

Risk No. 1136	LSF does not provide needed features
Risk owner	Kevin G. Thach
Status	Mitigating
Probability	Low
Impact	***Cost:*** Medium ***Schedule:*** Low ***Scope/Tech:*** Low
Mitigations	The OLCF-4 project team continues to work with IBM to ease the identification of problem areas associated with the use of early access, TDS, and early accepted hardware.
Triggers	Information from project, acceptance, early science, and operations teams that LSF is not performing as required.

Risk No. 1149	Burst buffer implementation does not perform as expected, resulting in low user adoption
Risk owner	Sudharshan S. Vazhakudai

Status	Mitigating
Probability	Medium
Impact	***Cost:*** Low ***Schedule:*** Low ***Scope/Tech:*** Low
Mitigations	Continuing user education about necessary modifications to applications required to utilize the burst buffer.
Triggers	Information from users about file system performance may point to lack of adoption of the burst buffer.

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.

5.6.2 Risks for OLCF-4 Operations That Will Begin in CY 2018

These risks that have already been identified but will not be in effect until the acceptance milestone is achieved in the OLCF-4 project in Q3 CY2018. They are listed with their mitigations and triggers.

Risk No. 1079	OLCF-4 post deployment hardware and software issues
Risk owner	Arthur S. Bland
Status	Accept
Probability	Medium
Impact	***Cost:*** Low ***Schedule:*** Low ***Scope/Tech:*** Low
Mitigations	The OLCF team will work closely with the vendor to track and identify the root cause of early failures as quickly as possible to develop plans for remediation if required.
Triggers	Problems during the installation and acceptance activities.

Risk No. 1157	Limited or incomplete information relative to system load and incomplete integration with control systems will contribute to higher operating costs
Risk owner	James H. Rogers
Status	Mitigating
Probability	Medium
Impact	***Cost:*** Medium ***Schedule:*** Low ***Scope/Tech:*** Low
Mitigations	During acceptance and early science activities, normal operating conditions will be identified that can be correlated to mechanical energy plant (MEP) flow rates; the control system will be adjusted to use these as the basis for operations. Consider mechanisms for introducing additional functionality over time, ensuring that the OLCF-4 system is adequately protected but seeking greater MEP efficiency.

Triggers	Monitor the development of the control system integration software stack, ensure a proper application programming interface (API) compatibility between layers. Ensure that monitoring and operations staff are in place for this new service that is well integrated with the programmable logic controller (PLC) system.

Risk No. 1193	Early Science Program for OLCF-4 terminated before completion
Risk owner	Jack C. Wells
Status	Mitigating
Probability	Low
Impact	***Cost:*** Low ***Schedule:*** Low ***Scope/Tech:*** Low
Mitigations	Continual evaluation of operational metrics, budget scenarios, and discussions with program sponsors. Annual user survey results and operational assessments of key user program delivery metrics.
Triggers	Indications of failure to achieve operational metrics and/or user dissatisfaction.

Risk No. 1194	Water quality declines during OLCF-4 operations
Risk owner	Stephen McNally
Status	Mitigating
Probability	Low
Impact	***Cost:*** Low ***Schedule:*** Low ***Scope/Tech:*** Low
Mitigations	The ORNL Facilities Management Division has a contract with a vendor to provide deionized water quality control from their local offices in Knoxville. Medium-temperature-water (MTW)loop and system components are all made from nonferrous materials, which greatly reduces the particulate matter risk. Additional mitigation is provided by straining 100% of the flow through the MTW loop.
Triggers	Measurements that the water quality has declined or that particulate is present in the strainers.

Risk No. 1225	20MW MEP PLC controls implementation requires significant improvement from baseline/project implementation
Risk owner	Stephen McNally
Status	Mitigating
Probability	Low
Impact	***Cost:*** Low ***Schedule:*** Low ***Scope/Tech:*** Low
Mitigations	Tightly integrated activity between ORNL Mechanical Utilities Division, the PLC programming subcontractor, ORNL Facilities Management Division, and the OLCF to monitor the performance of the PLC system as operating conditions become more broadly understood. Add systems and engineering that integrate information available to the computer (scheduler information, individual node power and temperature data) and information available to the PLC (cooling capacity, maintenance state). Continuously examine the effect of changes and improvements to the PLC programming to introduce improved methods for managing the different components of the MEP and improving the operating efficiency of the mechanical systems.
Triggers	Telemetry from the MEP and PLC systems.

6. SITE OFFICE SAFETY METRICS

CHARGE QUESTION 6: Does the facility exhibit a culture of continuous improvement in Environment, Safety, and Health (ES&H) practices to benefit staff, users, the public, and the environment? Has the facility implemented appropriate ES&H measures?

OLCF RESPONSE: Yes.

6.1 SUMMARY

ORNL is committed to operating under the DOE safety regulations specified in 10 CFR 851,[*] "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 and Integrated Performance Management and the 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 following feedback from the site office staff in the Operations and Oversight Division about OLCF's safety culture.

[*] 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.

- A review of the monthly safety charts and the total recordable cases and days away, restricted, or transferred (DART) summary documents indicated that the overall FY 2017 total recordable cases rate, and DART rates significantly decreased to the lowest rates since UT-Battelle began managing ORNL. Operations of the OLCF in the NCCS remained safe, efficient, and effective as there were zero total recordable cases and DARTs in FY 201.
- UT-Battelle continued to emphasize the subject of electrical safety. In addition, UT-Battelle emphasized safety culture through employee participation in the Battelle Laboratory Operations Supervisor Academy (LOSA) and through focusing on the effectiveness of the Lab Space Manager program, specifically safety culture in laboratory spaces. UT-Battelle provided focus on utilization of the Safe Conduct of Research principles for safe work in the laboratory. Finally, in FY 2017, emphasis was placed upon the UT-Battelle line mangers' use of the Safety Talks Program. Safety Talks is the tool managers should use to engage employees during walk-through visits. The Safety Talks system utilizes the Safe Conduct of Research principles as talking points. Safety Talk participation by OLCF line managers was encouraged by OLCF executive management. The OLCF statistics concerning participation were commonly discussed in the OLCF project and operations meetings.
- The OLCF-4 project Health and Safety Plan and the Hazard Analysis were reviewed and revised in 2017.
- In March 2017, the OLCF experienced an electrical safety event in one of the managed data center spaces. In this event, an OLCF subcontractor cut an energized electrical component cable. Work was immediately stopped, and a full fact finding/critique was conducted. The event was a reportable occurrence through the DOE Occurrence system. In addition to the occurrence, a corrective action plan and work plan were developed and communicated to OLCF onsite subcontractors. This work plan is the work control document for OLCF onsite subcontractors/vendors and sets appropriate expectations and requirements for work control and electrical safety. It is important to note that no one was harmed during this incident, and it provided a great opportunity to improve existing operations to help stave off a more serious event.
- In July 2017 the OLCF, through a purchase service agreement with the Safety Services Division, staffed the OLCF-4 project with an installation manager. This was an observed best practice and lessons learned from Argonne Leadership Computing Facility-2 (ALCF-2) supercomputer installation project (called a "floor manager" in the ALCF-2 project). The need for an installation manager was identified through the ALCF-2 lessons learned and from the March 2017 OLCF subcontractor electrical event. The OLCF installation manager chosen for the project has over 25 years of experience as a health and safety professional within DOE construction, demolition, and operations projects. The OLCF installation manager is responsible for the day-to-day safe installation of the OLCF Summit system. This role coordinates all installation activities with the installation vendor contractor, mechanical and electrical subcontractors, and UT-Battelle conventional craft through plan-of-the-day meetings, hazards analysis, and general work coordination within the construction and installation areas. The installation manager is also responsible for the orientation and mentoring of subcontractor personnel to the UT-Battelle safety culture and expectations. OLCF management is committed to providing funding for this role as part of their commitment for safe operations.

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.

7. CYBER SECURITY

CHARGE QUESTION 7: Does the facility exhibit a culture of continual improvement in cyber security practices? Have innovations been implemented to improve the facility's cyber security posture? Does the facility have a valid cyber security plan and Authority to Operate?

OLCF RESPONSE: Yes. The OLCF maintains a strong culture of continuous operational improvement, especially in the area of cyber security. The technical staff track and monitor for existing threats and vulnerabilities to assess the risk profile of the OLCF operation. The facility is committed to innovating in this area through development of open-source tools and employing cutting-edge practices that enhance the operation while not increasing the OLCF's risk profile. The most recent OLCF Authority to Operate was granted on February 21, 2017, and is managed through an ongoing authorization process and has no authorization termination date set (Figure 7.1).

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 ORNL certification and accreditation package currently uses *Recommended Security Controls for Federal Information Systems and Organizations*[*] 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 INSTALLATION OF A SCALABLE BRO-BASED INTRUSION DETECTION CAPABILITY

The OLCF has replaced its Snort intrusion detection system with a much more capable Bro installation. Using multiple in-line taps for OLCF's 100 and 40 gigabit links, traffic is aggregated and then dispatched from a set of Arista switches to a scalable set of Bro worker nodes. Bro provides rich session metadata to our existing security information and event management (SIEM) for security analysis and alerting capabilities. This infrastructure augments the existing NetFlow and

[*] National Institute of Standards and Technology, *Recommended Security Controls for Federal Information Systems and Organizations*, Special Publication 800-53, revision 3, US Department of Commerce, Joint Task Force Transformation Initiative, August 2009.

[†] *Federal Information Processing Standards Publication 199*, Standards for Security Categorization of Federal Information and Information Systems, Computer Security Division, Information Technology Laboratory, National Institute of Standards and Technology, Gaithersburg, MD, February 2004.

data transfer metrics gathering tools to provide a much clearer picture of how the external and internal OLCF networks are being used by user and system processes.

Department of Energy

ORNL Site Office
P.O. Box 2008
Oak Ridge, Tennessee 37831-6269

March 22, 2017

Mr. Kevin A. Kerr
Information Systems Security Manager
Oak Ridge National Laboratory
UT-Battelle, LLC
Post Office Box 2008
Oak Ridge, Tennessee 37831-6045

Dear Mr. Kerr:

AUTHORIZATION DECISION DOCUMENT FOR OAK RIDGE NATIONAL LABORATORY (ORNL) SUPERCOMPUTING ENCLAVE

Reference: Letter from Kevin A. Kerr to Johnny O. Moore, subject, *Contract DE-AC05-00OR22725, ORNL Supercomputing Enclave Approval to Operate*, dated February 21, 2017

As the Authorizing Official, I have reviewed the referenced request. The ORNL Supercomputing Enclave is authorized to operate. No additional conditions outside the substance of the request are required.

The information system is now being managed by an ongoing authorization process, thus an authorization termination date is not set. I accept the responsibility for performing all necessary activities associated with the ongoing authorization process.

If there are any questions or additional information is required, please contact John Young at (865) 576-7471 or youngjc1@ornl.gov.

Sincerely,

Johnny O. Moore, Manager
ORNL Site Office

Enclosure

cc w/enclosure:
Mike E. Bartell, ORNL
Amy D. Nuckols, ORNL
Neil Masincupp, SC-OR
Martha J. Kass, SC-OSO
John C. Young, SC-OSO

Document/Material Transmitted Contains Official Use Only Information.

When separated from enclosure, this document **does not** contain Official Use Only.

Figure 7.1. OLCF Authority to Operate.

7.3 DEVELOPMENT OF A PUBLIC KEY–BASED PASSWORD DISTRIBUTION SYSTEM

A Public Key-Based Password Distribution System (PKPass), as initially mentioned in the 2016 OLCF OAR, was submitted to UT-Battelle for approval to open-source the code base. It has also been modified to check certificate validity, as well as to provide electronic signature creation and validation when passwords or secrets are distributed to recipients. A strong unit test framework has also been developed.

7.4 UNDERPINNING AN ESI-DAK SYSTEM FOR EXTREME DATA CORRELATION

The OLCF needed a comprehensive, real time, highly performant, secure, and scalable information correlation system in order to meet the emerging needs of the HPC operations, cyber security, and user assistance functions. This system will augment strategic business analytics capabilities and reporting as well as real-time tactical response capabilities. OLCF has chosen the Apache Kafka distributed streaming platform, which will provide an ESI-DAK capability.[*] Among other features, researchers at ORNL who are interested in this event data will be able to directly access events without reaching to operations personnel for access or explanation. The early significance of this work has been demonstrated.[†] At this point, Bro IDS, Alpine, and Alpine-TDS are streaming system metrics to Kafka.

To organize and understand all of the information flowing through the system, a Data Dictionary has been created. This dictionary allows administrators to describe message schema and allows potential consumers of information to help themselves in making discoveries without any of the delays associated with contacting the subject matter experts for information.

7.5 OPENSHIFT PLATFORM HARDENING AND SECURITY RISK MITIGATION

OLCF's Platform As A Service offering using RedHat's Openshift platform underwent heavy development in 2017 and entered a pilot phase with multiple internal and external users. Because Openshift uses Kubernetes under the hood, the OLCF HPC Operations team was able to make several modifications to provide significant user-level boundary protections to prevent processes from accessing unauthorized data or submitting unauthorized jobs. The Openshift offerings will continue to expand in CY 2018 and look to be a source of tremendous innovation as the center continues to partner with new projects and experimental facilities.

7.6 CREATION OF AN OPEN PRODUCTION SECURITY ENCLAVE

In 2017, NCCS created more precise cyber security enclave descriptions that enabled the center to create new production and development enclaves to support additional scientific workflows. Figure 7.2 shows the new cyber security quadrant model that began operation in CY 2017. As shown, the model lists four enclaves: moderate production, moderate development, open production, and open development. Understanding the subtleties between both production and development as well as moderate and open was critical in differentiating the four enclaves. Each enclave supports a specific mission for the center. Moderate production is where Titan, Eos, Rhea, and other production OLCF systems reside. Identifying the support and services differences between moderate production and moderate development was where this new model began. Traditionally, the NCCS operated all moderate systems at the same service level. However, research needs drove a further examination into this traditional practice and revealed that

[*] US Department of Energy, ASCR Cybersecurity for Scientific Computing Integrity—Research Pathways and Ideas Workshop, June 2–3, 2015, LBNL-191105, 2015. https://www.orau.gov/integrity2015/ASCR_CYBERSECURITY_RESEARCH_PATHWAYS_AND_IDEAS_REPORT.pdf

[†] B. H. Park et al., "Big data meets HPC log analytics: Scalable approach to understanding systems at the extreme scale," *Proceedings of the IEEE International Conference on Cluster Computing, ICCC*, Article 8049013, Volume 2017, pp. 758–765, September 2017.

several systems such as the testbeds and other research systems do not require 24/7 support and heavy configuration and change management processes, as they need to be flexible to fulfill their needs. This separation became key to then further defining sub-levels within a newly discussed open enclave.

Matching the differing needs of moderate production and development, the open production and development enclaves were defined and put into operation. Within the open enclaves, the NCCS was able to provide rapid account provisioning leveraging ORNL's XCAMS account management infrastructure for external users. Users in the open enclaves require simple and highly flexible access to compute and data resources for unprecedented workflows within the center. These open enclaves have also enabled easy access to next-generation testbed resources for research teams to build libraries and software on systems such as Arm architectures as a possible disruptive computing platform for exascale or beyond.

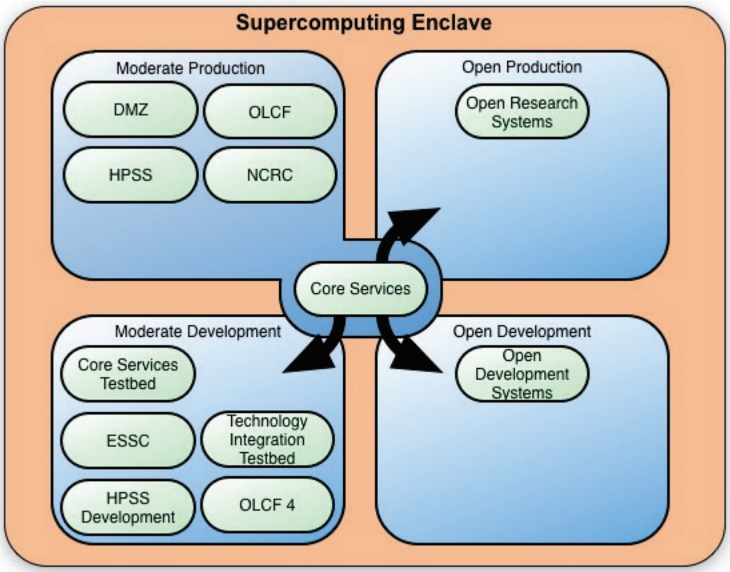

Figure 7.2. NCCS Supercomputing Enclaves.

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? Are there other metrics that would have helped you assess the facility's operational performance?

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

8.1 SUMMARY

Table 8.1 provides a summary of the metrics and actual data for the current reporting period and proposed metric targets for the following calendar year. It is important to note that Summit will not officially enter into production until January 1, 2019.

Table 8.1. OLCF metrics and actual data for 2017 and proposed metrics and target for 2018 (2016 data provided for context)

2016 metric and target	2016 target	2016 actual	2017 target	2017 actual	2018 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 2016 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: 93%.	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	2016 target	2016 actual	2017 target	2017 actual	2018 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 1 year following a major system upgrade, the targeted scheduled availability is 85% and overall availability is 80%)[a]					
Scheduled Availability. 2016 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: 99.64%; HPSS: 99.89%; /atlas0: 99.92%; /atlas1: 99.92%.	Titan: 95%; HPSS 95%; external file systems 95%. (Monthly)	The OLCF exceeded the metric target. Titan: 99.39%; HPSS: 99.46%; /atlas1: 99.59%; /atlas2: 99.59%.	Titan: 95%; HPSS 95%; external file systems 95%.
Overall Availability. 2016 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: 97.03%; HPSS: 98.54%; /atlas0: 98.2%; /atlas1: 98.2%.	Titan: 90%; HPSS 90%; external file systems: existing, 90%. (Monthly)	The OLCF exceeded the metric target: Titan: 98.09%; HPSS: 98.87%; /atlas1: 98.87%; /atlas2: 98.88%.	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 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. (Monthly)	The capability usage was 59.81%. 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]	56.96%	N/A[b]	60.46%	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 2016 OPERATIONAL ASSESSMENT REVIEW

In March 2017, the OLCF presented the 2016 operational activities of the center to the DOE sponsor.

During this review, the review committee did not identify any recommendations.

APPENDIX B. TRAINING, WORKSHOPS, AND SEMINARS

Event Type	Event Title	Date	Participants
Monthly User Conference Calls	2017 OLCF User Conference Call: VisIT	1/25/2017	35
Monthly User Conference Calls	2017 OLCF User Conference Call: Continuous Integration	2/22/2017	45
Monthly User Conference Calls	2017 OLCF User Conference Call: Getting Started with ADIOS: Part 2	3/29/2017	17
Monthly User Conference Calls	2017 OLCF User Conference Call: Preparing Applications for Summit: NUCCOR	4/26/2017	42
Monthly User Conference Calls	2017 OLCF User Conference Call: Rand pbdR	6/28/2017	22
Monthly User Conference Calls	2017 OLCF User Conference Call: Summit--Transitions to Operations (T2O) Plan	7/26/2017	69
Monthly User Conference Calls	2017 OLCF User Conference Call: OLCF Data DOI Service	8/23/2017	15
Monthly User Conference Calls	2017 OLCF User Conference Call: Using Containers	9/27/2017	58
Monthly User Conference Calls	2017 OLCF User Conference Call: CUDA 8 Coming to Titan	10/25/2017	23
Workshop/Training/Meeting	ALCF/OLCF Joint User Meeting at SC17	11/15/2017	19
Seminar Series	BEAM: An HPC Pipeline for Nanoscale Materials Analysis and Neutron Data Modeling	1/24/2017	-
Seminar Series	How to Create Digital Object Identifiers (DOI)	1/31/2017	-
Seminar Series	Parallel Processing of Video Decoder and Binary Translation	2/3/2017	-
Seminar Series	ORNL MiniApps Bluejeans Seminar Series: The XrayTrace MiniApp	2/9/2017	-
Seminar Series	The Multi-dimensional Character of Nucleosnthesis in Core-collapse Supernovae	2/10/2017	-
Seminar Series	The Divide-Expand-Consolidate (DEC) Local Coupled Cluster Scheme	2/13/2017	-
Seminar Series	Model Prediction with Effective Uncertainty Modeling, Quantification, and Bias Corrections	2/15/2017	-
Seminar Series	ORNL Deep Learning: Age, Gender and Fine-Grained Ethnicity Prediction for the East Asian Dataset	2/16/2017	-
Seminar Series	Designing Next-Generation Programming Models using Lessons Learned from PGAS	3/17/2017	-

Event Type	Event Title	Date	Participants
Seminar Series	ORNL MiniApps Bluejeans Seminar Series: MiniSweep Deep Dive	3/27/2017	-
Seminar Series	Study of Shock-Driven Ejecta Using Continuum Hydrodynamics Simulations	3/30/2017	-
Seminar Series	Imaging in the Information Dimensions	3/30/2017	-
Seminar Series	Thermal Nonequilibrium Models for High-Temperature Reactive Processes	4/7/2017	-
Seminar Series	ORNL MiniApps Webinar Series: Ziz-An Extensible MiniApp for Astrophysical Multiphysics	4/11/2017	
Seminar Series	ORNL MiniApps Webinar Series: Ziz Miniapp Deep Dive	4/18/2017	-
Seminar Series	Providing Customer Support with Custom Script Development in HPC Environments	4/24/2017	-
Seminar Series	Dynamics of Stratified Flow Past a Sphere: Simulations Using Body-inclusive Numerical Model	4/27/2017	-
Seminar Series	Development and Implementation of a Computationally Efficient Treatment of Three-Body Interactions in HPC Solid ^{4}He	4/28/2017	-
Seminar Series	Parallelization of Multireference Perturbation Corrections MRCISD(TQ) and GVVPT2 in Electronic Structure Theory	5/11/2017	-
Seminar Series	DBZ2: Parallel Compression and Decompression Algorithm using Distributed File Tree Walk	8/18/2017	-
Workshop/Training/Meeting	ORNL MiniApps Webinar Series: OpenACC Implementation of miniSweep	8/10/2017	-
Seminar Series	HPC Container Techology at the OLCF	8/23/2017	-
Seminar Series	Effect of Crystal Packing on the Electronic Properties of Molecular Crystals	9/1/2017	-
Seminar Series	Flash Memory Summit 2017 Recap	9/14/2017	-
Seminar Series	Project Deep Sea	9/29/2017	-
Seminar Series	OLCF Seminar Series: NVIDIA Volta Architecture and CUDA 9	10/18/2017	-
Seminar Series	TagIt: An Integrated Indexing and Service for File Systems	10/27/2017	-
Seminar Series	GUIDE: A Scalable Information Directory Service to Collect, Federate, and Analyze Logs for Operational Insights Into a Leadership HPC Facility	11/3/2017	-
Seminar Series	Service Oriented Architecture Approach to Supercomputing	11/6/2017	-
Seminar Series	Scientific User Behavior and Data-Sharing Trends in a Petascale File System	11/10/2017	-

Event Type	Event Title	Date	Participants
Seminar Series	Large Scale Modeling Organic Photovoltaics	12/1/2017	-
Seminar Series	Can Non-orthogonal Configuration Interaction Be Efficient?	12/1/2017	-
Seminar Series	Large-scale Matrix Diagonalization Methods for Studying Strongly Correlated Systems	12/14/2017	-
IDEAS/Facility Webinar Series	IDEAS/ECP Seminar Series (Session 1): Python in HPC	6/7/2017	129
IDEAS/Facility Webinar Series	IDEAS/ECP Seminar Series (Session 2): OpenMP	6/28/2017	59
IDEAS/Facility Webinar Series	IDEAS/ECP Seminar Series (Session 3): Intermediate Git	7/12/2017	79
IDEAS/Facility Webinar Series	IDEAS/ECP Seminar Series (Session 4): ECP Roofline Webinar	8/16/2017	71
IDEAS/Facility Webinar Series	IDEAS/ECP Seminar Series (Session 5): Scalable Node Programming with OpenACC	8/20/2017	46
IDEAS/Facility Webinar Series	IDEAS/ECP Seminar Series (Session 6): Managing Defects in HPC Software Development	11/1/2017	63
IDEAS/Facility Webinar Series	IDEAS/ECP Seminar Series (Session 7): Tau Performance System Webinar	11/8/2017	29
IDEAS/Facility Webinar Series	IDEAS/ECP Seminar Series (Session 8): Better Scientific Software Webinar	12/6/2017	108
Workshop/Training/Meeting	Summitdev Training	1/10-12/17	87
Workshop/Training/Meeting	OLCF Getting Started Tutorial #1	1/13/2017	28
Workshop/Training/Meeting	2017 Hackathon #1 - Jülich	3/6-10/17	59
Workshop/Training/Meeting	PanDA@ORNL Meeting	3/30-31/17	25
Workshop/Training/Meeting	INCITE Proposal Writing Webinar #1	4/19/2017	24
Workshop/Training/Meeting	Performance Portability with Kokkos Training	5/2-4/17	28
Workshop/Training/Meeting	INCITE Proposal Writing Webinar #2	5/18/2017	72
Workshop/Training/Meeting	OLCF Getting Started Webinar #2	5/22/2017	20
Workshop/Training/Meeting	2017 OLCF User Meeting	5/23-25/17	117
Workshop/Training/Meeting	2017 Hackathon #2 - Brookhaven	6/5-9/17	68
Workshop/Training/Meeting	INCITE Proposal Writing Webinar #3	6/6/2017	36
Workshop/Training/Meeting	Introduction to CUDA	6/19/2017	46
Workshop/Training/Meeting	Scaling to Petascale Institute 2017	6/26-30/17	27
Workshop/Training/Meeting	DOE Computation Science Graduate Fellowship: Mini-GPU Hackathon	7/26/2017	40
Workshop/Training/Meeting	2017 Hackathon #3 - NASA	8/21-25/17	67
Workshop/Training/Meeting	JSRUN Tutorial	8/30/2017	52

Event Type	Event Title	Date	Participants
Workshop/Training/Meeting	Accelerated Data Analytics and Computing Webinar: Tools for Moving to Exascale	9/5/2017	59
Workshop/Training/Meeting	DLI: Introduction to Deep Learning	9/18/2017	-
Workshop/Training/Meeting	2017 Hackathon #4 - ORNL	10/9-13/17	76
Workshop/Training/Meeting	NCCS OpenShift Introduction	10/12/2017	-
Workshop/Training/Meeting	Deep Learning Expo	12/14/2017	90
Workshop/Training/Meeting	Introducing the Summit Early Science Program Webinar	12/21/2017	-

APPENDIX C. OUTREACH PRODUCTS

Date	Type of Product	Title
5/24/2017	Article	Earth Scientists Shake Things up on Titan
8/1/2017	Article	Container Revolution Enables Science Breakthroughs
9/1/2017	Article	Mastering a Bacterial Photosynthetic System
12/20/2017	Article	As Exascale Frontier Opens, Science Application Developers Share Pioneering Strategies
5/9/2017	Fact Sheet	VA-MVP Program Fact Sheet
5/9/2017	Fact Sheet	Summit Fact Sheet
11/1/2017	Fact Sheet	INCITE Fact Sheets
12/18/2017	Fact Sheet	Update OLCF Fact Sheet
5/3/2017	Graphic	Puzzle Pieces for Suzy's slide
5/10/2017	Graphic	CCSD Exhibit Backdrop
5/11/2017	Graphic	Certificate of Appreciation for User Meeting
6/28/2017	Graphic	Caution Signs for Summit Area
7/6/2017	Graphic	Summit 3D Rendering
7/10/2017	Graphic	Awards Night Materials
7/13/2017	Graphic	Summit Logo and Branding Concept
7/13/2017	Graphic	Summit Wall Wrap "Overlook"
8/29/2017	Graphic	Smoky Mountain Conference Graphics
9/22/2017	Graphic	Science Images (and Text) for SC ORNL Poster
9/27/2017	Graphic	Summit Sticker for Tapia/GHC
10/20/2017	Graphic	OLCF 25th Anniverary Button
11/7/2017	Graphic	OLCF 25th Anniversary Promotion (web banners, social elements, etc.)
11/7/2017	Graphic	25th Anniversary Event Signage
11/13/2017	Graphic	Business Card for SC17 (Summit/OLCF)
12/8/2017	Graphic	Frontier Logo and Branding
1/4/2017	Highlight	ORNL–Oracle Collaboration Offers Data Safeguard
1/4/2017	Highlight	OLCF Active in New User-Oriented OpenSFS
1/4/2017	Highlight	UTRC Researchers Take Flight with Graphics Processing Units for Large Eddy Simulations
1/4/2017	Highlight	Mini Hackathon Offers Shorter Schedule, Introductory Focus
1/31/2017	Highlight	Researchers Flip Script for Li-Ion Electrolytes to Simulate Better Batteries
1/31/2017	Highlight	Longtime User Requests Added to Next Version of OpenACC
1/31/2017	Highlight	ORNL Resources Help BigNeuron Tackle Grand Challenges
1/31/2017	Highlight	The Shape of Melting in Two Dimensions
1/31/2017	Highlight	Straatsma Named AAAS Fellow in Chemistry
2/2/2017	Highlight	Supercomputing, Experiment Combine for First Look at Magnetism of Real Nanoparticle
2/28/2017	Highlight	OLCF-Connected Team Wins Best Paper at Industrial Conference
2/28/2017	Highlight	Modeling Autoignited Combustion to Drive Engines Forward
2/28/2017	Highlight	Ascending to Summit: Announcing Summitdev

Date	Type of Product	Title
3/28/2017	Highlight	Researchers Shoot for Success with Simulations of Laser Pulse–Material Interactions
3/28/2017	Highlight	OLCF Staff Develops Digital Object Identifier Framework to Facilitate Open Access to Datasets
3/28/2017	Highlight	Multitasking Framework Accelerates Scientific Discovery
3/28/2017	Highlight	A Seismic Mapping Milestone
4/18/2017	Highlight	A Real CAM-Do Attitude
4/18/2017	Highlight	Predictive Power
4/18/2017	Highlight	Chemistry Applications Get in Top Shape for Summit
4/18/2017	Highlight	OLCF Conference Calls Equip Users with System Knowledge
4/18/2017	Highlight	Building the Bridge to Exascale
4/18/2017	Highlight	Test System Arms OLCF with Experimental Technology
5/9/2017	Highlight	Containers Provide Access to Deep Learning Frameworks
5/9/2017	Highlight	OLCF Brings Petascale Computing to 2017 APS March Meeting
5/9/2017	Highlight	Supercomputing Mimics Berkelium Experiments to Validate New Find
5/9/2017	Highlight	Assembling Life’s Molecular Motor
5/19/2017	Highlight	OLCF Users to Elect New OUG Executive Board Members
6/5/2017	Highlight	OLCF testing new platform for scientific workflows
6/5/2017	Highlight	OLCF Explores Deep Learning with DGX-1
6/5/2017	Highlight	Jansen Recognized for Development of Method That Enhances Physics Research
6/5/2017	Highlight	Researchers Seek Sigma Meson on the Path to Heavier Hadrons
6/5/2017	Highlight	Putting the Pedal to the Metal in the Hunt for Alloys
6/27/2017	Highlight	Data-Driven Science for HPC Through CADES
6/27/2017	Highlight	Annual User Meeting Spotlights Titan, Summit, and Deep Learning
6/27/2017	Highlight	OLCF Staff Explores Future of Accelerated Computing at GPU Technology Conference
6/27/2017	Highlight	How Hot is Too Hot in Fusion?
7/18/2017	Highlight	Titan Simulations Show Importance of Close 2-Way Coupling between Human and Earth Systems
7/18/2017	Highlight	Foertter Elected SIG HPC Education Vice Chair
7/18/2017	Highlight	3-D Models Help Scientists Gauge Flood Impact
7/18/2017	Highlight	Budiardja Receives Runner-Up Best Paper at Cray User Group Meeting
7/18/2017	Highlight	OLCF Staff Reach Students through Coding Events
8/8/2017	Highlight	ALCC Program Awards 1 Billion Hours on OLCF Resources
8/8/2017	Highlight	Summer Internships Offer Students Hands-On Experience, Mentorship
8/8/2017	Highlight	OLCF Staff Member Shares Work on Improving Simulation Efficiency
8/8/2017	Highlight	OLCF Hosts CUDA Workshop for GPU Programming
8/8/2017	Highlight	Galactic Winds Push Researchers to Probe Galaxies at Unprecedented Scale
8/25/2017	Highlight	Titan Helps Researchers Suck Mystery Out of Cell’s ‘Vacuum Cleaners’
9/19/2017	Highlight	Sodium Shakedown in Dopamine Research
9/19/2017	Highlight	OLCF Provides Research Opportunities, HPC Experience for Postdocs
9/19/2017	Highlight	CADES Supplies Resources, Expertise to World’s Most Powerful Accelerator
10/17/2017	Highlight	Decades-Long Physics Mystery Elucidated with Titan
10/17/2017	Highlight	OLCF Helps GE Deliver Next-Generation Gas Turbines

Date	Type of Product	Title
10/19/2017	Highlight	Faces of Summit: Making Spaces
10/31/2017	Highlight	OLCF Staff Honored for Excellence in Science, Technology, and Support
10/31/2017	Highlight	OLCF Postdoc Fuses the Gap between Experiment and Computation
10/31/2017	Highlight	OLCF to Celebrate 25 Years of HPC Leadership
11/10/2017	Highlight	Ready for 25 More Years
11/13/2017	Highlight	INCITE Grants of 5.95 Billion Hours Awarded to 55 Computational Research Projects
11/28/2017	Highlight	Conferences Promote Greater Diversity in HPC
11/28/2017	Highlight	Teams Gear up for Summit at Fourth Annual GPU Hackathon
11/28/2017	Highlight	Faces of Summit: Modeling Safety
11/28/2017	Highlight	Scaling Deep Learning for Science
12/20/2017	Highlight	New Boeing Method Accelerates Turbulence Modeling Uncertainty Analysis
12/20/2017	Highlight	OLCF Team ‘Connects’ with the High-Performance Computing Community at SC17
12/20/2017	Highlight	As Exascale Frontier Opens, Science Application Developers Share Pioneering Strategies
12/20/2017	Highlight	Faces of Summit: Bursting with Ingenuity
2/1/2017	Poster	ECP Meeting Poster
3/22/2017	Poster	DAC Meeting – OLCF Overview
3/22/2017	Poster	DAC Meeting – Summit
7/13/2017	Poster	CSGF Poster
7/26/2017	Poster	SNS/HFIR/CNMS User Meeting Poster
9/15/2017	Poster	Summit Hallway Poster (near Buddy's Office)
9/15/2017	Poster	Overlook Summit Panel
11/1/2017	Poster	ALCC Poster
12/11/2017	Poster	DLUG Meeting Poster
12/15/2017	Poster	INCITE 2018 Projects Poster
1/4/2017	PPT Slide	ORNL–Oracle Collaboration Offers Data Safeguard
1/4/2017	PPT Slide	OLCF Active in New User-Oriented OpenSFS
1/4/2017	PPT Slide	Mini Hackathon Offers Shorter Schedule, Introductory Focus
1/31/2017	PPT Slide	Longtime User Requests Added to Next Version of OpenACC
1/31/2017	PPT Slide	ORNL Resources Help BigNeuron Tackle Grand Challenges
1/31/2017	PPT Slide	Straatsma Named AAAS Fellow in Chemistry
2/28/2017	PPT Slide	OLCF-Connected Team Wins Best Paper at Industrial Conference
3/28/2017	PPT Slide	OLCF Staff Develops Digital Object Identifier Framework to Facilitate Open Access to Datasets
3/28/2017	PPT Slide	Multitasking Framework Accelerates Scientific Discovery
4/18/2017	PPT Slide	Chemistry Applications Get in Top Shape for Summit
4/18/2017	PPT Slide	OLCF Conference Calls Equip Users with System Knowledge
4/18/2017	PPT Slide	Building the Bridge to Exascale
4/18/2017	PPT Slide	Test System Arms OLCF with Experimental Technology
4/18/2017	PPT Slide	A Real CAM-Do Attitude
4/18/2017	PPT Slide	Predictive Power
5/9/2017	PPT Slide	Containers Provide Access to Deep Learning Frameworks

Date	Type of Product	Title
5/9/2017	PPT Slide	OLCF Brings Petascale Computing to 2017 APS March Meeting
5/9/2017	PPT Slide	Supercomputing Mimics Berkelium Experiments to Validate New Find
5/9/2017	PPT Slide	Assembling Life's Molecular Motor
6/5/2017	PPT Slide	OLCF testing new platform for scientific workflows
6/5/2017	PPT Slide	OLCF Explores Deep Learning with DGX-1
6/5/2017	PPT Slide	Jansen Recognized for Development of Method That Enhances Physics Research
6/5/2017	PPT Slide	Researchers Seek Sigma Meson on the Path to Heavier Hadrons
6/5/2017	PPT Slide	Putting the Pedal to the Metal in the Hunt for Alloys
6/27/2017	PPT Slide	Data-Driven Science for HPC Through CADES
6/27/2017	PPT Slide	Annual User Meeting Spotlights Titan, Summit, and Deep Learning
6/27/2017	PPT Slide	OLCF Staff Explores Future of Accelerated Computing at GPU Technology Conference
7/18/2017	PPT Slide	Foertter Elected SIG HPC Education Vice Chair
7/18/2017	PPT Slide	Budiardja Receives Runner-Up Best Paper at Cray User Group Meeting
7/18/2017	PPT Slide	OLCF Staff Reach Students through Coding Events
7/18/2017	PPT Slide	Titan Simulations Show Importance of Close 2-Way Coupling between Human and Earth Systems
7/18/2017	PPT Slide	3-D Models Help Scientists Gauge Flood Impact
8/8/2017	PPT Slide	ALCC Program Awards 1 Billion Hours on OLCF Resources
8/8/2017	PPT Slide	Summer Internships Offer Students Hands-On Experience, Mentorship
8/8/2017	PPT Slide	OLCF Staff Member Shares Work on Improving Simulation Efficiency
8/8/2017	PPT Slide	OLCF Hosts CUDA Workshop for GPU Programming
8/8/2017	PPT Slide	Galactic Winds Push Researchers to Probe Galaxies at Unprecedented Scale
8/25/2017	PPT Slide	Titan Helps Researchers Suck Mystery Out of Cell's 'Vacuum Cleaners'
9/19/2017	PPT Slide	OLCF Provides Research Opportunities, HPC Experience for Postdocs
9/19/2017	PPT Slide	CADES Supplies Resources, Expertise to World's Most Powerful Accelerator
9/19/2017	PPT Slide	Sodium Shakedown in Dopamine Research
10/17/2017	PPT Slide	Decades-Long Physics Mystery Elucidated with Titan
10/17/2017	PPT Slide	OLCF Helps GE Deliver Next-Generation Gas Turbines
10/19/2017	PPT Slide	Faces of Summit: Making Spaces
10/31/2017	PPT Slide	OLCF Staff Honored for Excellence in Science, Technology, and Support
10/31/2017	PPT Slide	OLCF Postdoc Fuses the Gap between Experiment and Computation
10/31/2017	PPT Slide	OLCF to Celebrate 25 Years of HPC Leadership
11/10/2017	PPT Slide	Ready for 25 More Years
11/13/2017	PPT Slide	INCITE Grants of 5.95 Billion Hours Awarded to 55 Computational Research Projects
11/28/2017	PPT Slide	Conferences Promote Greater Diversity in HPC
11/28/2017	PPT Slide	Teams Gear up for Summit at Fourth Annual GPU Hackathon
11/28/2017	PPT Slide	Faces of Summit: Modeling Safety
11/28/2017	PPT Slide	Scaling Deep Learning for Science
12/20/2017	PPT Slide	New Boeing Method Accelerates Turbulence Modeling Uncertainty Analysis
12/20/2017	PPT Slide	OLCF Team 'Connects' with the High-Performance Computing Community at SC17

Date	Type of Product	Title
12/20/2017	PPT Slide	As Exascale Frontier Opens, Science Application Developers Share Pioneering Strategies
12/20/2017	PPT Slide	Faces of Summit: Bursting with Ingenuity
3/1/2017	PPT Slides	OA PowerPoint Slides
4/15/2017	PPT Slides	Media briefing package: Washington Post
5/8/2017	PPT Slides	Reserse Site Visit Slides
5/10/2017	PPT Slides	Media briefing package: *New York Times*
8/2/2017	PPT Slides	Congressional Staffers Briefing Package
11/6/2017	PPT Slides	Summit and Summitdev Overview Slides
12/22/2017	PPT Slides	Media briefing package: Bloomberg
9/1/2017	Publication	OLCF Annual Report
9/25/2017	Publication	OLCF 25th Anniversary Book
11/7/2017	Publication	OLCF 25th Anniversary Event Program
1/4/2017	Quad Chart	UTRC Researchers Take Flight with Graphics Processing Units for Large Eddy Simulations
1/31/2017	Quad Chart	Researchers Flip Script for Li-Ion Electrolytes to Simulate Better Batteries
1/31/2017	Quad Chart	The Shape of Melting in Two Dimensions
2/2/2017	Quad Chart	Supercomputing, Experiment Combine for First Look at Magnetism of Real Nanoparticle
2/28/2017	Quad Chart	Modeling Autoignited Combustion to Drive Engines Forward
3/28/2017	Quad Chart	Researchers Shoot for Success with Simulations of Laser Pulse–Material Interactions
4/18/2017	Quad Chart	A Real CAM-Do Attitude
4/18/2017	Quad Chart	Predictive Power
5/9/2017	Quad Chart	Supercomputing Mimics Berkelium Experiments to Validate New Find
5/9/2017	Quad Chart	Assembling Life's Molecular Motor
6/5/2017	Quad Chart	Researchers Seek Sigma Meson on the Path to Heavier Hadrons
6/5/2017	Quad Chart	Putting the Pedal to the Metal in the Hunt for Alloys
7/18/2017	Quad Chart	Titan Simulations Show Importance of Close 2-Way Coupling between Human and Earth Systems
7/18/2017	Quad Chart	3-D Models Help Scientists Gauge Flood Impact
8/8/2017	Quad Chart	Galactic Winds Push Researchers to Probe Galaxies at Unprecedented Scale
8/25/2017	Quad Chart	Titan Helps Researchers Suck Mystery Out of Cell's 'Vacuum Cleaners'
9/19/2017	Quad Chart	Sodium Shakedown in Dopamine Research
10/17/2017	Quad Chart	Decades-Long Physics Mystery Elucidated with Titan
11/28/2017	Quad Chart	Scaling Deep Learning for Science
12/20/2017	Quad Chart	New Boeing Method Accelerates Turbulence Modeling Uncertainty Analysis
3/1/2017	Report	Contributions for Operational Assessement
8/14/2017	Report	2015-2016 ALCC Closeout Reports
12/31/2017	Report	2016 INCITE Closeout Reports
12/31/2017	Report	2016-2017 ALCC Closeout Reports
11/7/2017	Video	OLCF 25th Anniversary Long
11/7/2017	Video	OLCF 25th Anniversary Short
11/7/2017	Web Article	OLCF 25th Anniversary Timeline

Date	Type of Product	Title
11/1/2017	Web Article	SC16 Event Details
11/1/2017	Website	SC17
11/1/2017	Website	INCITE 2018

APPENDIX D. BUSINESS RESULTS FORMULAS

2017 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 by more than 4 hours will be counted as an adjacent unscheduled outage, as an unscheduled availability, and as an additional interrupt.

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

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 of the full 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 of the full 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 2017

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Vahid Ranjbar	Brookhaven National Laboratory	1,500,000	1,621,666	Spin Tracking for RHIC
Remi Lehe	Lawrence Berkeley National Laboratory	500,000	12,173	Innovative Spectral Particle-In-Cell Algorithms for Laser-Wakefield Acceleration
Ji Qiang	Lawrence Berkeley National Laboratory	30,000	81,219	Particle Accelerator Modeling on Exascale Computers
Shailendra Kaushik	General Motors	4,000,000	960,746	Multi-Disciplinary Optimization of Automobile Aerodynamic Performance, Thermal Performance, and Aero-Acoustics
Dan Williams	ANSYS Inc.	4,000,000	0	ANSYS AIM Fluid Solver Scaling
Jaime Peraire	Massachusetts Institute of Technology	4,000,000	0	Scalable Discontinuous Galerkin Solvers for the Implicit Large-Eddy Simulation of Transitional Flows
Peter Edward Vincent	Imperial College London	37,500,000	0	Development and Benchmarking of an In-Situ Visualization Pipeline for Next Generation CFD Tools
Mujeeb Malik	NASA Langley Research Center	3,000,000	0	Wall-Resolved Simulations of Complex Separated Aerodynamic Flows
Nathan See	SmartTruck	-	0	Investigating matching Unsteady CFD to Real World testing of Class 8 trucks
Paul Robert Shapiro	University of Texas at Austin	17,000,000	141,035,662	Simulating Reionization of the Local Universe: Witnessing our own Cosmic Dawn
Eric J Lentz	University of Tennessee	35,000,000	0	Examination of Spatial and Network Resolution Core-Collapse Supernova Simulations
Brant Robertson	University of California - Santa Cruz	3,000,000	102,376	Extending the Physics of the GPU-Enabled CHOLLA Code to the Power of Titan
Philip Fajardo Hopkins	California Institute of Technology	4,000,000	1,568,273	Galaxies on FIRE: Shedding Light on Dark Matter
Charles Horowitz	Indiana University	1,000,000	1,466,047	Crystallization of Rapid Proton Capture Ash
Gonzalo Merino	University of Wisconsin	1,000,000	350	Management and Operation of the IceCube Neutrino Observatory (ICNO)

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Rahul Kumar	Princeton University	2,500,000	1,971,400	Kinetic Turbulence in Magnetized Plasmas
Alexander Chekhovskoy	Northwestern University	-	0	Parallel Scaling of GPU-accelerated Simulations of Compact Binary Merger Remnant Disks, Outflows, and Nucleosynthesis
Rod Ray Linn	Los Alamos National Laboratory	2,000,000	0	Scaling Wildland and Urban Fire Modeling Capabilities using HIGRAD/FIRETEC
Jesse Carman	NOAA-Geophysical Fluid Dynamics Laboratory (GFDL)	2,000,000	117,668	Air-Ocean-Land-Ice Global Coupled Prediction on Emerging Computational Architectures: A Framework for ESPC Coupled Models
Inanc Senocak	University of Pittsburgh	1,000,000	0	GEM3D: Open-Source Cartesian Adaptive Complex Terrain Atmospheric Flow Solver for GPU Clusters
Gur Pines	University of Colorado	80,000	12	Overlapping Genes to Limit Mutational Escape
Daniel Jacobson	Oak Ridge National Laboratory (ORNL)	20,000,000	0	Center For Bioenergy Innovation
Daniel Jacobson	Oak Ridge National Laboratory (ORNL)	4,000,000	29,810,476	Scaling Up Parallelized Ortholog Detection Algorithms for Comparative Genomics of Bacterial Genomes
Hwee Kuan Lee	Oak Ridge National Laboratory (ORNL)	100,000	108,638	Development of Probabilistic Neural Network for Computer Aided Diagnosis of Cancer Specimens
Volkhard Helms	Universität des Saarlandes	2,000,000	0	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	4,000,000	3,059,009	Molecular Dynamics Simulations of HIV-1 Envelope Glycoprotein Trimers
Lawrence Shapiro	Columbia University in the City of New York	4,000,000	10,285,999	Mechanisms of Functioning of Somatic Hypermutation on Antibody Affinity Maturation
Manjunath Gorentla Venkata	Oak Ridge National Laboratory (ORNL)	200,000	31,943	Understanding and Constructing Fine-scale Genetic Map of Rhesus Macaque
Frank Noe	Freie Universtat Berlin	12,000,000	1,623,738	Adaptive Molecular Simulation of the Immunological Synapse
Chongle Pan	Oak Ridge National Laboratory (ORNL)	5,000,000	244	Large-Scale Integrated Omics Analyses of Complex Microbial Communities in Plant-Microbe Interfaces and Tropical Rainforests
Harel Weinstein	Cornell University	3,000,000	12,356,386	Mechanisms and Allostery in Energy-driven Molecular Machines in Membranes

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Edward Lyman	University of Delaware	5,000,000	0	Subdiffusive Signatures of Nanoscale Lipid Bilayer Domains
Jens Meiler	Vanderbilt University	5,000,000	0	Structure Determination of Cellulose Synthase Using Neutron Scattering and High Performance Computing
David N. Beratan	Duke University	1,000,000	0	Computer Assisted Diversity Oriented Molecular Library Design
Pak-Lee Chau	Institut Pasteur	500,000	442,338	Simulation of the Nicotinic Acetylcholine Receptor
Joseph Curtis	National Institute of Standards and Technology (NIST)	500,000	0	Atomistic Modeling of Small-Angle Scattering Experimental Data on HPC Resources via SASSIE-web
Mitchel John Doktycz	Oak Ridge National Laboratory (ORNL)	5,000,000	0	Multiscale Simulations in Adaptive Biosystems Imaging
Remco Hartkamp	Vanderbilt University	500,000	60	Permeability of Multicomponent Gel-phase Bilayers
Miguel Fuentes-Cabrera	Oak Ridge National Laboratory (ORNL)	3,000,000	0	Theoretical Study of the 2D Self-Assembly of the Carboxysome
Yevgeny Moskovitz	Ben-Gurion University of the Negev	260,000	0	Validation of the Multi-Site Expansion Hypothesis of Neurological Impairments in N-methyl-D-aspartate Receptor
Jan-Michael Carrillo	Oak Ridge National Laboratory (ORNL)	2,000,000	1,195,035	Large Scale Coarse-Grained Molecular Dynamics Simulations of Lipid Bilayers
Ron Dror	Stanford University	1,500,000	239,134	Enabling Rational Design of Functionally Selective GPCR Drugs
Nikolay Smolin	Loyola University Chicago	700,000	0	Molecular dynamics simualtion of SERCA, SpeG and TRIM5a.
Aleksei Aksimentiev	University of Illinois at Urbana-Champaign	100,000	5,048	Benchmarking the Nuclear Pore Complex on NAMD
Julie C Mitchell	Oak Ridge National Laboratory (ORNL)	50,000	1,018	Co-evolutionary Networks: From Genome to 3D Proteome
Ivaylo Ivanov	Georgia State University	-	0	Integrative Modeling of Transcription Initiation Assemblies in Gene Regulation
Charlotte Barbier	Oak Ridge National Laboratory (ORNL)	2,000,000	2,891,594	Cavitation Bubbles near Solid Surfaces
Robert Vance Wilson	Oak Ridge National Laboratory (ORNL)	500,000	91,427	Biomimetic Vector Sensor Towed Arrays (BVSTA)

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Sumanta Acharya	Illinois Institute of Technology	500,000	103,230	Indirect Dry Cooling for Power Plants Using Encapsulated Phase Change Materials
Yue Ling	Baylor University	1,000,000	500,380	Direct Numerical Simulation of Spray Formation in a Two-phase Mixing Layer Between Two Parallel Gas and Liquid Streams
Jesse Ault	Oak Ridge National Laboratory (ORNL)	2,000,000	5	Using Parallel Computation to Improve Blood Cell Simulations in Cardiovascular Flows
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 (ORNL)	1,000,000	750,203	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
Nike Dattani	University of Oxford	1,000,000	242	Competing with Quantum Computers for Finding the Ground State Energy of the Iron-molybdenum Complex
John Anthony Parkhill	University of Notre Dame	2,000,000	65	Accelerating Computational Chemistry a Million Times Over by Feeding Data-hungry Algorithms
Remco Havenith	University of Groningen	4,000,000	1,097,864	GRONOR
Srikanth Allu	Oak Ridge National Laboratory (ORNL)	1,000,000	0	Consortium of Advanced Battery Simulation
Kwangho Nam	University of Texas at Arlington	500,000	1,824,733	Multiscale Modeling of Kinase Conformational Change
Deborah Penchoff	University of Tennessee	1,500,000	147,939	Selective Binding of Rare Earth Elements and Actinides for Nuclear Forensics Applications
Ada Anna Sedova	Oak Ridge National Laboratory (ORNL)	2,500,000	0	Creation of a Tested HPC-based Workflow to Calculate VISION INS Spectra From DFT-based Frequency Calculations
Warren Washington	University Corporation of Atmospheric Research (UCAR)	-	0	CESM Century-Scale Climate Experiments with High-Resolution Atmosphere
Katherine J Evans	Oak Ridge National Laboratory (ORNL)	9,000,000	10,787,461	The Computational Climate Science Integrated Allocation
Francois William Primeau	University of California - Irvine	150,000	0	Implementation and Testing of a Tracer Transport Matrix for the MPAS-O Ocean Model

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
William I Gustafson Jr	Pacific Northwest National Laboratory	-	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 (ORNL)	100,000	0	Workflow Automation and Advanced Analytics for Climate Simulations
Kiran Alapaty	United States Environmental Protection Agency	2,500,000	494,202	Next Generation Climate Modeling using the CESM and MPAS Modeling Systems
Cecilia Bitz	University of Washington	-	0	Climate Model Response to Ozone: HPSS Data Recovery
Dali Wang	Oak Ridge National Laboratory (ORNL)	500,000	5,138	Intelligent Knowledge Management Framework for Earth Science Data Initiative
Valentine Anatharaj	Oak Ridge National Laboratory (ORNL)	10,000	0	Provisioning of Climate Data
Venkat E Tangirala	GE Global Research	5,000,000	0	Unsteady Combustion Processes in a Gas Turbine Combustor
Eric Brown-Dymkoski	Space Exploration Technologies Corp	5,000,000	4,763,599	Scaling and Validation of Adaptive-Grid Turbulent Mixing Simulations
Josette Bellan	California Institute of Technology	5,000,000	0	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	3,241,683	Development of a Physics-Based Combustion Model for Engine Knock Prediction
Scott Michael Schnobrich	3M Company	-	0	Cobra
Clayton Naber	Pinnacle Engines, Inc.	4,000,000	5,047,205	Development of Opposed-Piston Variable Compression Ratio Engine for Automotive Applications using High-Performance Computational Combustion Methodologies
Karl Virgil Meredith	FM Global	-	0	Fire Suppression Modeling
Gary Cai	Fiat Chrysler Automobiles	-	0	A Fundamental Study of the Factors Affecting Adverse Autoignition (knock) in Internal Combustion Engines
Ramanan Sankaran	Oak Ridge National Laboratory (ORNL)	-	0	Level-set Modeling Simulations of Chemical Vapor Infiltration for Ceramic Matrix Composites Manufacturing
Seung Hyun Kim	Ohio State University Research Foundation	2,000,000	665,159	Development of a Physics-Based Combustion Model for Engine Knock Prediction: Phase II

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Dario Alfe	University College London	5,000,000	11,596,323	Non-Covalent Bonding in Complex Molecular Systems with Quantum Monte Carlo
Panchapakesan Ganesh	Oak Ridge National Laboratory (ORNL)	3,000,000	1,723,591	Data Driven Discovery by Design of Energy Materials
Andreas Glatz	Argonne National Laboratory	3,000,000	2,311,312	OSCon
Garnet Kin-Lic Chan	Princeton University	4,000,000	0	High Performance Quantum Chemistry and Tensor Networks
Alexander Alexeev	Georgia Institute of Technology	1,000,000	543,736	Mesoscale Modeling of Suspensions of Synthetic Microswimmers
Gabriel Kotliar	Rutgers, the State University of New Jersey	2,000,000	2,269,017	Application of GPU-accelerated Quantum Monte-Carlo Impurity Solver to Plutonium Compounds
Anton Kozhevnikov	Swiss National Supercomputing Centre (CSCS)	-	0	ADAC-CSCS
Jamison Daniel	Oak Ridge National Laboratory (ORNL)	1,000,000	0	Cluster-based Visualization of Tera- and Peta-Scale Datasets
Neena Imam	Oak Ridge National Laboratory (ORNL)	4,000,000	15,271,750	Durmstrang
David Ronald Pugmire	Oak Ridge National Laboratory (ORNL)	5,000,000	792,875	SDAV
Sergey Panitkin	Brookhaven National Laboratory	-	172,960,324	Next Generation Workload Management System
Fernanda Schafer Foertter	Oak Ridge National Laboratory (ORNL)	1,250,000	0	Developing Scalable Heterogeneous Computing Training Code Examples
Bronson Messer	Oak Ridge National Laboratory (ORNL)	1,000,000	0	CORAL Benchmarking
Scott A Klasky	Oak Ridge National Laboratory (ORNL)	25,000,000	5,076,874	ADIOS
Matthew Belhorn	Oak Ridge National Laboratory (ORNL)	-	0	Data Transfer Working Group
Robert M Patton	Oak Ridge National Laboratory (ORNL)	3,000,000	581,224	Scalable Deep Learning Systems for Exascale Data Analysis
Joshua Ryan New	Oak Ridge National Laboratory (ORNL)	2,000,000	1,220,731	Big Data Mining for Building Analytics

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Mark Christopher Miller	Lawrence Livermore National Laboratory	100,000	0	FASTMath Installation and Portable Performance Testing
Lubomir Riha	IT4 Innovations National Supercomputing Center	2,700,000	1,659,546	ESPRESO - ExaScale PaRallel FETI SOlver
Manuel Arenaz	Appentra Solutions S. L.	50,000	0	Porting Parallware Tool to Large HPC Installations Including Titan
Matthew Dearing Wolf	Oak Ridge National Laboratory (ORNL)	2,000,000	11,690	In Situ Analytics Infrastructures
Eric T. Phipps	Sandia National Laboratories	3,000,000	418,409	EQUINOX Embedded Uncertainty Quantification
Catherine Schuman	Oak Ridge National Laboratory (ORNL)	2,500,000	1,682,397	Scalable Evolutionary Optimization for Designing Networks
David Gutzwiller	Numeca USA, Inc.	2,000,000	447,614	Preparation of the NUMECA FINE/Open CFD Solver for Leadership Computing: Parallel I/O, Profiling, and Acceleration
Vivek Sarkar	Rice University	1,000,000	83,446	Unified Portable Programming System for Integrating Task Parallelism, Accelerator Parallelism, and Message Passing on Exascale Systems
Alan Gray	University of Edinburgh	3,000,000	0	Performance Portability for Large-scale Scientific Applications
Ramakrishnan Kannan	Oak Ridge National Laboratory (ORNL)	1,000,000	448,752	Mini-Apps for Big Data
Chad Allen Steed	Oak Ridge National Laboratory (ORNL)	1,000,000	0	In Situ Visual Analytics for Extreme Scale Simulation Science
Jeffrey Young	Georgia Institute of Technology	750,000	0	Evaluation and Porting of ORNL Mini-apps to Future Directive-Based Languages and Runtimes
Cory Hauck	Oak Ridge National Laboratory (ORNL)	800,000	434	Diagnostics for Data Compression at Scale
Jeff Clune	University of Wyoming	120,000	443,809	Improving Artificial Intelligence by Studying and Harnessing the Evolution of Structurally Organized Neural Networks
Terry Ray Jones	Oak Ridge National Laboratory (ORNL)	3,000,000	1,039,199	UNITY: Unified Memory & Storage Space
Alex Aiken	Stanford University	5,000,000	0	Scaling Legion S3D

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Manjunath Gorentla Venkata	Oak Ridge National Laboratory (ORNL)	1,000,000	1,212,775	SharP: Shared Datastructure Programming Paradigm for Extreme Scale Systems
Abhinav Vishnu	Pacific Northwest National Laboratory	3,500,000	4,524,128	Machine Learning on Extreme Scale GPU systems
Michael Joseph Brim	Oak Ridge National Laboratory (ORNL)	5,000,000	602,599	ExCDS (Exascale Computing Data Service)
Edmon Begoli	Oak Ridge National Laboratory (ORNL)	-	0	Advanced Data Architectures
Sherman Kisner	High Performance Imaging	240,000	0	Scalable parallelization for high performance tomography
Audris Mockus	University of Tennessee	1,500,000	65,389	Fingerprinting Online Users Using Doc2Vec Model and Bayesian Networks
Shantenu Jha	Rutgers, the State University of New Jersey	5,000,000	5,621,385	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	5,000,000	5,355,122	Understanding File System Performance in Production Supercomputers
Misun Min	Argonne National Laboratory	1,500,000	20,002	CEED
Jeremy Travis Johnston	Oak Ridge National Laboratory (ORNL)	3,500,000	2,996,712	Surrogate Based Modeling for Deep Learning Hyper-parameter Optimization
Arvind Ramanathan	Oak Ridge National Laboratory (ORNL)	1,500,000	2,820,847	ECP Cancer Distributed Learning Environment
Peter Zaspel	University of Basel	3,000,000	598,524	Extreme-scale Many-core Solvers for Data Analysis and Uncertainty Quantification
Jakub Kurzak	University of Tennessee	500,000	0	SLATE
Catherine Schuman	Oak Ridge National Laboratory (ORNL)	3,000,000	4,955,122	Scalable Neuromorphic Simulators: High and Low Level
Ian Foster	Argonne National Laboratory	300,000	946,863	CODAR
Sunita Chandrasekaran	University of Delaware	20,000	497	Migrating Legacy Code to Novel Directive-based Programming Models
Mark Berrill	Oak Ridge National Laboratory (ORNL)	-	0	Summit Acceptance Apps

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Robert Glenn Brook	University of Tennessee	-	0	JICS/AACE Emerging Architectures Activities
Wayne Joubert	Oak Ridge National Laboratory (ORNL)	-	0	ALExa ECP
George Martian Bosilca	University of Tennessee	500,000	206,388	Open MPI x
Hartwig Anzt	University of Tennessee	-	0	ECP-PEEKS
Allan Grosvenor	Microsurgeonbot	1,500,000	794	Intelligent Middleware, Making CFD Accessible Through Leadership-Scale Deep Learning Research
Galen Shipman	Los Alamos National Laboratory	-	0	Legion ECP
Nicholas J Zabaras	University of Notre Dame	1,000,000	919,135	Deep Gaussian Processes for Predictive Multiscale Modeling
Terry Ray Jones	Oak Ridge National Laboratory (ORNL)	2	0	ECP: Simplified Interface for Complex Memories
Judith C. Hill	Oak Ridge National Laboratory (ORNL)	500,000	29	Computational Science Graduate Fellowship Program
Lin-Wang Wang	Lawrence Berkeley National Laboratory	4,000,000	36,563	HPC4mfg: Making Semiconductor Devices Cool Through HPC Ab-initio Simulations
Lonnie D Crosby	University of Tennessee	3,484,800	431,738	UTK Benchmarking and Performance Projections
Alexander Joseph McCaskey	Oak Ridge National Laboratory (ORNL)	300,000	0	xacc-tnqvm-ibm-quantum-computing
Tyler Simon	National Security Agency (NSA)	-	0	Scalable Tensor Decomposition
William Scott Daughton	Los Alamos National Laboratory	-	0	Linkages Between Turbulence and Reconnection in Kinetic Plasmas
William M Tang	Princeton University	1,000,000	108,065	Extreme Scale PIC Research on Advanced Architectures
Judith C. Hill	Oak Ridge National Laboratory (ORNL)	1,500,000	1,782,342	Workflow Optimization and Processing of Complex Datasets for Off-site Fusion Energy Research
John Gary Shaw	University of Rochester	1,000,000	3,207	Laser Plasma Simulation Environment
Daniel Patrick Fulton	Tri Alpha Energy Inc.	6,000,000	4,211,644	Fusion Reactivity Enhancement: Numerical Beam-Driven FRC
William M. Tang	Princeton University	6,000,000	5,271,555	Big Data Machine Learning for Fusion Energy Applications
Igor Igumenshchev	University of Rochester	2,000,000	531,367	Development of Simulation Tools for Laser Fusion Experiments
Henri Francois Calandra	Total E&P	7,500,000	0	Advanced Computing for Geoscience Applications

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Dilip Reddy Patlolla	Oak Ridge National Laboratory (ORNL)	2,000,000	576,368	GPU Accelerated Settlement Detection
Terry Ray Jones	Oak Ridge National Laboratory (ORNL)	5,000,000	342,888	Foresight
George Serban Constantinescu	University of Iowa	2,000,000	495,745	Critical Assessment of a Fully 3-D Nonhydrostatic Flow Solver to Simulate Flood Propagation in Natural Environments
Kyle Withers	University of California - San Diego	1,000,000	255,391	Improving numerical simulations of earthquake ground motions using high-resolution 3-D seismic velocity and improved rupture models
Jibonananda Sanyal	Oak Ridge National Laboratory (ORNL)	500,000	0	GPU-Accelerated Settlement Detection
Dalton Lunga	Oak Ridge National Laboratory (ORNL)	200,000	0	FMOW
Tsuyoshi Ichimura	University of Tokyo	-	0	Fast & Scalable Finite Element Method with Low-order Unstructured Elements for Earthquake Simulation on Summit
Katrin Heitmann	Argonne National Laboratory	4,000,000	30,410,125	Sky Surveys
Gabriel Nathan Perdue	Fermi National Accelerator Laboratory	1,000,000	905,582	MACHINE Learning for MINERvA
Jason Newby	Oak Ridge National Laboratory (ORNL)	-	0	The COHERENT Experiment at the SNS
Jean-Roch Vlimant	California Institute of Technology	2,000,000	335,120	HEPDeepLearning
Boram Yoon	Los Alamos National Laboratory	2,000,000	14,748	Artificial Intelligence for Collider Physics
Malachi Schram	Pacific Northwest National Laboratory	1,000,000	0	Belle II
Andre Walker-Loud	Lawrence Berkeley National Laboratory	2,500,000	16,842,364	Scaling Lattice QCD Calculations for Leadership Computing Facilities
Richard Brower	Boston University	-	0	ECP Lattice Field Theory Solvers
Arvind Ramanathan	Oak Ridge National Laboratory (ORNL)	5,000,000	16,953,437	Integrating Neutron Scattering Experiments with Atomistic Simulations to Characterize Biophysical Mechanisms of Intrinsically Disordered Proteins
Liqun Zhang	Tennessee Technological University	1,500,000	13,719,010	Molecular Dynamics Simulations to Investigate the Structure, Dynamics and Functional Properties of Human Beta Defensin Type 3

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Amanda Randles	Duke University	2,000,000	2,106,071	Massively Parallel Hemodynamic Simulation
Vipin Sachdeva	Silicon Therapeutics	1,500,000	539,560	Large-scale Validation of Drug Discovery Pipeline Protocols
Byung Hoon Park	Oak Ridge National Laboratory (ORNL)	5,000,000	4,229,518	Accelerating Materials Modeling Loop of Leadership Computing and Spallation Neutron Source
Jerzy Jerry Bernholc	North Carolina State University	3,148,800	277,740	High Performance Simulations of Electron Spin Distribution and Dynamics in Low-Dimensional Materials
Dongwon Shin	Oak Ridge National Laboratory (ORNL)	8,700,000	3,918,454	High Performance Cast Aluminum Alloys for Next Generation Passenger Vehicle Engines
Trung Dac Nguyen	Vietnam Academy of Science and Technology	1,000,000	21,511	Implementing and Optimizing GPU-accelerated Molecular Dynamics Models
Venkat Padmanabhan	Indian Institute of Technology Kharagpur	500,000	367,852	Active Layer Morphology for Improving Efficiency of Organic Photovoltaics
Shanmugavelayut Kamakshi Sundaram	Alfred University	1,000,000	1,265,569	Structural and Mechanical Property Optimization of Photomultiplier Tube (PMT) Glasses for Neutrino Detection using MD Simulations
Yangyang Wang	Oak Ridge National Laboratory (ORNL)	3,000,000	11,173,264	Elucidating the Influence of Reversible Non-Covalent Interactions on Dynamic Properties for Rational Design of Soft Materials
Rick Archibald	Oak Ridge National Laboratory (ORNL)	2,000,000	82,636	CADES/OLCF Computational Workflows for Materials Science
Hendrik Heinz	Regents of theUniversity of Colorado	4,000,000	4,222,847	Understanding and Designing Alloy Nanocatalysts in Atomic Resolution
Saurabh Ghosh	Vanderbilt University	2,000,000	590,013	Engineering Multifunctionality in Oxide Heterostructures
Jeffrey C Grossman	Robert Bosch LLC	6,000,000	853,577	Accurate Simulation of Li-Rich Layered Oxide Materials
Srikanth Allu	Oak Ridge National Laboratory (ORNL)	4,000,000	0	A Large Scale Cycling Studies for Lithium-Ion Battery Packs
Michael J. Demkowicz	Texas A&M University - College Station	-	0	Phase-field Modeling of He Growth, Coalescence, and Stability at Cu-V Interfaces
Marcus Müller	Georg-August-Universität Göttingen	750,000	1,733,354	SOMA: Soft Coarse Grained Monte-carlo Acceleration
Peter Thomas Cummings	Vanderbilt University	4,000,000	3,832,615	Materials Genome Screening of Soft Materials

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Steven Hartman	Oak Ridge National Laboratory (ORNL)	-	0	ORNL Neutron Sciences Experiment Data
Adolfo Eguiluz	University of Tennessee	2,400,000	1,884,401	GW Many-body Solver for Ab-initio Multi-orbital Hamiltonians for Correlated Materials
Weiwei Sun	Oak Ridge National Laboratory (ORNL)	200,000	0	Electronic Structure of Defects in Two-dimensional Materials
Jeetain Mittal	Lehigh University	1,000,000	842	Biomolecular Assembly Processes in the Design of Novel Functional Materials
Emmanuel Vallejo	Universidad Autónoma del Estado de Hidalgo	100,000	1,713,522	Theoretical and Experimental Study on Polymeric Molecular Self-assembly on Metallic Substrates
Jacek Jakowski	Oak Ridge National Laboratory (ORNL)	3,000,000	1	Electronic Structure Modeling of Impurity States
Allison Lee Dzubak	Oak Ridge National Laboratory (ORNL)	1,700,000	502,787	QMC modeling of MOFs
Bobby Sumpter	Oak Ridge National Laboratory (ORNL)	10,500,000	5,430,854	Center for Nanophase Materials Sciences (CNMS)
Lin Lin	Lawrence Berkeley National Laboratory	1,000,000	153,010	ESMATH
Bo Kong	Ames Laboratory	-	0	Multi-scale Simulations of Gas Atomization Process for Metal Alloy Powder Production
Axiel Birenbaum	Oak Ridge National Laboratory (ORNL)	3,000,000	215,361	Reduced Symmetry in Oxide Heterostructures: Defects and Interfaces
Haixuan Xu	University of Tennessee	2,000,000	876,600	Dislocation Interaction with Radiation-Induced Defect in Structrual Materials
Mordechai C Kornbluth	Robert Bosch LLC	500,000	9,845	Conductivity Mechanisms in Glasses
Miguel Fuentes-Cabrera	Oak Ridge National Laboratory (ORNL)	1,000,000	66,716	Molecular Dynamics Simulations of Liquid Metal Assembly at the Nanoscale
Trung Dac Nguyen	Vietnam Academy of Science and Technology	1,000,000	35	Development of Models and Accelerated Sampling Techniques for Massively Parallel Molecular Dynamics Codes
Nouamane Laanait	Oak Ridge National Laboratory (ORNL)	-	0	Deep Learning Applications in Microscopy at Mixed-Precision

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Georgia Tourassi	Oak Ridge National Laboratory (ORNL)	4,000,000	7,031,165	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	3,000,000	4,154,819	Gas-Liquid-Solid Interfaces for Energy Applications
Bamin Khomami	University of Tennessee	1,000,000	2,840,329	Multiscale Modeling
Massoud Kaviany	University of Michigan	1,500,000	5,659,846	Human Thermosensation: Response of TRPV1 Domains to Temperature
Brendan Michael McLaughlin	Queen's University Belfast	2,000,000	0	T-Iron
Phillip C. Stancil	University of Georgia	1,200,000	809	Quantum Dynamics of Molecule-Molecule Collisions in Full-Dimensionality
Rajan Pandey	GlobalFoundries	3,000,000	5	Ab-Initio Materials Modeling for Next Generation CMOS Devices
Rachel Slaybaugh	University of California - Berkeley	4,000,000	27,622	Ex-core Dosimetry Calculations for a Pressurized Water Reactor and Comparison to Operational Data
Sacit Cetiner	Oak Ridge National Laboratory (ORNL)	3,000,000	545	Nuclear-Renewable Hybrid Energy Systems
Robert Lindsay Varner Jr	Oak Ridge National Laboratory (ORNL)	-	0	Majorama Demonstrator Secondary Data Archive
Rajan Gupta	Los Alamos National Laboratory	3,000,000	28,943,717	Probing Novel Physics using Nucleon Matrix Elements
Leah Broussard	Oak Ridge National Laboratory (ORNL)	100,000	20,353	nEDM at SNS
Dean Lee	North Carolina State University	600,000	85,790	Nuclear Lattice Simulations with GPU Acceleration on Titan
Kaushik Datta	University of Southern California	-	0	PRISMA
Hanchuan Peng	Allen Institute for Brain Science	2,000,000	863,633	BigNeuron
Fabien Jonathan Delalondre	Ecole Polythechnique Federale de Lausanne (EPFL)	3,000,000	0	Brain Simulation Modeling
Misun Min	Argonne National Laboratory	2,500,000	71	Nek-HOM (Codes for High Order Methods)
Alex Travesset	Ames Laboratory	45,000	42,125	Crystal Structure Prediction in Nanoparticle Superlattices
Brian Edwards	University of Tennessee	1,000,000	0	Nanoscale Ion Transport Through Carbon Nanotubes

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Jerzy Bernholc	North Carolina State University	2,000,000	6,245	Theoretical Investigations of Nanostructures
Peter Zaspel	University of Basel	2,000,000	0	Scalable Numerical Algebra for UQ and CFD
Helmut Katzgraber	Texas A&M University - Commerce	300,000	0	Benchmarking and Development of Quantum Annealing Architectures
Balakrishnan Naduvalath	University of Nevada - Las Vegas	500,000	1,037	Quantum Dynamics of Ultracold Chemical Reactions
Abhishek Kumar	Oak Ridge National Laboratory (ORNL)	-	0	Controlling the Movement of Leidenfrost Droplet: a Molecular Dynamics Study
Shiwei Zhang	College of William & Mary	4,000,000	2,677,481	Quantum Many-Body Computations of Strongly Correlated Systems
Jason Lee Pries	Oak Ridge National Laboratory (ORNL)	1,500,000	0	High Fidelity Electric Motor/Generator Modeling and Optimization
Tameem Albash	University of Southern California	8,000,000	7,518,281	True Scaling of Time to Solution on a Quantum Annealer
Gregory Gershom Howes	University of Iowa	-	0	Visualization and Analysis of High-Dimensional Kinetic Plasma Physics Data
Carlos Federico Lopez	Vanderbilt University	4,000,000	2,353,785	Reducing Unidentifiability in Cell-Signaling Network Models Through Multidimensional Analysis and Molecular Simulation
George Em Karniadakis	Brown University	2,000,000	6,312,445	Multiscale Simulation of Human Pathologies
Chongle Pan	Oak Ridge National Laboratory (ORNL)	2,000,000	329,966	Proteogenomics Analysis of Tropical Soil Communities Under Long-term Nitrogen and Phosphorus Fertilization
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
Sumanta Acharya	University of Memphis	1,707,110	1,221,477	Thermally Effective and Efficient Cooling Technologies for Advanced Gas Turbine Systems
John Schaefer	Boeing Company	-	0	Uncertainty Quantification and Sensitivity Analysis of Turbulence Model Coefficients for the Common Research Model

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Said Elghobashi	University of California - Irvine	8,000,000	31,465	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
Pui Kuen Yeung	Georgia Institute of Technology	5,000,000	2,609	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	522,766	Demonstration of Noise Suppression by Water Injection in High-Speed Flows with Direct Numerical Simulation
Joseph D Smith	Missouri University of Science and Technology	1,000,000	870	Numerical Investigation of Flow Instabilities Inside a Westinghouse SMR
Rachelle Lea Speth	Boeing Company	-	0	Noise Control Concepts for the Main Landing Gear of a Commercial Transport Aircraft
Chad Michael Winkler	Boeing Company	3,000,000	703,266	Large Eddy Simulations of Jet Noise and S-duct Diffusers
Raul Payri	Universitat Politècnica de València	1,500,000	630,908	DNS Atomization
Pino Martin	University of Maryland - College Park	2,000,000	245,341	Scaling the CRoCCo Code