ORNL/SPR-2019/1142

US Department of Energy, Office of Science High Performance Computing Facility Operational Assessment 2018 Oak Ridge Leadership Computing Facility

April 2019

Approved for public release.
Distribution is unlimited.

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

ORNL/SPR-2019/1142

Oak Ridge Leadership Computing Facility

US DEPARTMENT OF ENERGY, OFFICE OF SCIENCE HIGH PERFORMANCE COMPUTING FACILITY OPERATIONAL ASSESSMENT 2018 OAK RIDGE LEADERSHIP COMPUTING FACILITY

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

April 2019
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

Executive Summary

HIGH PERFORMANCE COMPUTING FACILITY
2018 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2019

EXECUTIVE SUMMARY

Oak Ridge National Laboratory's (ORNL's) Leadership Computing Facility (OLCF) continues to surpass its operational 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.

While delivering an exceptional operational year on production resources such as Titan, Eos, and Rhea, the OLCF pressed forward in reaching the peak of the petascale with a successful installation and deployment of the Summit supercomputer. Summit debuted as the most capable and efficient system in its class during the November 2018 International Conference for High Performance Computing, Networking, Storage, and Analysis (SC18). Summit represents the culmination of a multi-year effort between the OLCF and vendor partners IBM, Nvidia, and Mellanox, delivering a system that is unmatched in the modeling, simulation, and analysis communities. Application teams worked closely with the OLCF through the Center for Accelerated Application Readiness (CAAR) program for years in preparation for the Summit deployment to hit the ground running with science-ready applications on day one.

Calendar year (CY) 2018 was filled with outstanding operational results and accomplishments: a very high rating from users on overall satisfaction for the fifth year in a row; a tremendous amount of core hours delivered to research projects; the largest annual number of research publications since the deployment of Titan—for the third year in a row; and success in delivering on the allocation split of roughly 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 (see Operational Performance section). 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 2018 OLCF metric targets and the associated results. More information can be found in the Operational Performance section for each OLCF resource.

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*, *Nature Chemistry*, *Advanced Energy Materials, ACS Nano*, *Physical Review X*, *Nature Communications, ACS Catalysis, Proceedings of the National Academy of Sciences, Chemistry Materials, Journal of Physical Chemistry Letters, Green Chemistry, Journal of Materials Chemistry A, Physical Review Letters*, *ACS Applied Materials & Interfaces, Nanoscale, and Bioinformatics*. Crucial domain-specific discoveries facilitated by resources at the OLCF are described in the *High Performance Computing Facility Operational Assessment 2018 Oak Ridge Leadership Computing Facility* (OAR) Strategic Results section. For example, Titan enabled a molecular dynamics simulation that studied a new liquid electrolyte interaction with fluoride which may very well enable fluoride-based batteries (Section 8.2).

Table ES.1.2018 OLCF metric summary

Metric description	CY 2018 target	CY 2018 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%	92%
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%	51.09%
Scheduled Availability, TITAN: Sustain scheduled availability to users, measured as a percentage of maximum possible scheduled.	95%	98.97%
Overall Availability, TITAN: Sustain availability to users, measured as a percentage of maximum possible.	90%	98.51%
Overall Availability, External File System: Sustain availability to users, measured as a percentage of maximum possible.	90%	Atlas 1: 99.95% Atlas 2: 99.29%
Overall Availability, Archive Storage: Sustain availability to users, measured as a percentage of maximum possible.	90%	HPSS: 99.53%

The Titan system has been a highly capable system for the US Department of Energy all while pushing modeling and simulation boundaries since it was made available for production use in May of 2013. Calendar year 2018 marked the 5-year production anniversary (6-year component life) of Titan. It is a significant feat for Titan to still be one of the most powerful computational resources in the world 5 years after deployment considering how rapidly technology has advanced within that same timeframe. 2018 was a big year as the OLCF staff worked diligently to install, configure, accept, and deploy the Summit supercomputer. Summit debuted as the fastest and most powerful system in the world in June of 2018 with a large media event and announcement befitting of Summit's accomplishments. The event was headlined with the US Secretary of Energy Rick Perry, Tennessee Governor Bill Haslam, IBM Chairman, President, and Chief Executive Officer Ginni Rometty, and Nvidia co-founder and CEO Jensen Huang, as well as a host of state and local elected officials and esteemed business leaders. Those in attendance praised the OLCF and ORNL for their achievements in high-performance computing and celebrated the first application to break the exascale barrier, which was demonstrated by a team led by ORNL's Dan Jacobson and Wayne Joubert. The accolades didn't stop there. At the November SC18 conference, Summit was credited as having played a substantial role in the annual Gordon Bell application competition as five of the six science teams used Summit for all or part of their computational and analytic needs. The eventual winners credited Summit in their submissions for providing a system that brings an unapproachable scale to reality that has already and will continue to further their research.

The OLCF now has over 50,000 GPUs and just shy of 300 petabytes of storage available for use between the Titan and Summit systems. This successful deployment and operation of resources 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) and most capable computing center in the world. The OLCF 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.

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 fourth year in a row. Ninety-one percent of respondents (383) rated their overall satisfaction with communications from the OLCF as "satisfied" or "very satisfied." In addition, nearly all the respondents felt that the OLCF kept them well informed about changes (97%), events (98%), and current issues (96%). 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 2018 METRICS

In consultation with the DOE program manager and as proposed in the 2017 OAR, a series of metrics and targets were identified to assess the operational performance of the OLCF in CY 2018. The 2018 metrics, target values, and actual results as of December 31, 2018, are noted throughout this report and are summarized in the Operational Performance section. The OLCF exceeded all agreed-upon metric targets.

ES.3 RESPONSES TO RECOMMENDATIONS FROM THE 2017 OPERATIONAL ASSESSMENT REVIEW

The OLCF did not receive any recommendations from the 2017 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:

"The OLCF appears to be well-run, with a combination of mission focus, attention to detail, and proactive strategic action."

"The OLCF is clearly a well-run and customer-focused facility that enables users to produce outstanding science, and clearly demonstrates why it is a valuable national resource. This report provides a clear and detailed description of OLCF's operations strategies, plans, and goals for the facility. OLCF continues to implement innovations and effective management strategies to improve its facility operations and its users' experiences. We expect that OLCF will continue to enable highly impactful scientific achievements consistent with DOE strategic goals."

ES.4 OPERATIONAL REALIGNMENTS TO BETTER SERVE OLCF STAKEHOLDERS

The OLCF is housed within the National Center for Computational Sciences (NCCS) division at ORNL. In CY 2018 the NCCS underwent an organizational restructuring where the already large and growing HPC Operations group, responsible for the installation, deployment, and maintenance of the compute, data, and infrastructure systems, was split into two essential and more manageable functions. The resulting two groups were the High Performance Computing and Data Operations group and the High Performance Computing Core Operations group. The HPC Core Operations group is a new group that is responsible for all of the foundational infrastructure needed to operate and maintain world-class user facilities such as the OLCF. Essential functions of the group are security, networking, and other shared infrastructure such as configuration management, databases, and authentication systems that exist among many of the NCCS resources.

User Support Results

HIGH PERFORMANCE COMPUTING FACILITY
2018 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2019

1. USER SUPPORT 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 calendar year 2018, the Oak Ridge Leadership Computing Facility (OLCF) supported 1,443 users and 387 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 2018 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 2018 with 92% of tickets being resolved within 3 business days. The OLCF continued to enhance its technical support, collaboration, training, outreach, and communication and engaged in activities that promoted high performance computing (HPC) to the next generation of researchers.

1.1 USER SUPPORT 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 and next generation 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 2018 survey was launched on October 9, 2018 and remained open for participation through December 12, 2018. The survey was sent to 1,230 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, including the Exascale Computing Project (ECP) teams who logged into an OLCF system between January 1, 2018, and September 30, 2018. OLCF staff members were excluded from participation. A total of 422 users completed the survey, for an overall response rate of 34%. The results of the 2018 survey can be found on the OLCF website at https://www.olcf.ornl.gov/olcf-media/center-reports/2018-outreach-survey/.

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

Table 1.1. 2018 user support result metrics summary

Metric Description	2017 target	2017 actual	2018 target	2018 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 2017 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 2018 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%	93%	80%	92%
Average of all user support services ratings	3.5/5.0	4.6/5.0	3.5/5.0	4.5/5.0

1.2 USER SUPPORT METRICS

The OLCF exceeded all user support metrics for 2018. 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 2017 Target	CY 2017 Actual	CY 2018 Target	CY 2018 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	4.5/5.0

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	ECP
Overall satisfaction with the OLCF	4.6/5.0	4.6/5.0	4.6/5.0	4.7/5.0	4.5/5.0
Overall satisfaction with support services	4.5/5.0	4.5/5.0	4.6/5.0	4.5/5.0	4.4/5.0
Overall satisfaction with compute resources	4.6/5.0	4.5/5.0	4.4/5.0	4.6/5.0	4.4/5.0
Overall satisfaction with data resources	4.5/5.0	4.4/5.0	4.3/5.0	4.5/5.0	4.3/5.0

1.2.2 Average Rating across All User Support Questions

The calculated mean of answers to the user support services specific questions on the 2018 survey was 4.5/5.0, indicating that the OLCF exceeded the 2018 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 two open-ended responses to "What are the best qualities of the OLCF?"

The OLCF and its staff are a world-class resource, critical to the advancement of science in the USA. The power of the computing systems and the skills of the staff are the facilities' best qualities.

For decades, OLCF has been successful at fielding state-of-the-art HPC systems and HPC technologies to allow application science teams to productively pursue new discoveries. OLCF staff are friendly and helpful, and they take pride in the successes of the science teams that use their systems to advance knowledge. OLCF has been a key driver in helping the HPC user community make effective use of advanced technologies such as GPU computing, by organizing and supporting GPU Hackathons, by helping support productivity-oriented programming standards like OpenACC, supporting efforts such as OpenPOWER, and by organizing traditional HPC training events for users of OLCF systems such as Titan and Summit. OLCF is much more than just a "compute cycle shop," and their community building efforts are greatly appreciated. These are just a few of the activities that set OLCF apart from many of their peers. It has become particularly evident in the last several years that OLCF has outgrown the stereotypical role of traditional HPC centers and has become a much better partner to the science teams it serves. I would like to see other HPC centers follow the good example set by OLCF's community building efforts, and I urge OLCF to continue with these activities in the future.

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 2017 and 2018 surveys.

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.

ORAU invited all OLCF users, who had logged into an OLCF system between 1/1/2018 through 9/30/2018, to participate in the annual survey. The survey was hosted online beginning on October 9, 2018 and remained open for completion through December 12, 2018. A total of 422 users completed the survey, resulting in a response rate of 34.3%.

ORAU sent out the initial notification on October 9th and additional notices were sent by Ashley Barker, the User Assistance and Outreach Group Leader, the OLCF Executive Board, and the OLCF Leadership Team.

The survey was also 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 2018 of users who have been using OLCF resources more than 2 years (Table 1.4).

Table 1.4. User survey participation

	2017 survey	**2018 survey**
Total number of respondents (Total percentage responding to survey)	448 (40%)	422 (34%)
New users (OLCF user <1 year)	26%	23%
OLCF user 1–2 years	27%	22%
OLCF user >2 years	47%	55%

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 dissatisfied or very 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 shown, the metrics are very comparable across all four allocation programs, and the variations are statistically insignificant.

Table 1.5. Statistical analysis of survey results

Survey area	INCITE			ALCC			DD			ECP		
	Mean	Variance	Std. Deviation	Mean	Variance	Std. Deviation	Mean	Variance	Std. Deviation	Mean	Variance	Std. Deviation
Overall satisfaction with the OLCF	4.6	0.35	0.59	4.6	0.45	0.67	4.7	0.32	0.57	4.6	0.66	0.81
Overall satisfaction with support services	4.5	0.48	0.69	4.5	0.48	0.69	4.5	0.48	0.69	4.5	0.62	0.79
Overall satisfaction with user assistance team	4.6	0.29	0.54	4.6	0.34	0.58	4.6	0.29	0.54	4.6	0.34	0.58
Overall satisfaction with accounts services team	4.6	0.30	0.55	4.7	0.23	0.48	4.6	0.34	0.58	4.6	0.46	0.68
Overall satisfaction with compute resources	4.5	0.41	0.64	4.4	0.48	0.69	4.6	0.40	0.63	4.6	0.50	0.71

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 and 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 2,198 tickets in response to user queries for CY 2018. The center exceeded the problem-resolution metric and responded to 92% of the queries within 3 business days (Table 1.6).

Table 1.6. Problem resolution metric summary

Survey Area	CY 2017		CY 2018	
	Target	Actual	Target	Actual
Percentage of problems addressed in 3 business days	80%	93%	80%	92%
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 two reported categories in 2018 were NCCS access and running jobs (Figure 1.1), which is very similar to the results from 2017.

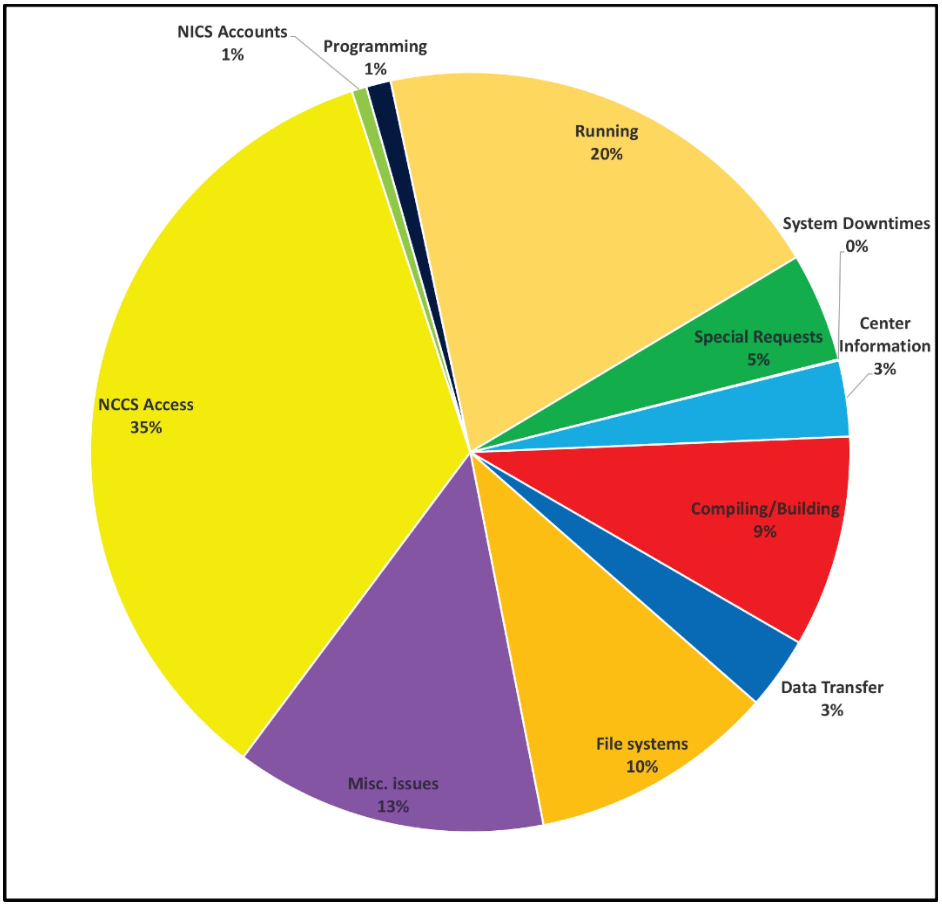

Figure 1.1. Categorization of help desk tickets.

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 2018 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. Below are some examples of UAO initiatives in 2018 that helped improve the overall user experience, albeit some them very much behind the scenes.

In 2018 members of the OLCF UA group continued to be heavily involved in the process to prepare for production and transition users onto the new IBM Power9 system, Summit. During this process, members of the UA group participated in CORAL working groups and also worked closely with members of the Summit acceptance team to understand issues, report bugs, and help drive development of IBM provided tools.

Job Launcher: The job launcher, jsrun, used on Summit was developed by IBM for the CORAL systems at Oak Ridge and Livermore. The new launcher provides similar functionality to mpirun, but also introduces resource organization techniques which are unique to jsrun. Because the job launcher plays a vital role in the efficient utilization of compute resources, the UA group focused additional resources on jsrun documentation and training. jsrun, has been among the largest drivers of user support tickets, so far on the system. In response, the OLCF developed and deployed jsrunVisualizer, an interactive online tool to aid users in understanding resource allocation and process placement for jobs launched using jsrun. The jsrunVisualizer tool allows for quick prototyping by graphically representing how multiple jsrun options would interact together, while also generating an example job script. It runs entirely in a web browser, allowing users to determine an appropriate resource layout for their application without waiting in a queue or spending their project allocation. Sharing capabilities were designed to improve collaboration and communication within project teams and facilitate conversation with OLCF support staff. While development to add new features is ongoing, jsrunVisualizer has proven useful for both new and experienced Summit users.

Batch Scheduler: The batch scheduler is another component all users of a system's compute resources must learn. Summit utilizes IBM's LSF batch scheduler to organize and distribute work onto the compute resources. Because LSF is new to the OLCF, the UA group placed additional focus on providing complete and detailed training and documentation to aid user transition from other batch schedulers. UA group members also worked with local staff and LSF developers to understand the scheduler logic, configuration, APIs, and tools needed to meet the center's scheduling policy needs. The UA group developed submission wrappers and verification codes to help enforce center policies and improve the user friendliness of the provided tools. UA also extended the scheduler configuration management features of the Resource and Allocation Tracking System (RATS) CRM so it can interface directly with IBM's LSF via LSF's Python-based API. This required code improvements to IBM's open-source Python LSF API library as well. With these improvements, limits and batch priorities on Summit for users, projects, and even entire allocation programs can be controlled easily from within the center's CRM software. This critically important functionality helps ensure that the center can meet operational goals for hours delivered to each allocation program.

Documentation: Web documentation is searchable and available 24-hours a day; it is one of the center's largest training, instructional, and reference tools. Web documentation is an important tool for users of a new system and continues to serve as reference for veteran users through the life of a system. Through testing, working group participation, and by working closely with OLCF staff, the UA group created a full and detailed web documentation for Summit. The web documentation was added to the OLCF public web site prior to early user access and will continue to be updated through the life of Summit as the system evolves.

Software: The OLCF user community is diverse and requires a diverse set of software. It is important for the center to help meet diverse software needs by providing multiple compilers, software packages, and multiple versions of each. Drawing from experience and user feedback, the UA group designed a software environment that utilizes lmod, Spack, and a custom continuous integration process to provide multiple software packages in an organized and familiar module system.

The User Assistance Development Team (UA Dev) deployed a number of software features to improve the center's internal Customer Relationship Management software (RATS CRM), leading to greater operational efficiencies for the center in calendar year 2018.

Support for New Project Automation Services: In order to support new offerings to end-users featuring on-demand containerized services via RedHat OpenShift, special "automation users" can now be easily created and managed in RATS CRM. OLCF staff can use these Project-centric UNIX users to run services like CI/CD, HPC workflows, etc., on behalf of projects.

New "Open Security Enclave" User Accounts: In order to accommodate greater flexibility in user authentication methods, UA Dev extended RATS CRM to create and manage "open" user accounts. These open accounts allow rapid onboarding of new users who only need access to certain resources that exist in an open security enclave, such as our new Ascent system. These accounts use 1-factor authentication methods, making them good candidates for situations that need to employ automated workflows like CI/CD and training. Open user accounts are backed by ORNL's laboratory-wide account and password management systems (XCAMS), which were extended with new APIs specifically to accommodate these new accounts for the OLCF. Using these new features, OLCF staff were able to process nearly a hundred accounts for new users during training sessions at SC'18, start to finish, in less than 15 minutes.

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 Development of Genomics Algorithm to Attain Exascale Speeds

The INCITE project, "Co-evolutionary Networks: From Genome to 3D Proteome," focuses on the analysis of Populus species and hybrids for their potential as renewable feedstocks for production of biofuels. Fuels converted from cellulosic biomass are an alternative energy source for supplementing fossil fuels with biofuels that can benefit economic growth and energy security of the U.S. This project studies cooccurrence patterns to discover interactions in the genetic materials in the Populus species which give rise to structural interactions of proteins at the cellular level.

To support this work, the CoMet (Combinatorial Metrics) application was developed at the OLCF by Wayne Joubert, INCITE project liaison, as lead developer. CoMet is able to compute 2-way and 3-way interactions within genetic data at a rate that is 20,000 to 300,000 times faster than previous state of the art methods, making it possible to solve problems in hours that would normally require years of runtime on leadership systems. In 2018, CoMet was also adapted to the Summit platform's Volta GPU tensor core architecture and became the world's first application to achieve Exascale performance, reaching 2.36 billion billion mixed precision floating point computations per second. Additional details of the CoMet application, a recipient of the 2018 Gordon Bell Prize, are highlighted in the following paper: https://dl.acm.org/citation.cfm?id=3291732.

1.4.3.2 Advanced Parallel Segmentation Algorithms for Faster Data Analysis

Two-fluid flow in porous media occurs routinely in natural and engineered geologic systems and is important in applications such as geologic carbon sequestration, oil and gas recovery, and contaminant transport. The goal of the INCITE project "Advancing Models for Multiphase Flow in Porous Medium Systems" led by James McClure, Virginia Tech, is to use digital rock physics techniques to advance multiscale models for two-fluid flow that are consistent across disparate length scales and resolve the operative physics with higher fidelity than existing models. Recent theoretical, experimental, and computational developments offer the opportunity to improve macroscale models by synthesizing information from different scales to better describe the physics.

Experimental data volumes that capture the movement of fluids within rock can be produced by synchrotron light sources such as the Advanced Photon Source (APS) at Argonne National Laboratory. "Fast" micro-tomography (uCT), which allows for 4D imaging of microscopic flow processes that occur within complex materials, collects 3D images (1920 x 1920 x 1200 voxels) every 45 seconds based on continuous scans. To identify the location of fluid within the experimentally-generated image, each 3D image needs to be segmented to produce a labeled map of the materials that are present. Segmentation routines often rely on computationally intensive algorithms such as the non-local means algorithm, and parallel implementations provide the only viable path to analyze data in real-time.

OLCF Computational Scientist Mark Berrill developed a parallel segmentation pipeline to provide this capability, reducing segmentation times to approximately one minute, a several orders of magnitude improvement. Based on this approach, experimental data collected at the GSECARS beamline at the Advanced Photon Source (APS) at Argonne National Lab could be reconstructed on-site at the APS, transferred to the OLCF, analyzed and visualized within 15 minutes of data collection. The ability to perform data analysis in real-time allows scientists to make on-the-fly adjustments to experimental approaches that can enhance the quality and value of the data collected. The segmentation tool is currently being applied to generate labeled data that will be used to train the next generation of AI approaches to further enhance experimental data processing capabilities.

1.4.3.3 Improved DCA++ Application Performance Enhances Scientific Productivity

The INCITE project "Computational Studies of Spin-Fluctuation Fingerprints in Cuprates" aims to understand the pairing mechanism in the cuprate high-temperature superconductors, in which the exchange of spin fluctuations is believed to give rise to superconductivity. The project team, led by ORNL's Thomas Maier, carried out theoretical calculations of the single-particle spectral function to identify the fingerprints of a spin-fluctuation interaction. This is an important step in advancing the theory of high-temperature superconductivity, which remains one of the most challenging problems in condensed matter physics.

The team uses the open-source scientific software, DCA++ (https://github.com/CompFUSE/DCA), that implements the dynamic cluster quantum Monte Carlo algorithm. The DCA++ code uses MPI and C++ Threads in the modern C++ standard library for parallelization, CUDA and MAGMA for GPU acceleration of linear algebraic operations. As an application that is being constantly and rapidly developed to tackle increasingly complex scientific problems, DCA++ underwent a major refactoring in the last few years to ensure sustainable software development. One of the major changes is the adoption of C++ Threads to replace the original Pthreads implementation. While the Pthreads code has been proven to run and scale efficiently on OLCF's Titan, the performance of the refactored code, along with many newly added capabilities, were not profiled and optimized for OLCF machines.

OLCF computational scientist Ying Wai Li, along with Arghya (Ronnie) Chatterjee and Oscar Hernandez of ORNL's Computer Science Research Group, performed a series of detailed profiling of the new DCA++ code on various OLCF machines. This involved obtaining profiles and traces on Titan and Summit using different profiling tools such as Rice's HPCToolkit, Score-P, and NVIDIA's nvprof for a comprehensive analysis. They compared the performance and code behavior on different platforms, identified the hotspots and load imbalances in the code for optimization opportunities, and diagnosed a major implementation flaw with mutex locks that dominated the computational time. These informative profiles and diagnoses enabled the DCA++ developers to improve the threading implementation and reduce a significant amount of runtime. This information also provided insight into different research directions for improving the time-to-solution of the DCA++ application under the support of the PI's multi-year SciDAC effort, which includes the investigation of new programming model strategies, as well as the development of more efficient Monte Carlo algorithms.

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 three 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 128 users voted in the 2018 election and selected Evan Schneider, Sarat Sreepathi, and Lin-Wang Wang to serve on the OUG executive board for 3 year terms. The current board chair is Joe Oefelein. The board elected Mike Zingale to serve as the vice chair this year, and he will automatically become the chair for the 2019–2020 cycle.

The OLCF hosted a total of nine monthly OUG conference calls in 2018, with a total of 447 attendees (up from 345 attendees in CY2017). Each call highlighted a topic selected by UAO and the OUG Executive board to be of potential interest to users. In addition, the UAO OUG coordinator presented general tips, tricks, and reminders on each call.

On May 15-17, the OLCF hosted its annual OLCF User Meeting to share selected computational science and engineering achievements emerging from the OLCF's user programs, enable direct interactions among users, advance OLCF's relationship with our user community, and highlight computational requirements for the future. To do so, the meeting featured plenary talks, panels, poster sessions, center updates from OLCF staff, and paired talks from OLCF users and application developers. The event featured user success stories on both Titan and Summit on day one and two. The meeting's final day focused on OLCF data services and began with an overview of the agenda given by OLCF's Arjun Shankar. Jason Kincl (OLCF) followed with a presentation on the different uses of containers at OLCF titled "OLCF Container Orchestration for HPC Middleware". Next, Dale Stansberry (OLCF) gave a talk titled "Constellation: Digital Object Identifiers for OLCF Data", which described an ISO standard for unique, persistent, and resolvable identifiers for data "objects" (to make data citable like publications). Visualization of scientific data was then discussed in the context of the SIGHT software by Ben Hernandez (OLCF) in a talk titled "SIGHT". Finally, Valentine Anantharaj (OLCF) closed the day with a talk exploring data needs and requirements with a talk titled "OLCF Research Data Initiatives". In total, 118 OLCF users and staff members attended the meeting to share their achievements or discuss upcoming technologies. For those users who could not attend the meeting in person, a BlueJeans link was supplied to allow remote participation and the slides and recordings for all talks were posted to the 2018 OLCF User Meeting event page (https://www.olcf.ornl.gov/calendar/2018-olcf-user-meeting/).

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

HPC education at OLCF is delivered through a combination of training events, such as workshops, webinars, user conference calls, and seminar series. Such training events serve to engage the public and user community, and they are an integral part of learning how to perform science on large-scale HPC systems. In addition, because most researchers are funded to perform science (not software development), poor software practices can make their way into user's applications. To help address this, the OLCF also participates in a collaborative webinar series on Best Practices for HPC Software Developers.

In 2018, the OLCF facilitated or collaborated on 7 week-long GPU hackathons, 9 user conference calls, 11 Interoperable Design of Extreme-Scale Application Software – ECP (IDEAS-ECP) webinars, 2 INCITE proposal writing webinars, a 3-day Introduction to HPC Workshop, an Introduction to Summit Webinar, a 5-day Summit Application Readiness Workshop, a 4-day Summit Training Workshop, the annual OLCF User Meeting, and many other training activities. See Appendix B for a complete summary of these events. A few of the notable 2018 events are highlighted in Sections 1.4.5.1-1.4.5.4.

1.4.5.1 Summit Application Readiness Workshop

On March 5-9, the OLCF hosted a Summit Application Readiness Workshop, with the primary objective of providing the detailed technical information and hands-on help required for select application teams to meet the scalability and performance metrics required for Early Science proposals. During the workshop, Summit Phase 1 (25% system acceptance) was available in a configuration that allowed the teams to demonstrate these metrics using representative science runs. Technical representatives from IBM's Center of Excellence delivered plenary presentations, but most of the time was dedicated to the extended application teams to carry out hands-on technical work on Summit. This was an important milestone for the projects, and a requisite for a successful Early Science proposal. In total, 77 early Summit users attended the event (https://www.olcf.ornl.gov/summit-application-readiness-workshop/).

1.4.5.2 Introduction to HPC Workshop

On June 26-28, the OLCF hosted a 3-day Introduction to HPC Workshop meant to help train those new to high performance computing. Each day consisted of presentations to deliver new information and hands-on sessions to help reinforce the concepts. The presentations covered basic skills, such as UNIX, the vim text editor, and c/FORTRAN programming, before moving on to cover parallel programming (using MPI and OpenMP) and GPU computing (CUDA and OpenACC). The presentations were given by OLCF staff members as well as our vendor partners, NVIDIA and Cray.

The 100+ participants consisted mostly of ORNL staff, University of Tennessee students, and ORNL summer interns interested in learning about HPC. For those users who could not attend the meeting in person, a BlueJeans link was supplied to allow remote participation and the slides and recordings for all talks were posted to the Introduction to HPC event page (https://www.olcf.ornl.gov/calendar/introduction-to-hpc/).

1.4.5.3 Summit Training Workshop

On December 3-6, the OLCF hosted a 4-day Summit Training Workshop meant to help new Summit users (or those intending to use Summit) get up and running on the system. Each day consisted of presentations in the morning followed by hands-on sessions in the afternoons. The presentations delivered relevant information about the system (hardware, job launcher, programming methods, etc.) as well as talks on porting experiences (from early Summit users), while the hands-on sessions gave participants the opportunity to get their own codes running on Summit (Ascent). OLCF staff and our vendor partners (IBM and NVIDIA) were available to help during the hands-on sessions.

The primary target audience for this event were new INCITE projects with upcoming allocations on Summit, which began on January 1, 2019. To give these users a head start, in-person attendees of the event received access to Ascent, an 18-node system with identical hardware/nodes and software as Summit, which continued until January 2019 when their actual Summit access began. For those users who could not attend the meeting in person, a BlueJeans link was supplied to allow remote participation (to watch presentations, but not participate in the hands-on session) and the slides and recordings for all talks were posted to the Summit Training Workshop event page (https://www.olcf.ornl.gov/calendar/summit-training-workshop/) to benefit attendees and future users as well. In total, 171 people attended the workshop either in-person or remotely

1.4.5.4 2018 OLCF GPU Hackathons

In 2018, the Oak Ridge Leadership Computing Facility (OLCF) and their partners continued to expand the annual OLCF GPU Hackathon series. These events taught new GPU programmers how to leverage accelerated computing in their own applications and helped existing GPU programmers to further optimize their codes. In total, seven hackathons were held by the following host institutions:

- Technische Universitat (TU) Dresden (Dresden, Germany)
- Pawsey Supercomputing Centre (Perth, Western Australia)
- University of Colorado (Boulder, CO)
- NCSA – National Center for Supercomputing Applications (Urbana-Champaign, IL)
- Brookhaven National Laboratory (Upton, NY)
- CSCS – Swiss National Supercomputing Centre (Lugano, Switzerland)
- Oak Ridge Leadership Computing Facility (Knoxville, TN)

These institutions hosted 64 teams (compared to 47 in 2017); bringing GPU-programming expertise to more than 400 attendees from universities, national laboratories, supercomputing centers, government institutions, and industry. The teams consisted of three or more developers of a particular application, along with two mentors with extensive GPU programming experience. Together, the teams worked to port or optimize their applications on GPU-accelerated systems, including local clusters and workstations, development systems such as SummitDev (OLCF) and Athena (Pawsey), as well as world-class HPC systems such as Piz Daint and Summit. During the events, teams developed applications spanning a wide range of scientific domains, including (but not limited to) astrophysics, climate modeling, combustion, computational fluid dynamics, machine learning, molecular dynamics, plasma physics, and quantum mechanics.

After the events, some participants acknowledged the benefits they obtained during the hackathons in articles and peer-reviewed journals, which helps to expand the use of GPU-accelerated computing within their specific scientific domains (as well as the larger scientific community). For example, a December 2018 article in HPCwire (https://www.hpcwire.com/2018/12/05/optimizing-key-plasma-physics-code-for-latest-gen-nvidia-gpus-yields-threefold-increase-in-processing-speed/) highlighted a 3X speedup that a team of General Atomics researchers achieved in their plasma physics code, thanks in part to work performed at the hackathon held at the University of Colorado – Boulder. This team worked with

NVIDIA mentors to help them take advantage of new capabilities of the Summit architecture (e.g. CUDA-aware MPI enhanced with GPUDirect – allowing data transfers directly between peer GPUs). Another example is an article (https://ieeexplore.ieee.org/document/8408902) that was published in August 2018, which specifically highlighted the OLCF GPU hackathon series. In addition to describing the goals, format, and community benefits of the events, the article includes case studies from previous successful hackathon teams.

The hackathons also serve as an outreach opportunity for the OLCF, introducing the participants to the OLCF resources, nurturing relationships with vendor partners and leaders of host institutions, and encouraging teams to submit Director's Discretionary (DD) proposals for new OLCF projects. Examples of such proposals from 2018 include an ORNL quantum chemistry project to prepare for an upcoming INCITE proposal, a GE computational fluid dynamics project for continued GPU development, and a University of Michigan molecular dynamics project intended to prepare for an upcoming INCITE proposal.

It is always important to note that mature, optimized applications require more development than can be accomplished during a single, week-long hackathon event; but the testimonials that we receive from previous attendees as well as their publication records continue to show the impact of these events.

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

1.4.6.1 CSGF Annual Program Review Brings Staff and Students Together

For the ninth 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 15-19, 2018, in Arlington, Virginia, provided CSGF fellows with an outlet for research discussion and one-on-one meetings with DOE laboratory staff. The OLCF and supercomputing centers at Lawrence Livermore and Lawrence Berkeley National Laboratories led a 2-day workshop with support from three other national labs with major supercomputing centers—Argonne, Los Alamos, and Sandia National Laboratories. The labs provided speakers, small-group mentors, and training resources to deliver this year's HPC Workshop which was a 2-day, hands-on opportunity for fellows to learn HPC programming basics and run a program on a DOE supercomputer. The HPC Workshop has evolved over the years from a lecture-based program to a hands-on experience. This was the first time that all fellows, in their first through fourth years of the fellowship, were invited to participate in this mini hackathon. After workshop presentations on HPC programming for scientific computing, fellows were provided a piece of code to program for GPU processors, such as those used in hybrid supercomputers like the 27-petaflop Titan. The OLCF provided the fellows access to Titan during the workshop and continued to provide this access for fellows who want to continue to work on their training program for an additional few months.

1.4.6.2 Introduce Your Daughter to Code

For the third 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 field of computational sciences with programming activities that centered around running code on the OLCF's flagship supercomputer, Titan. This year, 24 girls ages 10 to 18 participated in the labwide event. The team guided 24 girls and their parents through machine-learning activities to demonstrate how computers assimilate information by analyzing input data and producing output based on "learned" patterns. The girls first explored an interactive 3D visualization of a fully connected convolutional neural network, which let them draw numbers that the network would then "guess" based on a grid system. Convolutional neural networks are artificial neural networks that employ a sliding "window" to look at information; therefore, they are exceptional at analyzing visual input such as images. Using sketch-rnn, an interactive experiment created by Google research scientist David Ha, the girls learned about recurrent neural networks, which are designed to process sequential information. Sketch-rnn suggests ways to finish a drawing based on the initial input it is given. Having trained on millions of doodles from the Quick, Draw! game, sketch-rnn offers more than 100 models. The girls also navigated the TensorFlow playground—an educational visualization of an artificial neural network—and played a human neural network game. Katie Schuman and team colleague Steven Young collaborated with the Neutron Scattering Division's Thomas Proffen to develop the game, which demonstrates through play how a machine receives input, interprets it, and produces output. During the game, the girls each took on roles as different kinds of neurons, which are cells that transmit information in the brain. Neural networks build on this biological concept.

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 (see Appendix C). In 2018, the Outreach team focused significant energy on launching the Summit supercomputer, developing materials and stories that introduced the new machine and celebrated Summit users' early science achievements.

In June 2018, the OLCF collaborated with ORNL and DOE public affairs teams to host a Summit launch event, which included VIPs, media, and other guests. This event was the centerpiece to a year of strategic planning and execution. The team knew there was a possibility of Summit being named the fastest computer in the world at the end of June, a news-worthy achievement that would focus on the computer's performance but not its applications. The team felt that a public debut ahead of the June 25 announcement would give the OLCF and the Laboratory the chance to celebrate what Summit will be able to do for the global science community, beyond just being fast. The outreach team coordinated a meeting with the public relations and marketing teams of vendors supporting the supercomputer in February of 2018 to plan for the Summit launch. With input from those teams, the OLCF team developed a strategic plan that they pitched to OLCF, CCSD, ORNL Communications, and Laboratory leadership. Then, they pitched their plan to the DOE's Office of Science and worked with them to pitch it to US Secretary of Energy Rick Perry. They invited the Secretary to attend a launch event, a request he accepted. The CEOs of IBM and NVIDIA also attended the event along with the Tennessee governor, University of Tennessee president, and a number of key DOE representatives.

The outreach team created banners and displays for the launch event, including a dedication plaque unveiled by ORNL Director Thomas Zacharia and a banner announcing that Summit had broken the exascale barrier. Additionally, team members worked to produce a polished webpage with more detailed information about Summit's capabilities and deployed a social media strategy that included branded tiles with Summit facts and quotes from early science researchers. Their overall strategy was to establish Summit and the OLCF as the premier location for researchers seeking to run codes at unprecedented speeds to pursue scientific innovation not possible on Titan or any other previous systems.

Before, during, and after the June event, the team monitored and responded to hundreds of media queries and requests for interviews to facilitate coverage of the event and developed various press materials, from designing the display overlooking the supercomputer to writing press releases, feature stories, and fact sheets covering different aspects of the Summit story. The team coordinated on-site media, including CNBC and CNN, and logistics for satellite truck live shot interviews with media outlets across the country. They fielded interview requests from and were published in high profile outlets like Wired, New York Times, CNET, Wall Street Journal, and in more than 1,900 print and television outlets. They also developed videos, including a time-lapse of staff building Summit and a series of interviews with the principal investigators leading early science projects selected to run on Summit in the near future, including explorations into cancer surveillance, systems biology, and astrophysical phenomena. Social media analytics for the OLCF accounts during the Summit campaign period showed a 230%

increase in daily reach on Facebook and 324% increase in Twitter impressions over the same period in 2017.

Packaging all this information cohesively required extensive internal planning, but the team also sought to amplify the message from external sources. To ensure extensive media coverage, the team continued coordinating with public relations executives from vendor partners to notify national and international press of the launch event, which resulted in extensive coverage of Summit's debut from notable outlets such as *WIRED*, *CNBC*, and *The New York Times*.

From the launch event in early June to the Top500 announcement later in the month, Summit remained prevalent in national and international news cycles far beyond the initial announcement, meaning the outreach team successfully facilitated nearly a full month of media coverage for maximum exposure.

The DOE Public Affairs Office noted that the interactions and coordination of the Summit launch was the gold standard for interagency coordination of large-scale announcements, proving that the OLCF outreach team not only met expectations, but exceeded them.

In addition to the Summit-focused activities in 2018, the team was also responsible for the creation of 76 highlights—including science and technology highlights and features about Titan users, OLCF staff members, and OLCF resources—and more than 213 total outreach products. Another way the OLCF reaches the public is through tours to groups of visitors who range from middle-school students through senior-level government officials. The center conducted tours for 365 groups in 2018, which included more than 5,490 individuals.

In all, the Outreach team produced at least 28 science highlights in 2018. Those highlights touched on a variety of science domains, from energy materials ("Titan Takes Fluoride from Taps and Toothpaste to Batteries"), to combustion ("GM Revs up Diesel Combustion Modeling on Titan Supercomputer"), to nuclear physics ("Nuclear Physicists Wield HPC to Uncover Magic Isotopes"). In addition, the team completed 20 technology stories and 24 people features. The team continued the "Faces of Summit" profile series they began in 2017, producing 10 additional profiles in 2018.

1.5 LOOKING FORWARD

1.5.1 Application Readiness and Early Science

OLCF's Center for Accelerated Application Readiness (CAAR) program is a partnership of OLCF staff, scientific application teams, vendor partners, and tools developers with the goal of readying a set of applications for the Summit architecture. Thirteen CAAR applications were selected after a call for proposals in late 2014 and cover a broad range of scientific disciplines and employ a range of programming models and software designs. Progress of the CAAR projects was reviewed in the 2016 Independent Project Review (IPR), and all projects were found to be on track to meet the scalability and performance metrics as defined for the key performance parameters (KPP). The application readiness phase of the CAAR program ended at the end of CY2018 and the projects have transitioned into the Early Science phase, which will end June 30, 2019. 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/E3SM	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 Bykov	Chemistry
NAMD	Dr. James Phillips University of Illinois–Urbana-Champaign	Tjerk Straatsma	Biophysics
NUCCOR	Gaute Hagen Oak Ridge National Laboratory	Gustav Jansen	Nuclear physics
NWCHEM	Karol Kowalski Pacific Northwest National Laboratory	Dmitry Liakh and Tjerk Straatsma	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

The KPP metrics were defined to be achieving scalability, i.e. reduced time to solution, for a challenging scientific run beyond 20% of Summit, while demonstrating performance improvement of a factor of two from use of the GPUs. Recognizing software development and porting efforts focused on Summit outside of the thirteen CAAR applications, a call for proposals for Early Science projects was issued at the end of CY2017. More than 60 teams, including the 13 CAAR teams, submitted letters of intent expressing interest in submitting a full proposal to the OLCF Early Science program. These letters of intent were due December 31, 2017. The proposal teams were given access to a partition of Summit in the first half of 2018 for code benchmarking purposes to allow them to prepare a full Early Science proposal, which was due June 30, 2018. Selection was made based on the scientific challenge problem proposed and meeting the KPP metrics for scalability and GPU performance. A total of 33 projects were selected and were provided access to Summit for their science runs until June 30, 2019.

1.5.2 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 within 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 2018, the CSEEN Postdoctoral program at the OLCF supported five trainees:

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 was to develop workflows for efficient massively parallel implementation of seismological methodologies on current state-of-the-art accelerated computer architectures, and he was 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. Yangang joined the Zhejiang University as an Assistant Professor in April 2018.

Ana Maria De Carvalho Vicente Da Cunha joined the SciComp group in July 2018 after earning her PhD from the University of Groningen in Theory and Condensed Matter Physics with a focus on computational spectroscopy. She has experience in applying both classical molecular dynamics and quantum mechanics to the study of biomolecular systems, which will be the focus of her OLCF post-doc.

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 worked 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. Austin joined the OLCF Scientific Computing Group as a Performance Engineer in January 2019.

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 is 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.

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.

Operational Performance

HIGH PERFORMANCE COMPUTING FACILITY
2018 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2019

2. OPERATIONAL PERFORMANCE

CHARGE QUESTION 2: *Did the facility's operational performance meet established targets?*

OLCF RESPONSE: Yes. The OLCF provides a series of highly capable and reliable systems for the user community. The 2018 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 operational performance 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. In 2018, the OLCF once again delivered all of the compute hours committed to the three major allocation programs: INCITE, ALCC, and DD. The operational performance demonstrates that the OLCF delivered another prominent operational year of reliable and technically sufficient resources to the scientific research community. This was all accomplished while deploying the new Summit and Spider 3 systems for production in CY 2019.

2.1 OPERATIONAL PERFORMANCE SUMMARY

Operational performance measures the performance of the OLCF against a series of operational parameters. The two operational metrics relevant to the OLCF's operational performance 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 IBM AC922 (SUMMIT) RESOURCE SUMMARY

The OLCF installed and deployed an IBM AC922 system named Summit, which will become available for full production in January of 2019. Summit is comprised of 4,608 high-density compute nodes each equipped with 2 IBM POWER9 CPUs and 6 NVIDIA Volta GPUs. In total, the Summit system is capable of 200 petaflops of peak computational performance and was recognized as the most powerful system in the world for its performance on both the high performance linpack (HPL) and conjugate gradient (HPCG) benchmark applications.

2.3 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.

2.4 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.

2.5 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.

2.6 GPFS FILE SYSTEM (ALPINE) RESOURCE SUMMARY

In January 2019, the OLCF will release Spider III, its next-generation global file system to support the computational resources in the OLCF. Spider III is a single GPFS namespace named Alpine, with a usable capacity of 250 PB and a filesystem-level performance of 2.5 TB/s. The Spider III file system will be the default high-performance parallel file system for the OLCF's moderate compute resources.

2.7 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.

2.8 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.

2.9 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.

2.10 OLCF COMPUTATIONAL AND DATA RESOURCE SUMMARY

The OLCF provided the Titan and Eos computational resources and the Spider II and HPSS data resources for production use in 2018 (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, 2018

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

2.10.1 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 first 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	August 1, 2019	Delivered as two separate file systems, /atlas1 and /atlas2. 30+ PB capacity
Eos	Cray XC30	October 3, 2013	August 1, 2019	Production with 736 Intel E5, 2,670 nodes.
Titan	Cray XK7	May 31, 2013	August 1, 2019	Production with 18,688 hybrid CPU-GPU nodes (AMD Opteron 6274/NVIDIA K20X)
JaguarPF	Cray XK6	September 18, 2012	October 7, 2012	Production at 240,000 cores until September 18, when partition size was reduced to 120,000 AMD Opteron cores. Additional Kepler installation. TitanDev access terminated
JaguarPF	Cray XK6	February 13, 2012	September 12, 2012	Full production until September 12, when partition size was reduced to 240,000 AMD Opteron cores. Beginning of Kepler installation
JaguarPF	Cray XK6	February 2, 2012	February 13, 2012	Stability test. Restricted user access. 299,008 AMD Opteron 6274 cores. Includes 960-node Fermi-equipped partition
JaguarPF	Cray XK6	January 5, 2012	February 1, 2012	Acceptance. No general access 299,008 AMD Opteron cores
JaguarPF	Cray XK6	December 12, 2011	January 4, 2012	142,848 AMD Opteron cores
JaguarPF	Cray XT5	October 17, 2011	December 11, 2011	117,120 AMD Opteron cores
JaguarPF	Cray XT5	October 10, 2011	October 16, 2011	162,240 AMD Opteron cores
JaguarPF	Cray XT5	September 25, 2009	October 9, 2011	224,256 AMD Opteron cores
JaguarPF	Cray XT5	August 19, 2008	July 28, 2009	151,000 AMD Opteron cores

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

2.10.2 Operational Performance Snapshot

Operational performance metrics 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 Operational performance summary for Titan

	Measurement	2017 target	2017 actual	2018 target	2018 actual
Cray XK7 (Titan)	Scheduled availability	95%	99.39%	95%	98.97%
	Overall availability	90%	98.09%	90%	98.51%
	MTTI[a] (h)	NAM[c]	660.95	NAM	784.5
	MTTF[b] (h)	NAM	1,741.49	NAM	1,238.61
	Total usage	NAM	91%	NAM	90.6%
	Core-hours used[d]	NAM	4,389,163,123	NAM	4,382,897,140
	Core-hours available	NAM	4,817,215,104	NAM	4,838,033,536
	Capability usage				
	INCITE projects	NAM	68.22%	NAM	57.59%
	All projects	35%	59.81%	35%	51.09%

[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.

Table 2.4. OLCF Operational performance summary for Eos

	Measurement	2017 target	2017 actual	2018 target	2018 actual
Cray XC30 (Eos)	Scheduled availability	NAM[c]	99.61%	NAM	99.21%
	Overall availability	NAM	98.39%	NAM	98.86%
	MTTI[a] (h)	NAM	783.58	NAM	1,237.22
	MTTF[b] (h)	NAM	2,908.61	NAM	2,172.65

[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 Operational performance summary for HPSS

	Measurement	2017 target	2017 actual	2018 target	2018 actual
HPSS	Scheduled availability	95%	99.46%	95%	99.87%
	Overall availability	90%	98.87%	90%	99.53%
	MTTI[a] (h)	NAM[c]	541.3	NAM	792.6
	MTTF[b] (h)	NAM	1,244.73	NAM	1,749.71

[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 Operational performance summary for Spider II, the external Lustre file system

	Measurement	2017 target	2017 actual	2018 target	2018 actual
/atlas1	Scheduled availability	95%	99.59%	95%	99.95%
	Overall availability	90%	98.87%	90%	99.95%
	MTTI[a] (h)	NAM[c]	509.48	NAM	2,188.9
	MTTF[b] (h)	NAM	872.39	NAM	2,188.9
/atlas2	Scheduled availability	95%	99.59%	95%	99.29%
	Overall availability	90%	98.88%	90%	99.29%
	MTTI (h)	NAM	433.08	NAM	724.82
	MTTF (h)	NAM	671.11	NAM	724.82

[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 5 year production anniversaries in 2018. The reported results for each system measure are for CY 2018 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 2018, the data reflected here artificially assume calculation boundaries of 00:00 on January 1, 2018, and January 1, 2019.

2.11 RESOURCE AVAILABILITY

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

2.11.1 Scheduled Availability

The scheduled availability is described by Eq. (1). The OLCF has exceeded the SA targets for the facility's computational resources for 2017 and 2018 (Table 2.7).

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

Table 2.7. OLCF Operational performance summary: Scheduled availability

	System	2017 target	2017 actual	2018 target	2018 actual
Scheduled availability	Cray XK7	95%	99.39%	95%	98.97%
	Cray XC30	NAM[a]	99.61%	NAM	99.21%
	HPSS	95%	99.46%	95%	99.87%
	/atlas1	95%	99.59%	95%	99.95%
	/atlas2	95%	99.59%	95%	99.29%

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

2.11.1.1 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 2018, OLCF staff executed scheduled maintenance on the Cray XK7 a total of four times, which were all related to compute system software testing and updates. Six unscheduled outages were experienced in 2018, which included two power interruptions, three compute system interruptions, and one file system issue. Similarly, the OLCF performed scheduled maintenance on Eos three times in 2018, with three unscheduled outages that included the same power and file system interruptions that affected the Cray XK7 system.

In 2017, representatives from OLCF, ALCF and NERSC agreed that during a scheduled maintenance, a significant event that delays the return of a system to scheduled production by more than 4 hours will be counted as an adjacent unscheduled outage, as unscheduled availability, and as an additional interrupt. Prior to 2017 each facility was handling this situation in a different way and we were asked by ASCR to come up with a consistent definition.

2.11.2 Overall Availability

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

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

As shown in Table 2.8, the OLCF exceeded the OA targets of the facility's resources for 2017 and 2018.

Table 2.8. OLCF Operational performance summary: Overall availability

	System	2017 target	2017 actual	2018 target	2018 actual
Overall Availability	Cray XK7	90%	98.09%	90%	98.51%
	Cray XC30	NAM[a]	98.39%	NAM	98.86%
	HPSS	90%	98.87%	90%	99.53%
	/atlas1	90%	99.95%	90%	99.95%
	/atlas2	90%	98.88%	90%	99.29%

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

2.11.3 Mean Time to Interrupt (MTTI)

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

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

The MTTI summary is shown in Table 2.9.

Table 2.9. OLCF Operational performance summary: Mean time to interrupt (MTTI)

	System	2017 actual	2018 actual
MTTI (h)	Cray XK7	660.95	784.5
	Cray XC30	783.58	1,237.22
	HPSS	541.3	792.6
	/atlas1	509.48	2,188.9
	/atlas2	433.08	724.82

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

2.11.4 Mean Time to Failure (MTTF)

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

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

Table 2.10. OLCF Operational performance summary: Mean time to failure (MTTF)

	System	2017 actual	2018 actual
MTTF (h)	Cray XK7	1,741.49	1,238.61
	Cray XC30	2,908.61	2,172.65
	HPSS	1,244.73	1,749.71
	/atlas1	872.39	2,188.9
	/atlas2	671.11	724.82

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

2.12 RESOURCE UTILIZATION 2018

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.

2.12.1 Resource Utilization Snapshot

For the Cray XK7 during the operational assessment period January 1–December 31, 2018, 4,382,897,140 Titan core-hours were used outside of outage periods from an available 4,838,033,536 Titan core-hours. The total system utilization for the Cray XK7 was 90.6%.

2.12.2 Total System Utilization

2.12.2.1 2018 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.

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

The measurement period is for 2018, 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 2018 was 90.6%, which marks the sixth 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 2018. Figure 2.1 represents the three major OLCF user programs and usage by the Exascale Computing Project (ECP) but does not include consumed core-hours from staff or vendor projects.

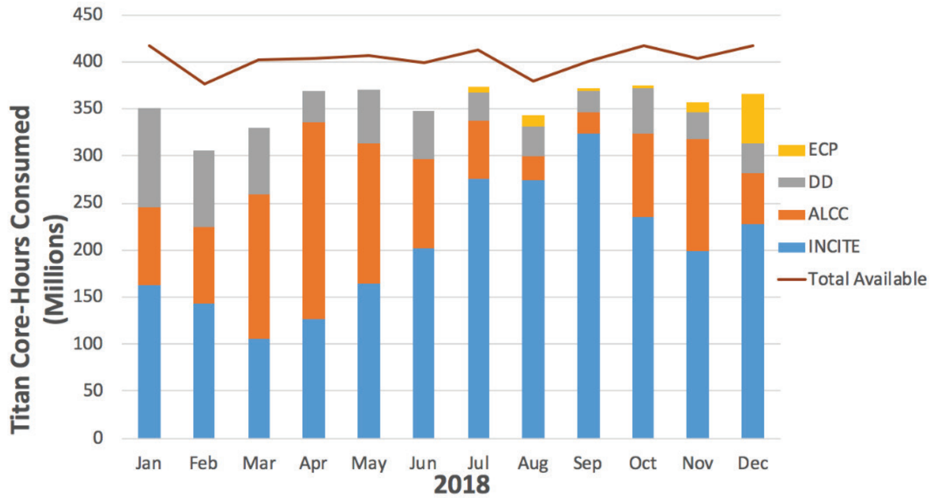

Figure 2.1. 2018 XK7 resource utilization—Titan core-hours by program.

2.12.2.2 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 2018 INCITE allocation program was the largest program in 2018, with a commitment for 2.25 billion Titan core-hours. The consumption of these allocation programs is shown in Table 2.11. As shown, all commitments were exceeded for each allocation program on Titan for 2018. This programmatic overachievement is assisted by the high uptime and diligent work of the OLCF operational staff for 2018.

Non-renewed INCITE projects from 2017 continued running through January 2018 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 2018 allocation period (ending June 30, 2018) were also granted extensions where appropriate.

Table 2.11. The 2018 allocated program performance on Titan

Program	Allocation	Hours consumed	Percent of total
INCITE[a]	2,250,000,000	2,438,640,309	58.42%
ALCC[b]	Allocation spans multiple CY	1,141,541,949	27.35%
DD	—	593,986,326	14.23%
Total		4,174,168,584[c]	100%

[a] Includes all INCITE program usage for CY 2018
[b] Includes all ALCC program usage for CY 2018
[c] Does not include usage outside of the three primary allocation programs

2.13 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 2018 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 2017 target	CY 2017 actual	CY 2018 target	CY 2018 actual
INCITE	NAM[a]	68.2%	NAM	57.6%
ALCC	NAM	56.7%	NAM	49.8%
All projects	35%	59.8%	35%	51.1%

[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, 51.1%, was once again well above the 2018 target of 35%. Although this number is down from CY 2017, it is still above the target metric. This slight dip in capability usage could be due to the age of Titan or could be attributed to the significant ramp-up of application work by project teams on the ECP or OLCF-4 (Summit) projects. In all, it is not surprising to see this slight decline in capability usage for a system the age of Titan (5-year production anniversary in 2018). 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.

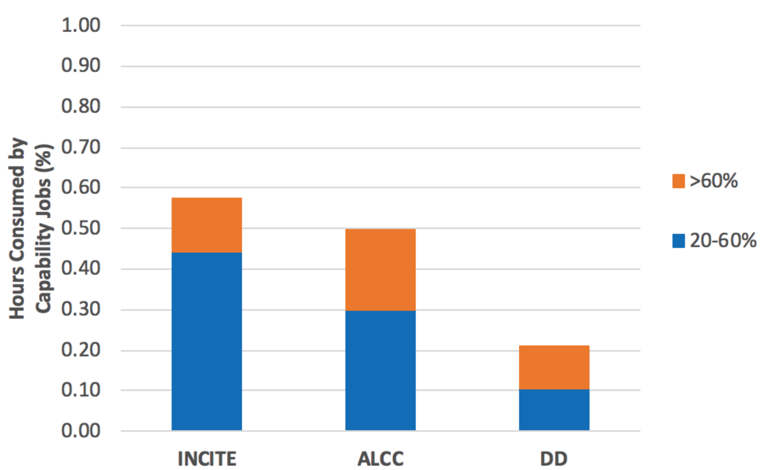

Figure 2.2. Capability usage by job size bins and project type.

2.14 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 2018, 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 a split of GPU and CPU enabled jobs. The DD program totaled roughly 30% usage for GPU-enabled compute hours. When compared with CY 2017, the programmatic breakdowns for GPU-enabled usage is almost identical for each of the user programs with the exception of the ALCC program. In 2017 the ALCC program boasted a 72% GPU-enabled usage number but the 2018 numbers show a fairly significant decline among the ALCC program's GPU-enabled usage on Titan at 50%. Even with this decline in the ALCC program, 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. The ALCC program does not require computational readiness reviews specifically for GPU usage and therefore the user facility cannot accurately predict the expected GPU-enabled usage for the ALCC program. With this in mind, a 50/50 split is very reasonable. For most months in CY 2018, GPU-enabled INCITE applications were consistently responsible for more than half of the delivered hours to those projects and all user programs averaged 56% in 2018.

Table 2.13. 2018 GPU-enabled and CPU-only usage by program

Program	Percentage	Hours
INCITE (GPU-enabled)	64.89	1,582,346,591
INCITE (CPU-only)	35.11	856,293,718
ALCC (GPU-enabled)	49.93	569,972,917
ALCC (CPU-only)	50.07	571,569,032
DD (GPU-enabled)	30.36	180,315,994
DD (CPU-only)	69.64	413,670,332

Approximately 56% of all delivered compute time on Titan was consumed by GPU-enabled applications in CY 2018. Figure 2.3 shows the percentage of GPU-enabled compute time by month and user program.

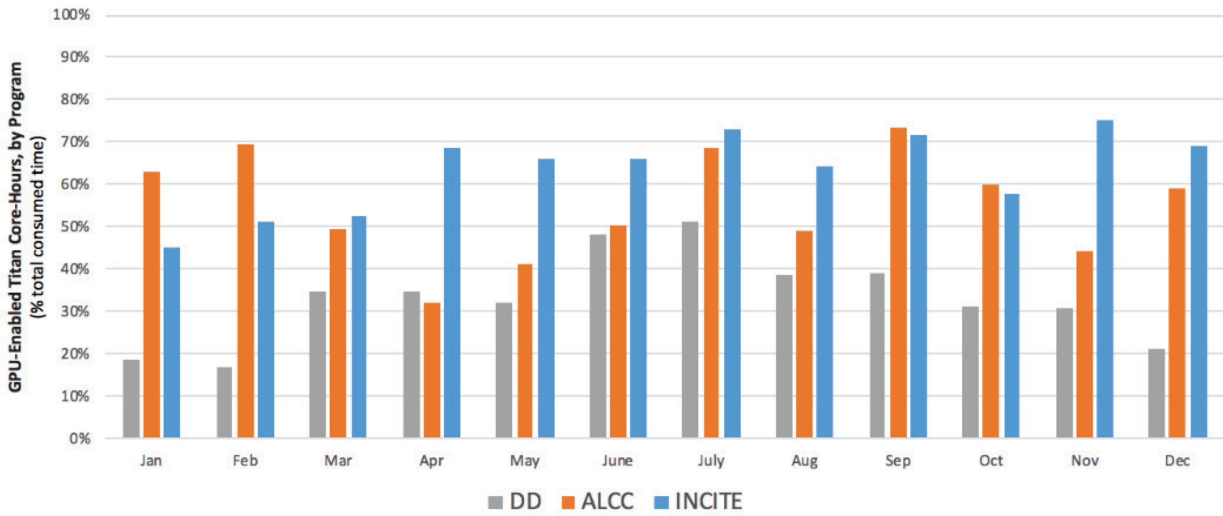

Figure 2.3. GPU-enabled percentage of compute time for the DD, ALCC, and INCITE user programs.

2.15 CENTER-WIDE OPERATIONAL HIGHLIGHTS

2.16 LUSTRE CORRUPTION ISSUE

The OLCF operations staff were able to very quickly address and engage toward a resolution of a hardware failure toward the end of August 2018. The tools developed and honed at the OLCF identified 10,279 affected files within 48 hours of the incident and users were immediately notified of the problem files. These tools allowed the OLCF to group the files into categories of severity to pass along impact assessments to the users. These categories are identified below.

Type of Impact	Number of impacted files	Recommendation/Action taken
Damaged Metadata	2,203	Can be returned to the users - placed in a directory for the users to examine whether to keep or discard the data
Data and/or objects are lost - total file loss	1,563	Paths to the files were deleted, users were given the path and filename that was deleted.
Data files were removed by automated purging	1,048	As the tools were running, the automated purger was enabled because the filesystem was nearing the capacity where performance would be impacted. These damaged files were deleted by the purger, not by system administrator intervention. Users were given the path and filename of the deleted data.
Damaged directory block for parent directory	5,465	These files are inside of a directory that is damaged and cannot be reached through the filesystem. Users were notified of the path and filename before the data was deleted.

The experience of the OLCF staff in using file system repair and inspection tools allowed a very quick response to the issue, prompt return to service of the compute platforms that rely on the file system, and directed communication to impacted users that reduced confusion on whether or not their data sets were impacted in this event. In total, the event resulted in a significant impact to the users and a 2-day degraded state for the Spider II filesystem. The file systems team within the HPC & Data Operations group performed exceptionally in identifying, communicating, and returning the file system back to the users. Previous incidents of this magnitude have taken much longer to identify and resolve at both the OLCF and other facilities. This shortened time-to-resolution is a direct result of the hard work of the OLCF file system experts.

2.17 NETAPP UPGRADE

During the 2018 operational year, the HPC Core Operations group initiated an upgrade of the storage that serves the user home areas, project areas, and software/compilers/debuggers repository. This purchase expanded the capability, capacity, and bandwidth of the storage system to support the increase in demand and diversity of network file system (NFS) storage needs. This allowed the OLCF to provide a performance tier backed by solid state drives (SSDs) for workloads that are critical to end-user science and operations of the NCCS, and a capacity tier backed by SATA drives for workloads that don't fall into the performance tier. All nodes in the cluster are connected by a redundant 40 gigabit network to allow dynamic movement of workloads to different nodes and tiers as needed without user impact. The storage system is architected using redundant high-availability pairs that support software updates and hardware maintenance without user impact. The increased storage capacity allowed increases to the user home and project area quotas and prepares the OLCF for future NFS storage needs and growth.

The performance tier is comprised of 4 nodes with approximately 440TB of SSD storage, 2 nodes will be dedicated to end-user workloads (home areas, project areas, software/compilers/debuggers), and the remaining nodes will support internal NFS workloads that are critical to the operations of the OLCF. This split is designed to prevent spikes that occur from end-user NFS workloads, which are becoming increasingly common with home areas available on the compute nodes, from impacting the critical internal applications that utilize NFS (staff home areas, LSF, Gitlab, OpenShift, etc.). The capacity tier will be a minor hardware refresh of the existing nodes, augmenting the existing SATA shelves with 2 additional storage shelves (600TB). This capacity tier will be used for lower-performance workloads such as persistent storage for containers in OpenShift, supporting development machines, application and database backups, and backups of volumes in the high performance tier.

The entire upgrade increased the OLCF's storage capacity from 440TB of SATA and 60TB of SSD to 1440TB and 660TB of SSD and increased the capability to process NFS traffic by 5x.

3. ALLOCATION OF RESOURCES

CHARGE QUESTION 3: *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. The OLCF enables compute and data projects through the Director's Discretionary user program. This program seeks to enable researchers through goals that are strategically aligned with ORNL and DOE, as described in section 3.1.

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

3.1.1 2018 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.

3.1.2 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.

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. More information about the ACCEL program and a science achievement highlight from SmartTruck is described in Section 8.3.2.

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, 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 2018, the OLCF DD program participants used 14.23% of total user resources consumed for these DD program goals, consuming 594 million Titan core hours.

Innovation

HIGH PERFORMANCE COMPUTING FACILITY
2018 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2019

4. INNOVATION

CHARGE QUESTION 4: *(a) Have innovations been implemented that have improved the facility's operations? (b) Is the facility advancing research, either intramurally or through external collaborations, that will impact 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 2018, the OLCF pursued innovative technological solutions 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 areas of operations in 2018: operational innovations, research activities for next generation systems, cross-facility integration, strategic external collaborations, and visualization approaches.

4.1 OPERATIONAL INNOVATION

4.1.1 FPROF

fprof (Large scale file system profiling tool) continues to expand its feature set and impact. In 2017, the OLCF reported the use of fprof by LLNL. LLNL's use case prompted new development involving the detection of files with "holes" in it, as LLNL's backend disk file system, ZFS, has compression turned on. Without proper detection and calibration, the overall file system size would not be entirely accurate. PNNL also made a request and provided a patch to detect other file system types such as PIPE and socket files. Cameron Harr (LLNL) and ORNL presented the new findings and lessons learned during the LUG 2018 conference, titled "High-performance Profiling of High-performance Parallel File Systems", which was received very positively by the community.

4.1.2 Life-cycle I/O (LCIO)

Evaluating file system performance over time is a non-trivial task on large scale, parallel file systems. The complexity arises due to the interplay between two variables; namely, application workload (reads and writes) and the state of the file system. Many benchmarks test how the file system performs within a given state, but to observe the change over time necessitates that the file system state ages between benchmarks. For an HPC file system, the sheer scale compounds these challenges. LCIO (Life-cycle I/O) is designed to efficiently age a large file system - the aging process must respect the standard HPC guidelines for preventing bottlenecking and serialization, but also create something resembling a realistic file system at the end of the process. LCIO can also take into account an existing file size and cardinality distribution, as generated by tools like fprof as a starting point and populate a file system toward that state.

This latter capability proved to be invaluable when the OLCF began acceptance activities for the 250 PB Summit file system. According to the Summit acceptance criteria, certain file system scaling and performance tests have to be met when the file system is in a certain state, e.g., populated with 15 billion files, to be exact. A prototype version of LCIO was leveraged to tackle this challenge. LCIO was able to use 1/10th of the Summit resource, the fprof-generated file size, and file distribution statistics as guidance that led to large file population generation in less than two days. The OLCF expects that this file system aging tool will play important roles in the future for system evaluation, acquisition, and operation. More information about LCIO, including build files and documentation can be found at: https://github.com/olcf/lcio.

4.1.3 Constellation DOIs

The Constellation DOI service, as reported in previous OLCF OAR's, continued to operate in CY 2018. The OLCF has seen an increase in the number of DOIs that were published. Among those was a new DOI for neutron spectroscopy data from a user at the Spallation Neutron Source. This is the first time the OLCF has published SNS data through the Constellation data portal and DOI service. Four other DOIs were added for datasets in electron microscopy, two of which were referenced in articles published in the journals “Nature Communications” and “Applied Physics Letters”.

4.1.4 Gordon Bell Targeted Innovations

In 2018, OLCF staff worked closely with several of the Gordon Bell submissions to aid them in tuning their applications to take advantage of new technologies on Summit. OLCF staff worked with the U. Tokyo team helping them take advantage of advanced network resources. Specifically, the use of Scalable Hierarchical Aggregation Protocol (SHARP) and GPU Direct technologies to reduce the overhead of collective and point to point communications. The use of SHARP resulted in significant speed ups by reducing expensive MPI Allreduce collectives.

Five of the six Gordon Bell finalists utilized machine learning, leading to significant interest in using the node-local NVMe devices on Summit to accelerate training. While IBM had not provided stable provisioning software for the drives, members of the OLCF developed the means necessary to reliably provision the NVMe devices for use in the machine learning workloads. These drives, due to their large capacity (1.6 TB) and high random read rate (1,000,000 IOPS), are of interest to machine learning workloads looking to train on datasets that are larger than memory.

4.1.5 GPU Age-aware Scheduling

In 2018, the OLCF completed the deployment of GPU Age-aware scheduling on Titan with the addition of CPU-only job demarcation. This technique prioritized the mapping of CPU-only jobs and small jobs to the older pool of GPUs. The original technique deployed in July of 2017 sought to increase the stability of leadership-class jobs on Titan. The technique, as previously reported in the 2017 OLCF OAR consisted of a reorganization to the default ALPs ordering scheme and an update to scheduling that prioritized known stable GPUs for leadership jobs.

The operational impact from the production environment is clear. From the period of Jan 2016 through Nov 2018, GPU Double Bit Error (DBE) based failures have shifted away from leadership-class jobs. Results, in Figure 4.1 show a month by month analysis of this impact. As shown by the trend line, in the period prior to the deployment of GPU Age-aware scheduling, failures occurred in leadership jobs over 55% of the time. However, from the period of Jan 2018 through Nov 2018, leadership jobs represented only 35% of the failures. The details of this effort were published and presented at Supercomputing in November of 2018.

Christopher J. Zimmer, Don Maxwell, Scott Atchley, Stephen McNally, Sudharshan S. Vazhkudai, "GPU Age-Aware Scheduling to Improve the Reliability of Leadership Jobs on Titan," *Proceedings of Supercomputing 2018 (SC18): 31th Int'l Conference on High Performance Computing, Networking, Storage and Analysis*, Dallas, TX, November 2018.

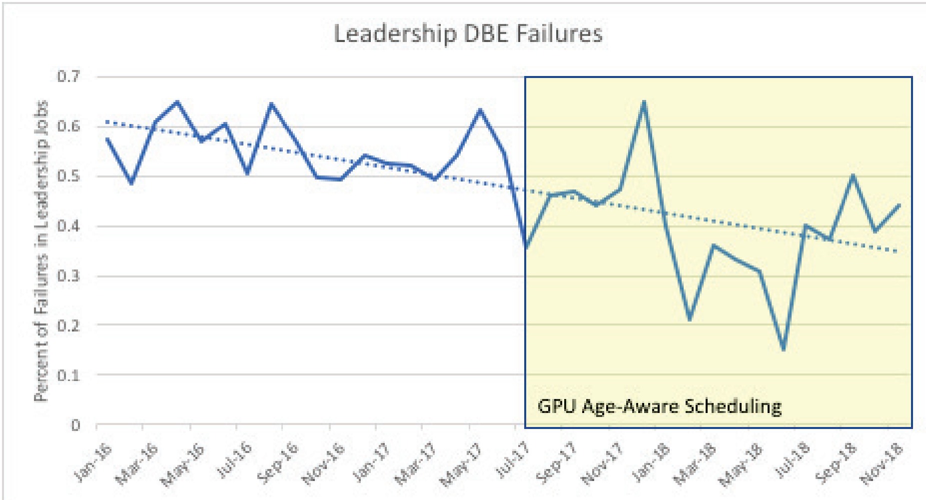

Figure 4.1. Leadership job DBE failures over time.

4.1.6 Distributed Archival Storage Request for Information (RFI)

OLCF, along with NERSC, ALCF, BNL and LLNL, conducted a market survey of distributed storage offerings via an RFI. The purpose of the RFI was to study the viability of connecting the DOE Office of Science (SC) simulation and experimental facility archives in order to cater to future cross-facility workflows and distributed data sharing. To this end, the RFI requested vendor proposals addressing several key DOE SC requirements such as distributed data access for cross-facility workflows, long-term data retention, a single unified namespace, support for federated identity management, performance (metadata rates and data ingest/egress bandwidth), redundancy, disaster recovery, availability, high capacity, efficient interfaces, smart metadata management, overall manageability/flexibility, and the need to handle diverse media types.

Eighteen vendors responded with potential solutions for a scalable storage system across geographically distributed facilities. Proposed solutions can be grouped into several categories ranging from turnkey local POSIX file systems and turnkey local object storage systems both with tape backends, cloud-hosted data and metadata, and co-development opportunities with or without a vendor. The labs collectively reviewed the proposals at ORNL, discussing the pros and cons of each solution. The exercise was extremely beneficial to understand the state of the industry in the area of distributed storage, vendor market presence, and the relevance and timeliness of the DOE SC requirements. This exercise was led by Sudharshan S. Vazhkudai, the Technology Integration group leader, for the OLCF.

4.1.7 Summit I/O Network Design

The OLCF improved the design of the network connectivity for Summit's Spectrum Scale (GPFS) file system for better performance and reliability. The design was modified from the original vendor proposed layout. The modified design improves Summit's performance by moving storage server-to-server traffic off of Summit's network and adds redundancy to the center-wide file system connections that serve supporting OLCF systems Rhea and the data transfer cluster. The original IBM design for the storage system would have met the performance needs of the file system, but it had two major drawbacks including: (1) the possibility of storage and MPI traffic congestion and (2) a lack of resiliency with center-wide connectivity.

First, in any distributed file system, such as GPFS, the storage servers occasionally need to communicate between themselves. IBM's Burst Buffer (BB) design would dramatically increase this communication when migrating data from compute node SSDs to GPFS. The design would have a BB process running on the GPFS server that would copy the data from compute nodes and then write that data into GPFS. Because GPFS stripes data over multiple servers, the vast majority of that data would go back out on the network to reach the other servers. The IBM design would have sent that data back to Summit's core network switches where it could potentially congest MPI traffic.

Second, because OLCF's users expect a single center-wide file system to serve all of OLCF's systems, it is critical that the file system remain operational and accessible even when Summit is down. The IBM design only provided a single link into the OLCF network. If that single link goes down (e.g., port failure, cut cable), then the OLCF clients would report the server down and request that operations fail over to the other server in the system. This would cause a performance degradation even if the server's other four ports on the Summit fabric were operating normally. This would then negatively impact Summit's performance as well. Adding a fourth HCA to provide a redundant connection to the OLCF fabric was not an option due to the lack of PCIe slots.

Fortunately, IBM had over-provisioned the network connectivity within the storage servers. Based this insight, the OLCF staff proposed an alternate design that removed the single-port HCA. For the remaining two dual-port HCAs, one port per HCA would connect to Summit's fabric and the other port would connect to the OLCF fabric. This configuration provides 25 GB/s per fabric, more than enough to meet Summit's needs while providing redundant connectivity for both fabrics.

This successful co-design effort between IBM and the OLCF has improved the performance and resiliency of Summit's center-wide file system.

4.1.8 ARM System Deployment, Community Software Development and Technical Evaluation

Over the last year, the OLCF has been active in the (relatively new, but growing) ARM HPC space.

In CY 2018 the OLCF deployed Wombat, a 16-node cluster from Hewlett Packard Enterprise (HPE). Since this hardware was experimental, HPE organized a collaboration consisting of the purchasers of this hardware as well as the hardware and software vendors (such as Mellanox and RedHat). This collaboration worked closely together to sort out various issues with the hardware and software deployments and operation. As of the end of 2018, there were approximately 40 users on Wombat (not counting NCCS staff accounts) from a variety of HPC specialties. For the most part, these users were using Wombat to test and evaluate applications on the ARM architecture.

In addition to the deployment of Wombat, the OLCF participated on the Technical Advisory Team for Sandia National Lab's (SNL) Astra acquisition, made several contributions to the Spack community (primarily fixing build issues with various Spack packages on the ARM architecture) and also made Wombat available as a training resource when ARM (the company) wanted to hold a workshop on SVE programming at SC'18.

Finally, the OLCF brought to bear Lustre expertise to fill one of the major deficiencies of the ARM platform for HPC users: the lack of support for a viable high-performance parallel filesystem. While the Lustre client software for ARM has existed for a few years, it was experimental and not production-quality. The OLCF contributed developer time to test the code and fix bugs that were found. The improved code was pushed upstream and merged into the main Lustre source tree. The practical upshot is that as of Lustre v2.12, ARM architecture is fully supported for both clients and servers.

4.2 RESEARCH ACTIVITIES FOR NEXT GENERATION SYSTEMS

4.2.1 Notable Data and Visualization Software Enhancements

ADIOS version 2

The ADIOS software has been redesigned with the support of the ECP project. The first production version of the new framework, ADIOS 2.3, was released at the end of 2018. The new ADIOS framework focuses on sustainability and flexibility of the framework, on the unified programming approach to in situ and file-based data processing, and on performance on the new supercomputers like Summit at OLCF. In testing with applications, ADIOS showed excellent I/O performance on Summit (e.g. 30 TB/s checkpoint writing speed in the XGC application to the local NVRAM on 2000 nodes). It was also used in the ECP Whole Device Modeling application project to couple two separate physics codes (XGC and GENE) as well as in writing checkpoints to disk and in half a dozen various in situ visualization and analysis pieces. This workflow demonstrated in Figure 4.2 the first-ever successful kinetic coupling of this kind. It showcases the unified publish-subscribe programming approach to file I/O, in situ analysis and code coupling, and the easiness of integrating multiple teams' codes into an in situ workflow and executing it on an HPC platform. Currently, existing ADIOS users are assisted to move to the new version of ADIOS, while new applications are starting out with the new version.

[1] J. Dominski; S. Ku; C.-S. Chang; J. Choi; E. Suchyta; S. Parker; S. Klasky; A. Bhattacharjee; *Physics of Plasmas* 2018, 25, DOI: 10.1063/1.5044707

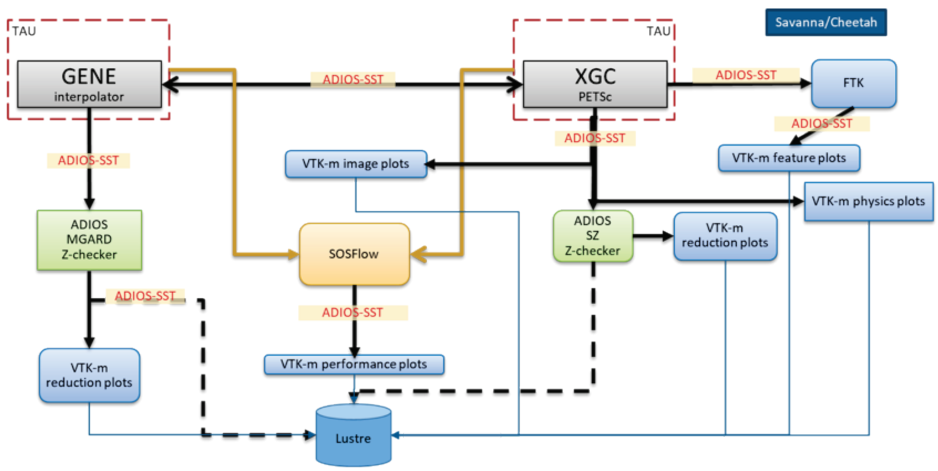

Figure 4.2. The ECP Whole Device Modeling coupled application workflow. XGC and GENE are two fusion codes, exchanging data in every computation step. Blue boxes denote in situ visualization services. Green boxes denote in situ data reduction services. Besides application data, performance monitoring information is also being collected and visualized (by TAU and SOSFlow). All data between the various applications and services are transferred by ADIOS.

Visualization work for Computational Seismology (Jeroen Tromp)

In CY 2018, staff from the OLCF provided support for reading and subsequently visualizing ADIOS output files from the SPECFEM3D_GLOBE application in VisIt. This support includes updating VisIt for ADIOS releases and improving support for seismologically specific visualizations (such as hotspots or areas of interest identification). The OLCF staff worked with Ebru Bozdag (member of the Tromp team) in performing analysis and visualization of their latest model for Global Adjoint Tomogography (M25). This work resulted in two conference submissions as well as ongoing experiments:

1. SIAM Submission: "A 3D visual tour of Earth's interior based on global adjoint tomography", Ebru Bozdag, David Pugmire, Cagda Demirkan, Sebnem Duzgun, Ergin Isleyen, Wenjie Lei, Ridvan Orsvuran, Youyi Ruan, Jeroen Tromp

2. AGU Fall 2018 Poster: "Exploring Earth's Interior in Immersive VR Environments", E. Isleyen, C. Demirkan, R. Orsvuran, E. Bozdag, S. Duzgun, D. Pugmire, W. Lei, Y. Ruan, J. Tromp https://agu2018fallmeeting-agu.ipostersessions.com/default.aspx?s=C6-E2-3A-A8-E6-67-F8-AF-58-63-5B-64-4B-38-D8-8B

Some of the recent explorations of the newly resolved features in the M25 model have led to visualization and performance analysis optimization discussion between the M25 and M15 models. Visualizations of the M25 model are below. The images in Figure 4.3 and Figure 4.4 below show the negative and positive perturbations of the Tahiti/Samoa regions, respectively. The improved resolution in the new model is giving insight into the seismological structure of the earth, and provides guidance on further work. These images provide scientists with valuable insights into the complex 3D nature of these plumes and subduction zones and their relationships to the structure of the earth's interior.

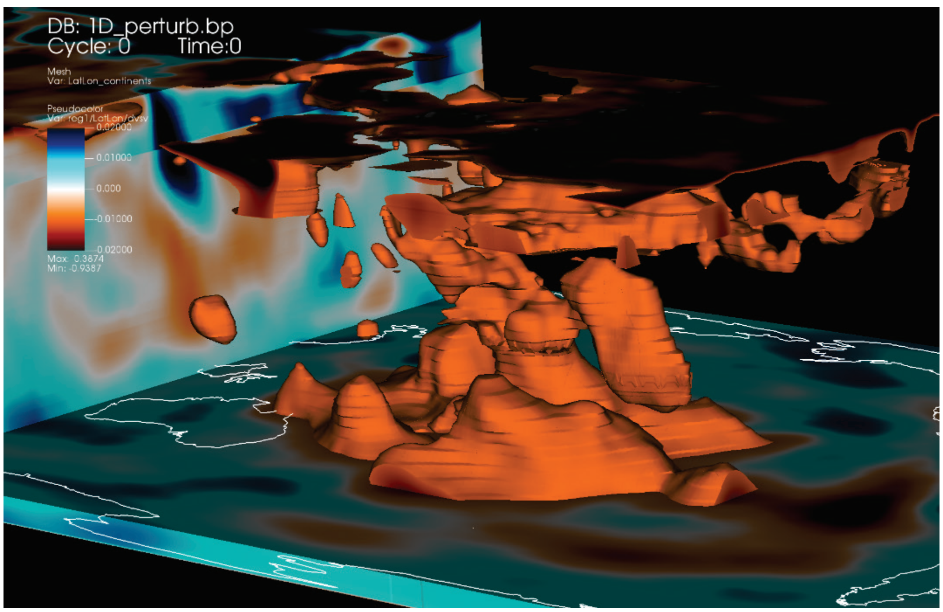

Figure 4.3. Negative perturbations of the Tahiti/Samoa regions

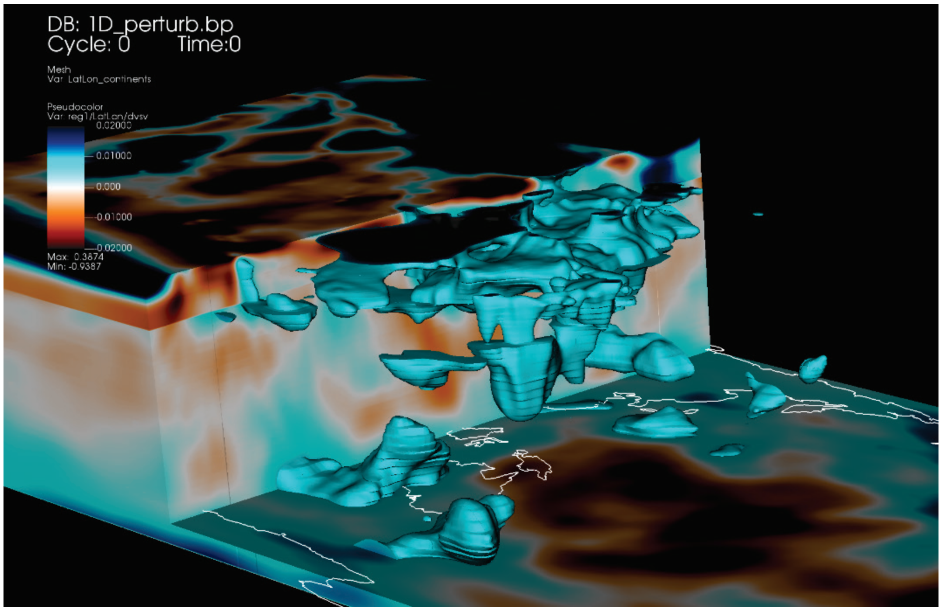

Figure 4.4. Positive perturbations of the Tahiti/Samoa regions.

The team has also VR exploration using immersive goggles. The features from the visualizations are exported and then viewed within the goggles. A video of this interaction can be found here: https://www.youtube.com/watch?v=12XJ2Nh_9IM&feature=youtu.be

NEW DISCOVERIES WITHIN 'SIGHT', An OLCF-developed visualization tool offers customization and faster rendering

Exploratory visualization can enable researchers to improve models before starting a simulation; make previously unseen connections in data that can inform modeling and simulation; and more accurately interpret computational results based on experimental data. OLCF computer scientist Benjamin Hernández incrementally developed the exploratory visualization tool SIGHT by working with OLCF users to fold in the specific features they needed for their projects. The following text highlights a successful implementation of this technology.

The computational prediction of the possibility of the direct ejection of large nanoparticles into cold water environments, above the boundary of the cavitation bubble, inspired an experimental study in the research group of Dr. Bilal Gökce and Prof. Stephan Barcikowski from the University of Duisburg-Essen, Germany, who performed a series of stroboscopic videography experiments showing the emergence of small satellite bubbles surrounding the main cavitation bubble generated in single picosecond laser pulse experiments. Carefully timed double pulse irradiation triggers expansion of secondary cavitation bubbles indicating, in accord with the simulation results, the presence of localized sites of laser energy deposition (possibly large nanoparticles) injected into liquid at the early stage of the bubble formation.

An image illustrating the results of the simulation was featured on the back cover of the April 15, 2018 Issue of Nanoscale is shown below in Figure 4.5. The image is prepared by Benjamín Hernández of the Oak Ridge Leadership Computing Facility (OLCF), who provided significant assistance with visualization of large atomic configurations using the visualization tool SIGHT he developed.

Currently the team is analyzing and visualizing simulation results of a atomistic simulation for a Ag bulk target of ~ 400 nm*400 nm*300 nm, made of ~2.8 billion atoms irradiated by a 100 fs laser pulse at 0.3 J/cm^2. The simulation will allow the team to compare the atomistic simulation results directly, for the first time, with the experimental observations on the similar length scale and reveal the processes responsible for the laser-induced modification of surface microstructure and morphology.

The visualization provides a comprehensive picture of generation of surface nanostructures in short pulse laser processing, such as the size, shape, and density of the nano-spikes as well as their internal structure like grain size distribution. The results of this study are reported in the following publications:

- C.-Y. Shih, R. Streubel, J. Heberle, A. Letzel, M. V. Shugaev, C. Wu, M. Schmidt, B. Gökce, S. Barcikowski, and L. V. Zhigilei, Two mechanisms of nanoparticle generation in picosecond laser ablation in liquids: the origin of the bimodal size distribution, Nanoscale 10, 6900-6910, 2018.

- B. Hernández, Heterogeneous Selection Algorithms for Interactive Analysis of Billion Scale Atomistic Datasets, NVIDIA GPU Technology Conference, San Jose California, 2018

- B. Hernández, C.-Y. Shih, R. Streubel, J. Heberle, A. Letzel, M. V. Shugaev, C. Wu, M. Schmidt, B. Gökce, S. Barcikowski, and L. V. Zhigilei, Inside back cover: Two mechanisms of nanoparticle generation in picosecond laser ablation in liquids: the origin of the bimodal size distribution, Nanoscale 10, DOI: 10.1039/C8NR90083C

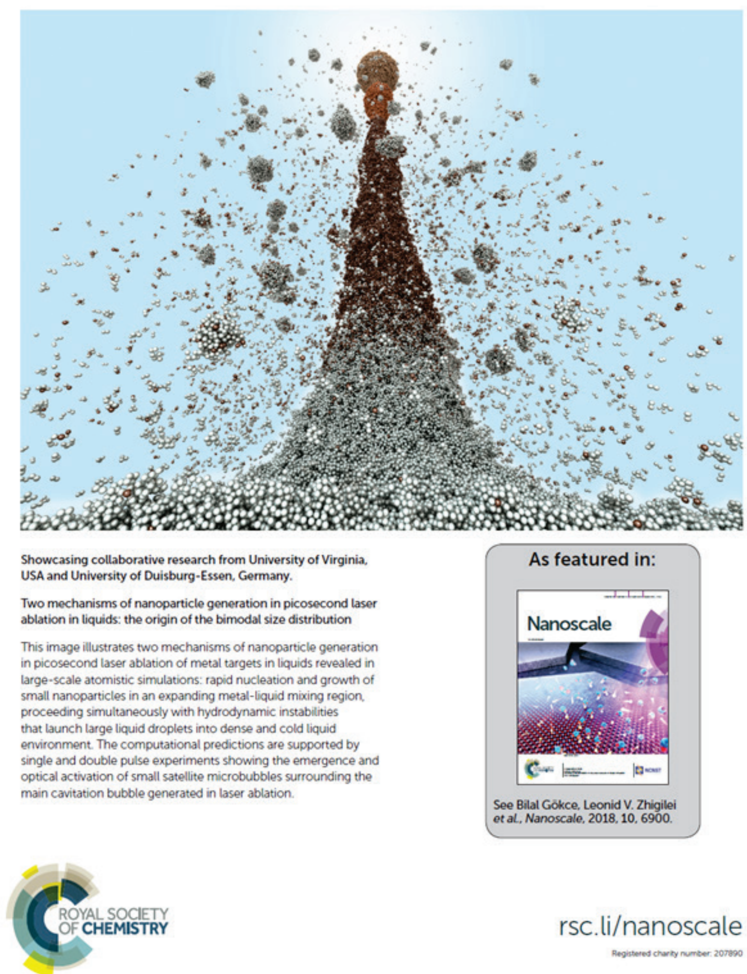

Figure 4.5. Visualization of surface nanostructures in short pulse laser processing. Back Cover of Nanoscale, April, 2018 Issue.

Augmented Reality (AR) with Summit

Under the mentorship of staff computer scientist Jamison Daniel, two interns – Cooper Colglazier (Georgia Institute of Technology) and Jessee Vomfell (The University of Tennessee) participated in the Higher Education Research Experiences program. This summer visualization project investigated how Augmented Reality (AR) could be effectively applied to monitor various aspects of the IBM AC922 Summit supercomputer at the OLCF.

AR augments existing environments with virtual data, typically via holograms overlaid onto physical structures, that enables interaction with tangible and intangible information simultaneously to obtain a more complete understanding of the data. This technology could provide a cost-efficient, portable alternative to large-scale visualization laboratories. The project resulted in a prototype monitoring tool to observe real-time temperature and GPU load on the Summit supercomputer.

The software tool leverages the Microsoft HoloLens and has two distinct modes. The machine room mode covers Summit with a holographic replica of the supercomputer to help ORNL staff identify and address maintenance issues onsite. For users unable to access the machine room directly, the overview mode allows them to remotely monitor the system by inspecting a miniature AR model of Summit and marking areas of concern for others onsite to examine.

Figure 4.6 Interactive Summit Nodes.

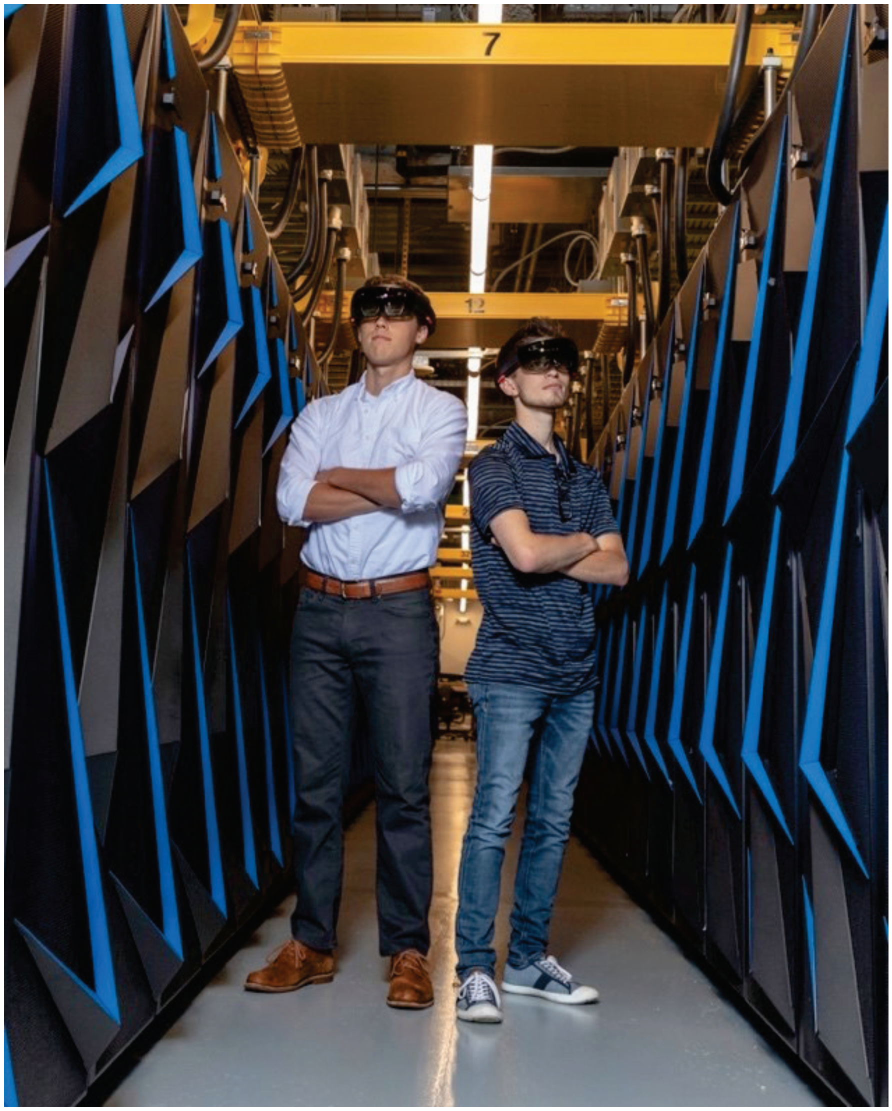

Figure 4.7. Cooper Colglazier and Jessee Vomfell demonstrate the use of AR in the Summit machine room.

Risk Management

HIGH PERFORMANCE COMPUTING FACILITY
2018 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2019

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 operational risks for the OLCF in CY 2018 are listed and described in this section. Planned mitigations and implementations are included in the subsequent descriptions. As of this writing, the OLCF has zero high-priority operational risks but a few major project risks to be tracked in CY 2019 for the OLCF-5 project.

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 (OLCF-4 and OLCF-5 during CY 2018). 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, 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 operations risks to the DOE program office as part of its monthly operations report. At the time of this writing, 38 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, 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 87 tracked risks 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 2018

Table 5.1 contains the major risks tracked for OLCF operations in 2018. The full OLCF operations risk register is available on request. The selected risks are all rated medium or high in impact.

Table 5.1. 2018 OLCF major risks

Risk ID/ description	Probability/ impact	Action	Status
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.
917: Robust support will not be available to ensure portability of restructured applications	Medium/medium	Mitigating	The OLCF deployed multiple compilers which maximizes the exposure of multiple levels of concurrency in user applications. The OLCF involvement in the standards bodies such as the OpenACC consortium continues to assist in mitigating this risk.
948: Lack of adequate facilities for an exascale system	Low/high	Accept	Plans to house the OLCF-5 project's Frontier system are in Building 5600 are awaiting the final approval in the Critical Decision 2/3 review in 2019.
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 2018. The number of open positions carried over from CY 2017 was 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. Additionally, CUDA 9.1 was deployed on Titan as a mitigation to this risk for codes to transition from Titan to Summit in CY 2019.
1142: OLCF cost increases because fewer computer room customers to distribute maintenance and operation costs among	Low/high	Mitigating	In 2018, the data center customer base remained static, and slightly fewer systems were deployed. The clean-out of portions of building 5600 Room E204 increased the available data center space. As 2019 begins, new customers and projects are anticipated. The clean-out of and facility upgrades to accommodate the OLCF-5 Frontier system, awaiting the CD2/3 approval in Q1 CY 2019, will create a data-center floor space shortage until the project is finished in room E102.

Table 5.1. 2018 OLCF major risks (continued)

Risk ID/ description	Probability/ impact	Action	Status
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: Lack of available spare parts as Titan ages	Medium/medium	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. During the operational year, the OLCF reduced the probability and impact for this risk's occurrence due to the observed failure rates and spare parts cache.

5.3 NEW OR RECHARACTERIZED RISKS SINCE LAST REVIEW

5.3.1 Recharacterized Risks

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

Risk No. 361	Scientists decline to port to heterogeneous architectures
Risk owner	Jack C. Wells
Status	Mitigating—Current
Probability	Medium → Low
Impact	***Cost:*** Medium ***Schedule:*** Low ***Scope/Tech:*** Low
Mitigations	Marked improvement of compiler directive technology from CRAY, CAPS and PGI (including OpenACC) is overcoming some technical barriers for computational scientists to port and achieve acceptable performance running on hybrid accelerated architectures.
Triggers	A decrease in the number and/or quality of proposals submitted to "headline user programs, e.g., INCITE.

Risk No. 948	Lack of Facilities for exascale system
Risk owner	James H. Rogers
Status	Accept—Current
Probability	Medium → Low
Impact	***Cost:*** High ***Schedule:*** High ***Scope/Tech:*** medium
Mitigations	As the operational year passed, the design maturity of the OLCF-5 system improved and it became apparent that this risk would not be realized unless there are significant changes to the vendor's proposal. This risk remains active until the facility enhancement scope of the OLCF-5 project is substantially complete and the delivery of the OLCF-5 machine is much closer.
Triggers	Information from the OLCF-5 project that there are power/space/cooling requirement changes coming from the vendor.

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 → Medium
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. As Titan's operational lifetime nears completion the probability of this risk impacting the OLCF's operations continues to decrease.
Triggers	Intelligence on actual or likely problems will cause this to occur.

Risk No. 1154	Lack of available spare parts causes issues as Titan ages
Risk owner	Don E. Maxwell
Status	Mitigating
Probability	High → Medium
Impact	***Cost:*** Medium → Low ***Schedule:*** Low ***Scope/Tech:*** High → Medium
Mitigations	GPU Parts that were predicted to fail in the NVIDIA model were replaced with new parts in CY 2017. Parts that had not yet failed were placed in a spare parts cache and have been used to replace actual failures. The status of this parts cache is continually evaluated and that evaluation shows that it will be sufficient to meet the operational lifetime requirements for Titan.
Triggers	Increase in application and/or GPU part failures

Risk No. 1240	Failure to handle Export Controlled Information (ECI) properly
Risk owner	Ryan Adamson
Status	Mitigating
Probability	Medium → Low
Impact	***Cost:*** Low ***Schedule:*** Low ***Scope/Tech:*** Low
Mitigations	Project Principal Investigators (PIs) and members participate in an initial project briefing where an ORNL Export Control Analyst describes the categorization of the project based on the information provided in the project's application form. OLCF Cyber Security staff contribute to this briefing and outline what storage resources are capable of handling ECI and what the requirements are of the project's team members. The project's PI is responsible for briefing this information to new members as part of the approval process to be added to the project. The OLCF has automated permission enforcement controls on resources/areas where ECI is capable of being stored.
Triggers	Information/indications that ECI has been released.

5.3.2 New Risks in This Reporting Period

The following risks were created and tracked during CY 2018. They are included with their risk creation date, mitigations, and triggers.

Risk No. 1244	Summit acceptance delays OLCF user programs start
Risk creation date	2018-09-12
Risk owner	Jack C. Wells
Status	Retired
Probability	Medium
Impact	***Cost:*** Low ***Schedule:*** Low ***Scope/Tech:*** Low ***Other:*** Low
Mitigations	There are a few mitigations that exist. First is that INCITE teams could begin working on Titan and transition over to Summit as the machine is accepted and Early Science is performed. The second is that there are currently not full promises to the ALCC and ECP programs and we could reduce the allocation to these programs to fulfill the INCITE commitment.
Triggers	Acceptance of Summit activity delayed past November 15, 2018.

5.4 RISKS RETIRED DURING THE CURRENT YEAR

Risk No. 906	Programming environment tools may be insufficient.
Risk owner	David E. Bernholdt
Status	Retired
Retirement Comment	The deployment of the final XK7 Programming Environment in March 2018 supported CUDA (9.1), creating the path for users to easily transition from Titan to Summit.
Impact	***Cost:*** Low ***Schedule:*** Low ***Scope/Tech:*** Medium
Mitigations	The Software Tools Group within the division owned this risk. They worked closely with users, developers, and standards bodies to deploy compilers, tools, and libraries to meet the requirements users have for the Programming Environment on Titan .
Triggers	The risk was triggered in April 2017 and the response of working with Cray to deploy an updated version of the PE resolved the trigger condition.

Risk No. 1146	Suboptimal resource utilization in HPSS causes lack of storage and greater expense
Risk owner	Sudharshan S. Vazhakudai
Status	Retired
Retirement Comment	This risk was not realized, the resource utilization matched the expectations when the additional bandwidth and capacity were added to the HPSS archival.
Impact	***Cost:*** Low ***Schedule:*** Low ***Scope/Tech:*** Low
Mitigations	Reserve capacity in newly deployed resource to make a shift in resource allocations should the usage not meet the design expectations.
Triggers	HPSS Storage Classes become highly utilized and there are trapped resources due to the configuration implemented to meet design expectations.

Risk No. 1193	Early Science Program for Summit terminated before completion
Risk owner	Jack C. Wells
Status	Retired
Retirement Comment	The Early Science period is being performed in conjunction with the INCITE 2019 allocation until the ALCC 2019 allocation period begins in July 2019.

Impact	***Cost:*** Low ***Schedule:*** Low ***Scope/Tech:*** Low
Mitigations	Ongoing discussions with the OLCF-4 project and DOE program sponsors to determine feasibility and need for Early Science Program's execution
Triggers	This risk was triggered and mitigated as the original Early Science Program had 3 months of execution before INCITE 2019 was set to begin.

Risk No. 1244	Summit acceptance delays user programs start
Risk owner	Jack C. Wells
Status	Retired
Retirement Comment	INCITE 2019 program began on January 1, 2019 as planned
Impact	***Cost:*** Low ***Schedule:*** Low ***Scope/Tech:*** Low
Mitigations	There are a few mitigations that exist. First is that INCITE teams could begin working on Titan and transition over to Summit as the machine is accepted and Early Science is performed. The second is that there are currently not full promises to the ALCC and ECP programs and we could reduce the allocation to these programs to fulfill the INCITE commitment
Triggers	This risk was triggered as system acceptance was later than November 15, 2018. The OLCF had planned for responses that impacted user programs but enacted a response that impacted Early Science over impacting the user programs.

5.5 MAJOR RISKS FOR NEXT YEAR

With the successful CD-2/3 review for the OLCF-5 project in February 2019, at least two new operational risks are being analyzed that will provide significant impact on the OLCF's operations in CY 2019. The first is that there is significant OLCF operational infrastructure in ORNL building 5600 Room E102 - where Titan and Spider2 (Atlas) reside. This is the targeted facility for Frontier's installation in 2021. The Facility Enhancement and Data Center Preparation scope from the project requires this space to be empty by October 2019. To meet this schedule will require operational infrastructure investments across NCCS customers - of which OLCF is one - as well as decommissioning and proper disposal of equipment in this room. There is significant schedule risk involved for both OLCF operations and the OLCF-5 project in achieving this goal.

Secondly, the OLCF-5 CD-1/3A review in 2018 approved work to begin building out the transformer room for Frontier which consumes almost all of the available office space on the 1st floor of building 5600. Moving the personnel that support the OLCF's resources could cause problems in responding to issues and/or system outages while the moves are occurring.

Finally, with the tight schedule and increased personnel in the area for demolition, construction, and moving of equipment and offices, there is an increased probability of a safety incident that could delay work in any of these projects.

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

The following risks were encountered and effectively mitigated in 2018. A short summary of the status and impact of the risk on the operations of the OLCF is included.

Risk No. 386	File system data integrity failure
Risk owner	Sudharshan S. Vazhkudai
Status	Mitigate
Probability	Low
Impact	***Cost:*** Low ***Schedule:*** Low ***Scope/Tech:*** Low
Mitigations	The OLCF will work closely with others in the Lustre community via OpenSFS to reduce the probability of data corruption via improved resiliency mechanisms. The OLCF will work on improved detection of data corruption once occurred and develop tools to quickly identify data within the file system that could be impacted by a component failure. The OLCF was able to very quickly detect a failure on 08/29/2018. The tools developed and honed at the OLCF identified the affected files in 48 hours and users were immediately notified of the problem files.
Triggers	Intelligence on failures from testing, monitoring systems indicate a critical failure, or user feedback.

Risk No. 1154	Decreasing availability of spare parts as Titan ages
Risk owner	Don E. Maxwell
Status	Mitigate
Probability	Medium
Impact	***Cost:*** Low ***Schedule:*** Low ***Scope/Tech:*** Medium
Mitigations	The OLCF will stockpile more spare parts and monitor failure rates to stay ahead of issues. Additionally, jobs that require GPUs and impact the capability performance metric are organized to run on the area of the machine with the GPUs that were replaced in CY 2016 and 2017 such that they are less likely to be subjected to a GPU failure.
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.

Environment Safety and Health

HIGH PERFORMANCE COMPUTING FACILITY
2018 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2019

6. ENVIRONMENT SAFETY AND HEALTH

CHARGE QUESTION 6: *(a) 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? (b) Has the facility implemented appropriate Environment, Safety, and Health measures?*

OLCF RESPONSE: Yes.

ORNL is committed to operating under the DOE safety regulations specified in 10 CFR 851, which 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. These safety requirements are incorporated into ORNL contracts 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 line managers, qualified safety and health professionals, 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 having 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 field implementation and documentation of all operational safety reviews, and it partners with the ORNL Offices of Institutional Planning, 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 seeks 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 OLCF Federal Project Director, who solicited the following feedback from the site office staff in the Operations and Oversight Division about OLCF's safety culture.

- Operations of the OLCF in the NCCS remained safe, efficient, and effective as there were zero total recordable cases and DARTs in FY 2018.

The following activities are ES&H highlights from CY 2018:

- The OLCF-5 project Health and Safety Plan and the Hazard Analysis were reviewed and revised in 2018.

- SUMMIT Installation and acceptance were completed in 2018. The project was completed safely with zero total recordable cases and DARTs. However, the project did sustain one first aid involving a small laceration and one electrical event caused by faulty equipment. The first aid case was a result of inappropriate use of a utility knife. The electrical event investigation concluded that the event was due to a fault with both the design and assembly of a plug and cord assembly. The subcontractor replaced all cords (previously installed or yet-to-be installed – approximately 10,000 cords) with uniquely-identified replacement cords manufactured with a corrected design, a revised assembly process 100% x-ray inspection for precursors to the failure, and 100% functional testing.

- The K100 data center, that houses Summit was transitioned from a construction area to a laboratory space and data center. A research safety summary (RSS) was drafted, reviewed, and approved to control activities in the space. In addition, participants were required to review the RSS and the area was posted in accordance with applicable ORNL requirements.

- Occupational exposure monitoring was conducted for noise in all three OLCF data centers. The results of these surveys indicate that occupational exposure to noise above the ACGIH TLV is highly unlikely. Therefore, the requirement for employees to be in the ORNL Hearing Conservation Program was removed. Communication between center management and staff informed the staff of the deletion of the requirement but allowed the employee to volunteer to participate in the hearing conservation program. However, the OLCF decided to still require hearing protection for room access as a best management approach.

- Occupational exposure monitoring was conducted in the Summit Data Center (K100) for electromagnetic fields around the electrical distribution panels in the space. Results of the survey indicated that there are no concerns for electromagnetic fields and interference with medical implants.

- The Authorized Access to ORNL Computing Centers access training was revised. Review of the existing module indicated that it had not been revised for several years. The training was updated to reflect new systems, personnel, and access requirements.

- The Computational Sciences Building Computer Center Operations Emergency Response Plan was updated to include emergency response specification for the K100 data center. In addition, the data center manager conducted briefings for all operation center operators, tier one support, and utility support personnel.

- Strides were made in the E102 and E204 data centers to eliminate excess materials and outdated systems.

- In January of 2018, the OLCF hired the OLCF-4 project installation manager as a new data center manager (DCM) to manage the three OLCF data centers. The new data center manager has over

25 years of experience as a health and safety professional within DOE construction, demolition, and operations projects. The new DCM is responsible for managing the data center day to day operations like a normal DCM but is also responsible for monitoring all work for compliance with safety standards and proper work control. In addition, the DCM is also responsible for the orientation and mentoring of subcontractor personnel to the UT-Battelle safety culture and expectations. Hiring a DCM with significant experience as a health and safety professional has brought a non-traditional viewpoint into the OLCF, which ultimately strengthened and positioned the OLCF to continue safe and efficient operations.

Cyber Security

HIGH PERFORMANCE COMPUTING FACILITY
2018 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2019

7. CYBER SECURITY

CHARGE QUESTION 7: *(a) Does the facility exhibit a culture of continual improvement in cyber security practices? (b) 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 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). 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.

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*[*] 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. 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*[†] *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.), 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.

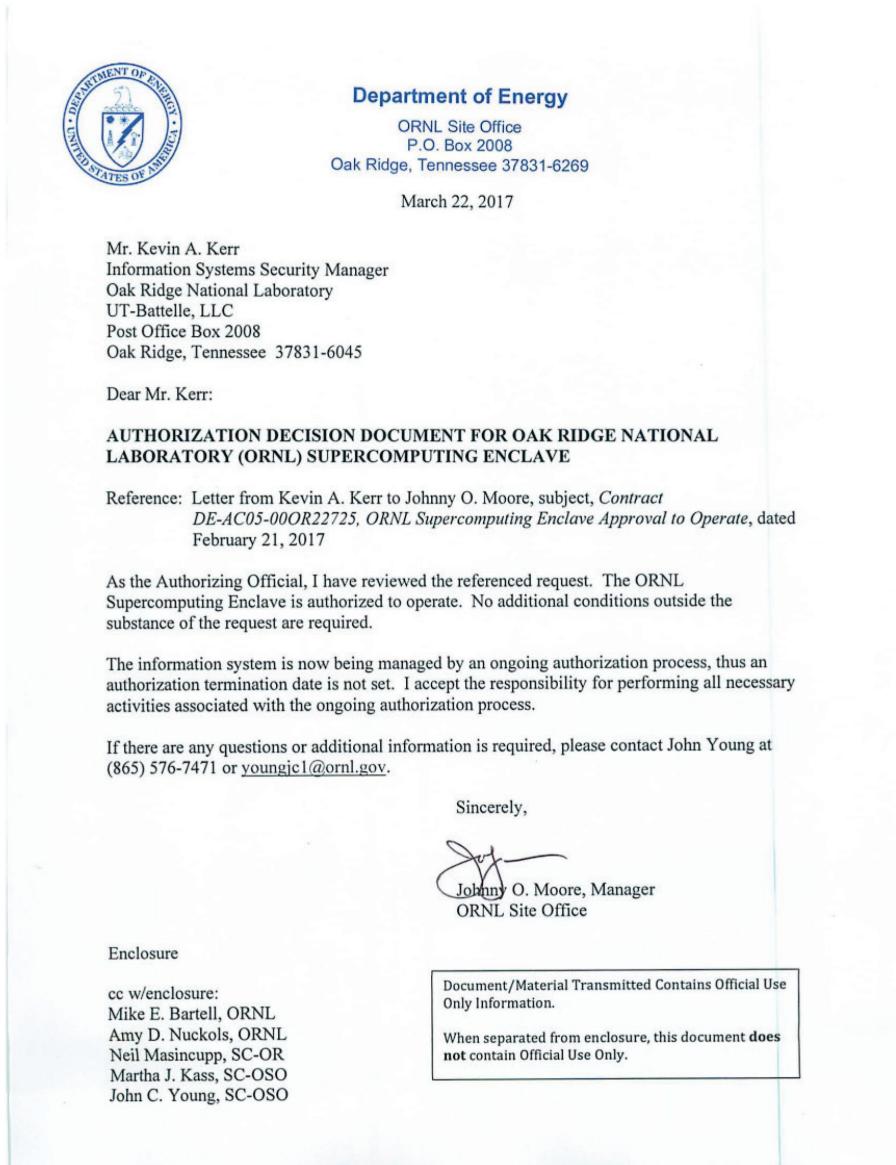

Department of Energy
ORNL Site Office
P.O. Box 2008
Oak Ridge, Tennessee 37831-6269

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.

7.2 OSQUERY/FLEET

OLCF is modernizing the Host based Intrusion Detection (HIDS) security tooling, Open Source HIDS SECurity (OSSEC), with a more modern and functional set of tools called osquery and Fleet. osquery and Fleet allow real time querying of hosts across OLCF infrastructure to provide unrivaled situational awareness. Further, it allows OSSEC-like functionality by logging historic forensic information to Splunk. osquery has use cases for security and system administration alike.

osquery and Fleet are designed to be:

- Resilient to a variety of failure modes: leverages OLCF's OpenShift, which provides redundancy and the ability to rapidly redeploy to address stability issues.

- Highly reproducible: re-provisioning a client involves configuration management software such as Puppet or BCFG2. Re-provisioning the server component can be easily performed using YAML config files to create new OpenShift containers/Pods.

- Sustainable: patching, for example, utilizes a docker build process which is integrated into the OpenShift Platform.

- Integration with existing OLCF services: OpenShift, Puppet, BCFG2, and Splunk.

- Optimized: Utilization of OpenShift allows rapid elastic scalability (can scale up and down rapidly).

7.3 SITU

Situ is an application developed by the ORNL Cyber Research group that uses statistical analysis of Netflow data from network devices to determine anomalous activity between internal OLCF devices and external devices. Netflow data from OLCF devices is collected on a central server, and then Situ ingests that data to create statistical models. Those models analyze traffic over a period of time and the statistical anomalies are quickly ascertained. The OLCF security team integrated information feeds from Situ into the Splunk SIEM system, making it simple to investigate anomalous events via a Splunk dashboard interface and correlate with other already existing security-related information sources. Further investigation using Splunk assists in identifying harmful traffic. Situ was deployed using the OLCF OpenShift platform for sustainable management of the service.

7.4 SUMMIT PENETRATION TESTING

The OLCF Cyber security team employed cutting-edge penetration testing techniques to discover several significant privilege escalation and denial of service vulnerabilities in the design and software of the Summit supercomputer and Alpine filesystem. Vulnerabilities were disclosed securely to IBM's PSIRT (Product Security Incident Response Team) where all vulnerabilities were quickly patched and verified as fixed. Information about these vulnerabilities were also shared with other DOE laboratories using systems based on IBM's Power 9 architecture for security awareness purposes. Specific Common Vulnerabilities and Exposures (CVEs) for discovered vulnerabilities include:

- CVE-2018-1776– https://www-01.ibm.com/support/docview.wss?uid=ibm10733555
- CVE-2018-1724– https://www-01.ibm.com/support/docview.wss?uid=ibm10734767
- CVE-2018-1703– https://www-01.ibm.com/support/docview.wss?uid=ibm10734239
- CVE-2018-1703– https://www-01.ibm.com/support/docview.wss?uid=ibm10734239
- CVE-2018-1782– https://www-01.ibm.com/support/docview.wss?uid=ibm10730967

Strategic Results

HIGH PERFORMANCE COMPUTING FACILITY
2018 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2019

8. STRATEGIC RESULTS

8.1 SCIENCE OUTPUT

CHARGE QUESTION 8: (a) *Are the methods and processes for monitoring scientific accomplishments effective? (b) Has the Facility demonstrated effective engagements with strategic stakeholders (i.e., beyond the user population)? (c) Is the Facility operating in a manner that enables delivery of facility mission and Department of Energy mission including maintaining a vibrant US effort in science and engineering?*

OLCF RESPONSE: Yes. OLCF projects and user programs are advancing DOE's mission to maintain a vibrant US effort in science and engineering and ASCR strategic goals via science accomplishments, as shown in section 8.2. The OLCF appropriately engages stakeholders through programmatic and industry partnerships, highlighted in section 8.3.

2018 Operational Assessment Guidance

The Facility collects 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.

8.1.1 OLCF Publications Report

In 2018, 469 publications resulting from the use of OLCF resources were published, based on the data collection completed on April 9, 2019, representing a 3.7% increase in publications over the previous year. In this document, "year" refers to the calendar year unless it carries the prefix "FY", indicating the fiscal year. 452 publications were reported in the 2017 OLCF OAR. A list of 2014–2018 publications is available on the OLCF website. https://www.olcf.ornl.gov/leadership-science/publications/

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 8.1 provides the updated, verified, and validated publications count for the 2012–2018 period, showing continued growth in both the total publications count and the number of publications in journals with high impact factors.

Table 8.1. Summary of unique OLCF publications for 2012–2018

Year	Unique, confirmed OLCF publications	High-impact publications with JIF* >10
2018	469	20
2017	467	27
2016	458	33
2015	352	21
2014	296	16
2013	359	9
2012	334	20

*JIF = Journal impact factor

8.2 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 2018, 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 2018, as shown in Table 8.2.

Altogether in 2018, OLCF users published 56 papers in journals with a journal impact factor (JIF) of greater than 7 and 20 papers with a JIF greater than 10. Also noteworthy is the fact that OLCF staff and users published in 216 independent journals or conference proceedings, which is characteristic of the wide breadth of research topics enabled on leadership resources at OLCF.

Table 8.2. Publications in high-impact journals in 2018

Journal	Number of publications
Nature	3
Science	3
Nature Chemistry	1
Advanced Energy Materials	1
ACS Nano	2
Physical Review X	1
Nature Communications	7
ACS Catalysis	2
Proceedings of the National Academy of Sciences	3
Chemistry of Materials	2
Journal of Physical Chemistry Letters	1
Green Chemistry	2
Journal of Materials Chemistry A	1
Physical Review Letters	21
ACS Applied Materials & Interfaces	2
Nanoscale	3
Bioinformatics	1

Titan Helps Scientists Fine-Tune Laser Interactions to Advance Cancer Treatments

PI: Michael Bussmann, Helmholtz-Zentrum Dresden-Rossendorf, INCITE
Allocation Program: INCITE

The Science

Until recently, most radiation for cancer therapy used photon beams in the form of x-rays to kill cancer cells, but scientists have learned these rays wreak havoc on healthy cells, too. A team studying a type of ion acceleration driven by high-intensity lasers recently performed a simulation of a novel laser target that describes the physics behind the acceleration and shows significant agreement with experiments. In the simulation, the target was expanded by the laser pulse's early start, allowing the laser to interact with more of the electrons than it would have without the expansion. The resulting proton beam consisted of ions with similar energies, which is crucial in radiation therapy to ensure that all ions will penetrate tissue to the same depth.

The Impact

More than ever, scientists are seeking to leverage heavier particles such as protons and ions because these particles can reach deep tumors and reduce the amount of radiation exposure to healthy tissues. The largest barrier to using ion beams is the size of the particle accelerators required to get these ions up to speed, but laser-driven ion acceleration holds promise for making accelerators compact enough to fit into medical environments. Simulations can lead scientists to an understanding of how to optimize the laser-driven ion acceleration process by altering the initial laser and target conditions.

Summary

A team used the PIConGPU (particle in cell on graphics processing units) code on Titan to study the use of ion acceleration via high-intensity lasers because laser-driven particle accelerators don't require the large magnets used by magnetic ion accelerators. The team performed a simulation of a novel laser target that not only describes the physics behind experiments performed at the GSI Helmholtz Centre for Heavy Ion Research but shows significant agreement with the experimental results. The use of HPC resources such as Titan—and the Adaptive I/O system (ADIOS)—can allow scientists to vary the experimental parameters in simulations and give them a better understanding of the physics behind laser-driven ion acceleration. Because laser shape, intensity, and length as well as target surface texture, density, and size contribute to the physics, researchers must study comprehensive ensembles of simulations to gain deep scientific understanding of this process.

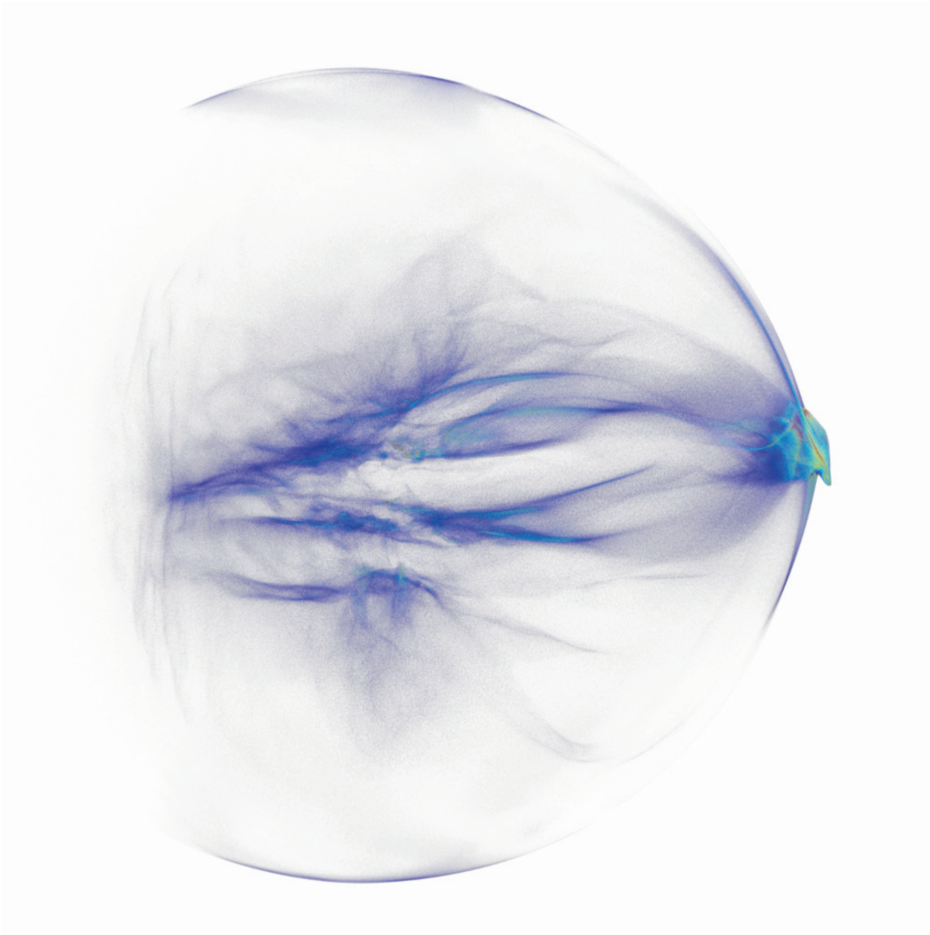

Figure 8.1. A 3D visualization of the HZDR team's final simulation of their expanded plastic target. The protons (blue) can be seen traveling along the laser axis from left to right (laser not shown). A particle bunch (red) of high-density protons can be seen at the right of the image. *Image Credit: Axel Huebl, HZDR (Simulation); Peter Hilz, LMU (Experiment); Michael Matheson, ORNL (Visualization)*

Funding

This research used resources of the Oak Ridge Leadership Computing Facility located in the Oak Ridge National Laboratory, which is supported by the Office of Science of the Department of Energy under contract DE-AC05-00OR22725.

Publications

P. Hilz, T. M. Ostermayr, A. Huebl, V. Bagnoud, B. Borm, M. Bussmann, M. Gallei, et al., "Isolated Proton Bunch Acceleration by a Petawatt Laser Pulse." *Nature Communications* 9, no. 423 (2018), doi:10.1038/s41467-017-02663-1.

Related Links

" Titan Helps Scientists Fine-Tune Laser Interactions to Advance Cancer Treatments," OLCF News (July 17, 2018)

Titan Takes Fluoride from Taps and Toothpaste to Batteries

PI: Thomas Miller, California Institute of Technology Allocation Program: INCITE

The Science

A collaboration of researchers recently discovered a new liquid electrolyte material that conducts fluoride in fluoride-based rechargeable batteries, which pack a major energy punch compared to lithium-ion batteries. As part of the project, a team led by the California Institute of Technology's (Caltech's) Thomas Miller used the 27-petaflop Cray XK7 Titan supercomputer at the OLCF to understand and refine the electrolyte's properties and confirm its unprecedented ability to conduct fluoride ions and retain chemical stability at room temperature, making the breakthrough material the first of its kind in the battery world.

The Impact

More energy-dense fluoride batteries could hold up to 8 times more charge than lithium batteries. Simulations such as the ones performed by Miller's team can help scientists understand battery interactions and improve existing and newly discovered electrolytes. The understanding of the new electrolyte, called BTFE, gives researchers a map for the mechanisms involved in stabilizing fluoride batteries and could aid in the development of new kinds of batteries with applications in cars, cell phones, or other electronic devices.

Summary

After a team at the Jet Propulsion Laboratory (JPL) discovered a new liquid electrolyte named BTFE, a team at Caltech used the Large-Scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) code to conduct molecular dynamics simulations on Titan to look at BTFE's reactions with fluoride. The new solvent worked and made the fluoride ions stable enough to shuttle from the battery's anode to the cathode, but the JPL team didn't quite know why before the simulations. Miller's group found that positively charged regions in BTFE were strongly interacting with the negatively charged fluoride to dissolve it. His team simulated many different solvents and positively charged ions—including a solvent that could potentially expand the voltage window of BTFE and make it even more stable. Follow-up studies on the project can aid in the optimization of new liquid electrolytes such as BTFE, enabling the possibility of fluoride-based batteries.

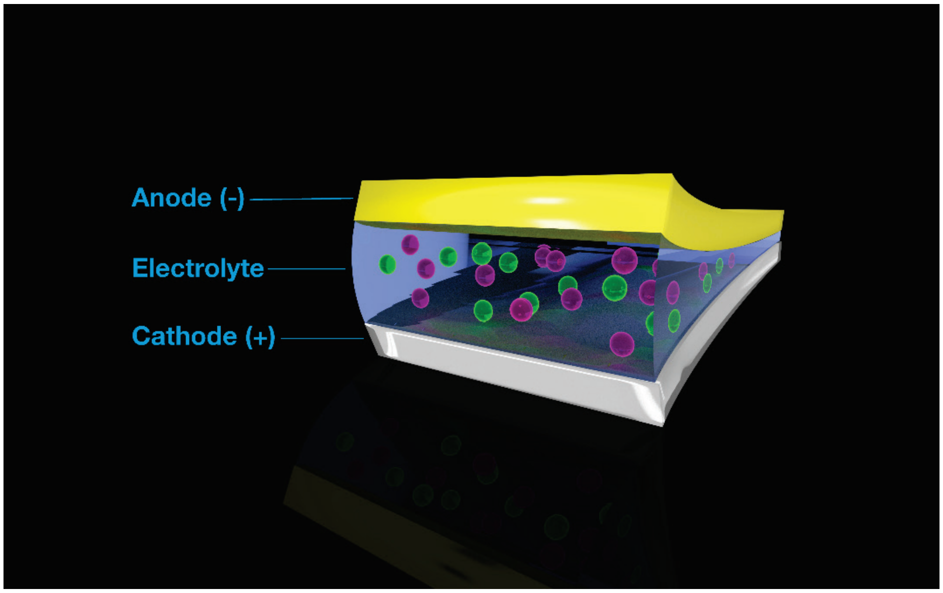

Figure 8.2. An illustration of a battery. In order for batteries to generate electricity, charged atoms, called ions (pink and green), travel between a negative node (anode) and a positive node (cathode) with the help of a liquid electrolyte solution. *Image Credit: Brett Savoie, Purdue University*

Funding

This research used computational resources from the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy (DOE) under contract DE-AC05-00OR22725.

Publications

Davis, V. K., C. M. Bates, K. Omichi, B. M. Savoie, N. Momčilović, Q. Xu, W. J. Wolf, M. A. Webb, K. J. Billings, N. H. Chou, S. Alayoglu, R. K. McKenney, I. M. Darolles, N. G. Nair, A. Hightower, D. Rosenberg, M. Ahmed, C. J. Brooks, T. F. Miller III, R. H. Grubbs, and S. C. Jones, "Room Temperature Cycling of Metal Fluoride Electrodes: Liquid Electrolytes for High-Energy Fluoride Ion Cells." *Science* 362, no. 6419 (2018): 1144–1148, doi: 10.1126/science.aat7070.

Related Links

" Titan Takes Fluoride from Taps and Toothpaste to Batteries," OLCF News (December 21, 2018)

Titan Powers Data Dive into Tropical Soil Microbes

PI: Chongle Pan, Oak Ridge National Laboratory Allocation Program: ALCC

The Science

In tropical forests, microbes manage to subsist on poor soil that oftentimes lacks sufficient levels of phosphorous, an essential nutrient for survival. To find out how, researchers carried out a detailed investigation of soil microbial genetics using high-performance computing and advanced data analysis. The team sifted through massive amounts of genetic data to analyze how the availability of phosphorous affects microbes' foraging strategies, finding that microbes in deficient soils possess significantly more genes dedicated to phosphorous extraction.

The Impact

Microbes play a major role in global nutrient cycles. Improved understanding of the interplay between microbes and soil could help scientists model processes that affect Earth's climate. For example, scientists could use the soil-microbe interplay to study the carbon cycle. In this cycle, microbes exert influence by breaking down organic carbon and releasing carbon dioxide into the air. The team's approach could be applied to other ecosystems to study nutrient limitations and inform agriculture and terrestrial biosphere modeling.

Summary

Using Titan and Rhea, the team carried out analysis techniques that allowed for microbial DNA fragments to be sorted into complete genomes and for the identification of more than 7,000 proteins in collected soil samples. This metagenomics and metaproteomics data enabled the comparison of genes and proteins from microbes in phosphorous-rich and phosphorous-deficient soils. The team discovered several key differences between the two groups. Among microbes in the phosphorous-deficient soil, the team found more than four times as many genes responsible for encoding enzymes—molecules that accelerate chemical reactions—dedicated to acquiring phosphorous from the environment. The team also identified more than 100 genes among the phosphorous-constrained microbes that work to pull phosphorus from phytate, an organic compound in plant tissue that microbes cannot easily access. Conversely, in the phosphorous-rich environment, the team found that microbes devoted more genes to acquiring other nutrients such as carbon and nitrogen.

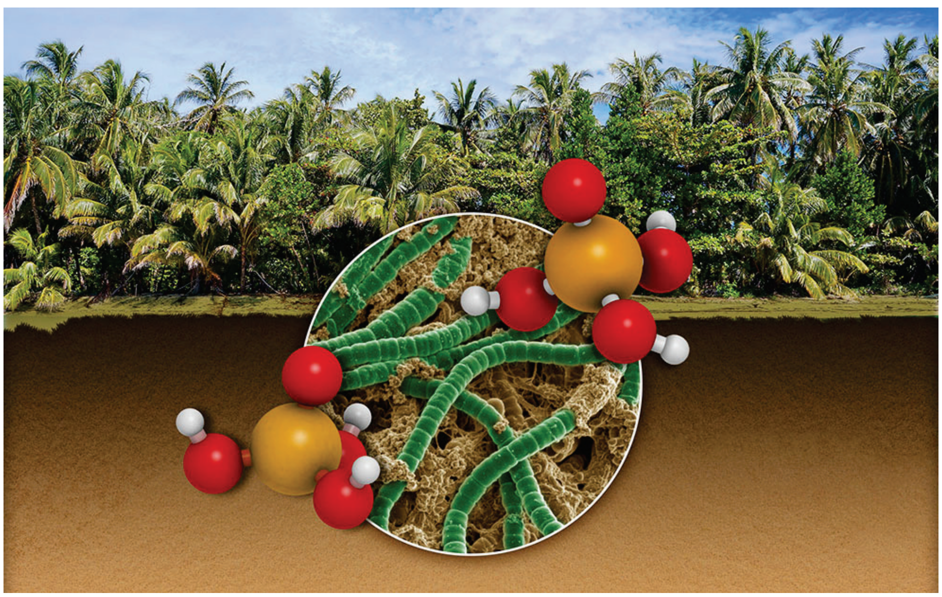

Figure 8.3. Scientists are studying how microbes in soil use nutrients like phosphorus at the molecular level, helping better model efficient land use and terrestrial processes. Image courtesy of Environmental Molecular Sciences Laboratory

Funding

This work was supported by the Department of Energy's (DOE's) Office of Science, Office of Biological and Environmental Research, including support by the Environmental Molecular Sciences Laboratory and the Joint Genome Institute, both DOE Office of Science user facilities. Funding from Oak Ridge National Laboratory's Laboratory Directed Research and Development program also supported this

work. Computing resources were provided by Oak Ridge Leadership Class Facility under the ALCC (ASCR Leadership Computing Challenge) program.

Publications

Q. Yao, Z. Li, Y. Song, S.J. Wright, X. Guo, S.G. Tringe, M.M. Tfaily, L. Paša-Tolić, T.C. Hazen, B.L. Turner, M.A. Mayes, and C. Pan, "Community proteogenomics reveals the systemic impact of phosphorus availability on microbial functions in tropical soil." Nature Ecology & Evolution 2 (2018), 499 (2018). DOI: 10.1038/s41559-017-0463-5.

Related Links

Oak Ridge Leadership Computing Facility highlight: https://www.olcf.ornl.gov/2018/02/27/titan-powers-data-dive-into-tropical-soil-microbes/ (February 27, 2018)

Uncovering the Genetics of Disease at Exascale Speeds

PI: Dan Jacobson, Oak Ridge National Laboratory Allocation Program: Early Science

The Science

Genes, DNA segments that encode for proteins and determine everything from eye color to height, also play a role in more complex traits such as addiction and disease. Deciphering which genes contribute to which traits, however, can be perilously complicated. Nevertheless, an ORNL team is poised to lead a revolution in our understanding of genes' role in disease—such as opioid addiction—by leveraging population-scale genomic datasets, state-of-the-art computing, and algorithmic improvements to uncover hidden networks of genes at record-breaking speeds.

The Impact

By identifying previously unknown relationships between genes, the team's work could help inform treatment for patients predisposed to substance abuse. For example, a quick genetic test at the doctor's office could one day indicate who should be prescribed alternative forms of pain management. The same research could also guide treatment for other conditions, including diabetes, heart disease, Alzheimer's, and dementia, and could improve practices to prevent disease from occurring in the first place.

Summary

As scientists have learned more about how genetic interactions, or epistasis, give rise to observable traits, their computing needs have grown exponentially. For each genetic variation within a population, a million billion correlations must be tested. When paired with a well-crafted application, pre-exascale supercomputers like Summit make these problems feasible. Thanks to some clever coding that exploited Summit's mixed-precision capabilities, Jacobson's team broke the exaflop, achieving a peak throughput of 2.36 mixed-precision exaops with their comparative genomics application Combinatorial Metrics (CoMet). The performance mark—the fastest science application ever recorded—represents a science output four to five orders of magnitude beyond the current state of the art. Through a partnership between ORNL and the US Department of Veterans Affairs (VA), the team will use CoMet to analyze approximately 600,000 genomes—one of the largest human genome datasets in the world—compiled under the VA's Million Veteran Program, a voluntary research program focused on studying how genes affect health. As epistatic networks are identified within population datasets, they can then be tested against known phenotypes, including diseases.

Funding

This research was supported by the Center for Bioenergy Innovation and the Plant-Microbe Interfaces Science Focus Area—both supported by the Office of Biological and Environmental Research in the DOE Office of Science. Computing resources were provided by the Oak Ridge Leadership Class Facility under the Summit Early Science program.

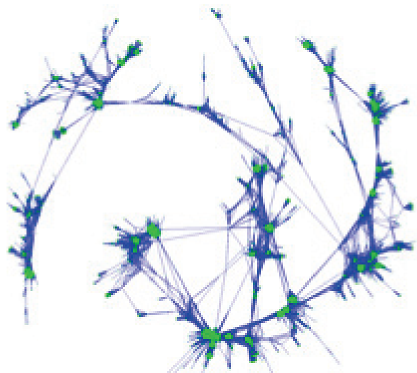

Figure 8.4. A visualization of a network depicting correlations between genes in a population. These correlations can be used to identify genetic markers linked to complex observable traits. Image courtesy of Dan Jacobson, ORNL

Publication

Joubert, Wayne, Deborah Weighill, David Kainer, Sharlee Climer, Amy Justice, Kjiersten Fagnan, and Daniel Jacobson. "Attacking the opioid epidemic: determining the epistatic and pleiotropic genetic architectures for chronic pain and opioid addiction." In Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis, p. 57. IEEE Press, 2018.

Related Links

Oak Ridge Leadership Computing Facility highlights: https://www.olcf.ornl.gov/2018/10/15/computing-genes-to-support-living-clean/ (October 15, 2018)

https://www.olcf.ornl.gov/2018/10/09/mixed-precision-a-strategy-for-new-science-opportunities/

https://www.olcf.ornl.gov/2018/06/08/genomics-code-exceeds-exaops-on-summit-supercomputer/

Identifying Extreme Weather Patterns

PI: Prabhat, Lawrence Berkeley National Laboratory Allocation Program: Early Science

The Science

As the Earth's climate warms, weather patterns have continued to deviate from their historical norms. Machine learning, a form of artificial intelligence, holds the potential to enable scientists to extract valuable insights from large, complex datasets, including climate data. Using a climate dataset on the Summit supercomputer, a team trained a deep neural network to identify extreme weather patterns from high-resolution climate simulations at unprecedented scale.

The Impact

Extracting pixel-level classifications of extreme weather patterns could aid in the prediction of how extreme weather events are changing. Though the team applied its work to climate science, many of its innovations, such as pattern-recognition algorithms, high-speed parallel data staging, an optimized data ingestion pipeline, and multichannel segmentation, lay the groundwork for future exascale deep-learning applications.

Summary

While machine-learning techniques have been successful in a variety of commercial applications, data limitations of existing tools have limited their use for scientific computing. By tapping into Summit's mixed-precision capabilities, the researchers achieved a peak performance of 1.13 exaops and a sustained performance of 0.999 exaops—the fastest deep-learning algorithm reported to date. The team's deep learning architecture is the first to be able to solve segmentation problems in climate science. Traditional image segmentation tasks work on three-channel red/blue/green images. But scientific datasets often comprise many channels; in climate, for example, these can include temperature, wind speeds, pressure values, and humidity. Running the optimized neural network on Summit allows the team to use all 16 channels and dramatically improve model accuracy.

Funding

This research used resources of the National Energy Research Scientific Computing Center (NERSC) and was supported by a grant from the Swiss National Supercomputing Centre. Computing resources were provided by the Oak Ridge Leadership Class Facility under the Summit Early Science program.

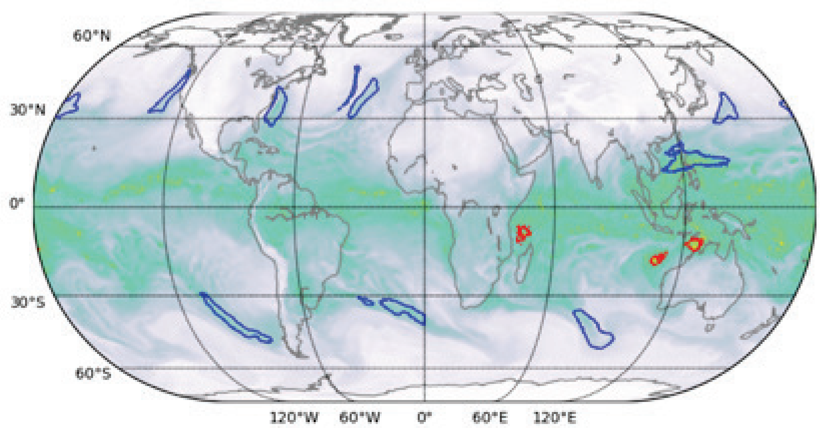

Figure 8.5. High-quality segmentation results produced by deep learning on climate datasets. Image credit: NERSC

Publication

Kurth, Thorsten, Sean Treichler, Joshua Romero, Mayur Mudigonda, Nathan Luehr, Everett Phillips, Ankur Mahesh et al. "Exascale deep learning for climate analytics." In Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis, p. 51. IEEE Press, 2018.

Related Links

NERSC highlight: http://cs.lbl.gov/news-media/news/2018/berkeley-lab-oak-ridge-nvidia-team-breaks-exaop-barrier-with-deep-learning-application/

8.2.1 INCITE 2018 Allocation/Utilization

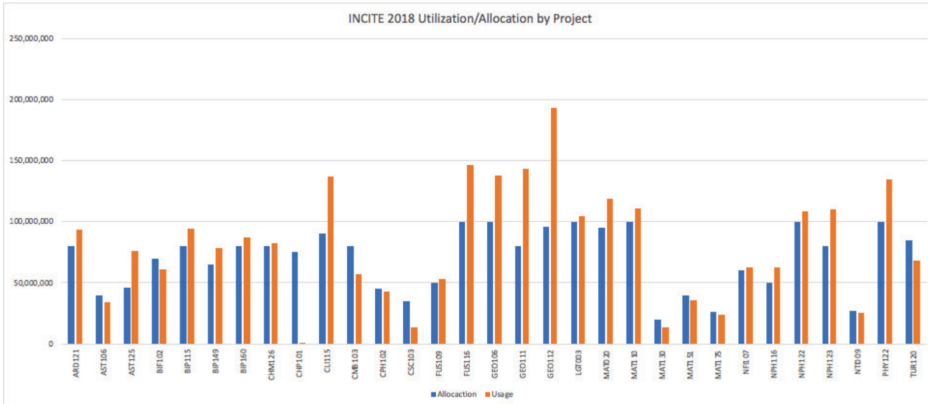

8.2.2 ALCC Allocation/Utilization for Calendar Year 2018

The ALCC allocation year is July 1–June 30 so below the OLCF reports the usage against the ALCC 2018 and ALCC 2019 Allocations separately. Usage for both programs is reported below against the full allocation amount.

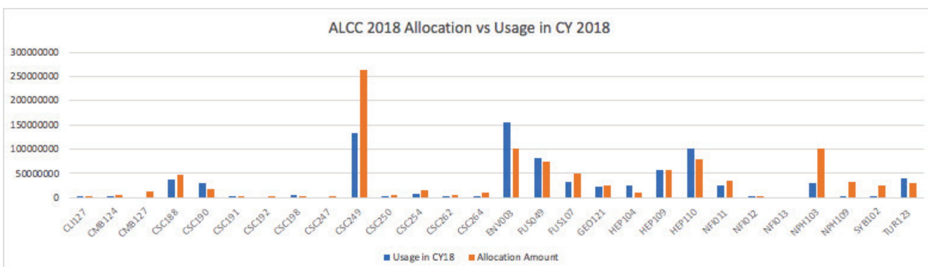

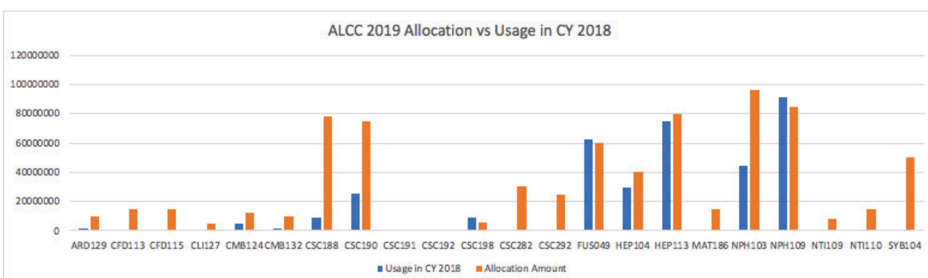

8.3 STAKEHOLDER ENGAGEMENT AND OUTREACH

8.3.1 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 continues to fund efforts at 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. (new) During 2018, the OLCF engaged with ECP and the E6 Laboratories to define and develop a security plan under which a Continuous Integration service will be deployed. The OnyxPoint service was selected, delivered, and tested at the OLCF with the ECP GitLab instance at the DOE Office of Scientific and Technical Information (OSTI) using ECP ST products and the OLCF system "Ascent".

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 Summit, as seen in the makeup of the 2018 Gordon Bell finalists. These methods required incorporating best-in-class software with cutting-edge libraries and deployment mechanisms.

8.3.2 Industry Engagement

Accelerating Competitiveness through Computational ExceLlence (ACCEL) maintained its steady course, attracting new industrial users as well as firms that are returning to solve new and larger problems, expand their experience using GPUs at scale, and grow their internal computational expertise. Project results affirm that ACCEL is helping OLCF deliver high impact science and engineering results. It also is enabling Oak Ridge National Laboratory and DOE to contribute to national competitiveness and grow the community of researchers able to use leadership computing resources, thereby strengthening the nations innovation infrastructure.

Thirty-six industrial projects were under way during 2018, which represented 9% of the total number of projects provided to external user programs (Early Science, INCITE, ALCC, and DD). These projects used 97,496,120 hours, representing approximately 2% of the total Titan hours delivered in 2018.

- In 2018, 15% of the total industrial project hours were allocated through INCITE, 55% via ALCC, and 30% through the OLCF DD program.

- Of 36 projects, 22 were new. These new projects received awards via ALCC (7 projects), and DD (15 projects).
    - Five of the new ALCC projects and three of the new DD projects were awardees through *HPC4Manufacturing* (HPC4Mfg), a program within DOE's HPC4 Energy Innovation initiative. 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.
    - One of the new DD projects was an awardee through *HPC4Mobility*, a new program within DOE's HPC4 Energy Innovation initiative. HPC4Mobility facilitates access to HPC capabilities and expertise across the DOE National Laboratories to deliver solutions that could revolutionize mobility and transportation.
    - Four companies were new to ACCEL and received DD awards: Masten Space Systems, Xanadu Quantum Technologies (quantum computing), SABIC (petrochemicals manufacturing), and ENSCO (aerospace/avionics/transportation).
    - ACCEL continued to support small businesses, granting four firms DD awards to address innovative science and engineering problems: Masten Space Systems, Xanadu Quantum Technologies, Tech-X and SmartTruck Systems.
    - ACCEL also continues to support the collaborative research programs between ORNL's National Transportation Research Center (NTRC) and the auto industry. GM received an ALCC award and partnered with NTRC to leverage recent advancements in CFD to improve the accuracy of CFD computations used in diesel passenger car engine calibrations, streamlining that process. This is enabling GM to accelerate innovative engine designs while meeting strict emissions standards.

SmartTruck Steps Up Simulations for Certification by Computation

PI: Nathan See, SmartTruck Allocation Program: DD

The Science

In 2016 the US Environmental Protection Agency (EPA) and the National Highway Traffic Safety Administration issued the Phase 2 Greenhouse Gas Emissions and Fuel Efficiency Standards for Medium- and Heavy-Duty Engines and Vehicles, which offered industry the option—for the first time—to certify accessory products for long-haul trucks by computational fluid dynamics (CFD) calculations rather than through physical testing. Using Titan, SmartTruck conducted a full unsteady analysis of its TopKit Aero add-on product and demonstrated a reduction in the parachute-like area of drag on the back of the truck. In 2018, SmartTruck used the results as proof of its compliance with Phase 2 regulations, becoming the first company to request certification through computational analysis instead of physical testing.

The Impact

Historically, companies that manufacture the trucks and their add-on parts have been permitted to certify their products only through physical testing, often using a wind tunnel in which air flows past a scaled-down model in a tunnel or using a special track for road testing. Certification by CFD would be convenient and, ultimately, cheaper by eliminating uncontrollable variables that hamper physical testing (e.g., weather, test track availability, and the cost of wind tunnel testing). The Phase 2 program requires the more rigorous unsteady (versus steady) CFD analysis, but through this kind of analysis, the program can help determine best practices that could extend to a range of CFD challenges in the industry.

Summary

The introduction of the Phase 2 Greenhouse Gas Emissions and Fuel Efficiency Standards for Medium- and Heavy-Duty Engines and Vehicles allowed companies to certify add-on products for long-haul trucks by computational fluid dynamics (CFD) calculations. To meet the rigorous EPA standards, companies need to introduce new modeling techniques to capture higher order physics. A team at SmartTruck used NASA's FUN3D code on Titan to successfully complete an unprecedented detailed unsteady analysis via modeling and simulation, and the analysis demonstrated a reduction in the drag on the back of the truck. Also called the delta drag-area improvement, this is the main criterion used by EPA for certification and, in this case, equated to an approximate 6.6 percent reduction in drag and a roughly 4–5 percent increase in fuel efficiency. In 2018, SmartTruck used these simulations and the delta drag-area improvement as proof of its compliance with Phase 2 regulations, becoming the first company to submit to EPA an aerodynamic device for certification through computational methods alone.

Figure 8.6. By forcing airflow to veer off the trailer towards the wake behind it, the TopKit Aero system (red) reduces the base wake (blue) and increases the pressure on the truck rear, creating a significant reduction in overall vehicle drag. *Image Credit: SmartTruck*

Funding

This research used resources of the Oak Ridge Leadership Computing Facility located in the Oak Ridge National Laboratory, which is supported by the Office of Science of the Department of Energy under contract DE-AC05-00OR22725.

Related Links

"SmartTruck Steps Up Simulations for Certification by Computation," OLCF News (August 29, 2018)

APPENDIX A. RESPONSES TO RECOMMENDATIONS FROM THE 2017 OPERATIONAL ASSESSMENT REVIEW

In March 2018, the OLCF presented the 2017 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 Con Call	2018 OLCF User Conference Call: Overview of the New OLCF Website	1/31/18	20
Monthly User Con Call	2018 OLCF User Conference Call: Launching Jobs with JSRUN	2/28/18	89
Monthly User Con Call	2018 OLCF User Conference Call: Profiling with Forge	3/28/18	31
Monthly User Con Call	2018 OLCF User Conference Call: CUDA 9	4/25/18	80
Monthly User Con Call	2018 OLCF User Conference Call: Kokkos Trainng Overview	6/20/18	55
Monthly User Con Call	2018 OLCF User Conference Call: New Python Infrastructure at OLCF	7/25/18	52
Monthly User Con Call	2018 OLCF User Conference Call: PBDR	8/29/18	23
Monthly User Con Call	2018 OLCF User Conference Call: JSRUN Visualizer	9/26/18	54
Monthly User Con Call	2018 OLCF User Conference Call: Summit Storage	12/12/18	43
IDEAS-ECP/Facility Webinar Series	IDEAS/ECP Seminar Series: Bringing Best Practices to a Long-Lived Production Code	1/17/18	81
IDEAS-ECP/Facility Webinar Series	IDEAS/ECP Seminar Series: Jupyter and HPC: Current State and Future Roadmap	2/28/18	112
IDEAS-ECP/Facility Webinar Series	IDEAS/ECP Seminar Series: Scientific Software Development with Eclipse	3/31/18	45
IDEAS-ECP/Facility Webinar Series	IDEAS/ECP Seminar Series: Software Citation Today and Tomorrow	4/18/18	26
IDEAS-ECP/Facility Webinar Series	IDEAS/ECP Seminar Series: On-Demand Learning for Better Scientific Software: How to Use Resources & Technology to Optimize Your Productivity	5/9/18	44
IDEAS-ECP/Facility Webinar Series	IDEAS/ECP Seminar Series: Popper: Creating Reproducible Computational and Data Science Experimentation Pipeline	6/13/18	42
IDEAS-ECP/Facility Webinar Series	IDEAS/ECP Seminar Series: How Open Source Software Supports the Largest Computers on the Planet	7/18/18	32
IDEAS-ECP/Facility Webinar Series	IDEAS/ECP Seminar Series: Introduction to Modern Cmake	9/19/18	122
IDEAS-ECP/Facility Webinar Series	IDEAS/ECP Seminar Series: From Continuous Integration to Static Linters	10/18/18	60
Workshop/Training/Meeting	Introduction to the NVIDIA Visual Profiler	2/8/18	12
Workshop/Training/Meeting	Git it Together – Using Version Control for Scientific Collboration	2/9/18	10
Workshop/Training/Meeting	Getting Started at OLCF Webinar	1/10/18	7
Workhop/Training/Meeting	Summit Application Readiness Workshop	3/5-9/2018	77
Workshop/Training/Meeting	DLI: Introduction to Deep Learning	2/28/18	84
Workshop/Training/Meeting	INCITE Webinar #1	5/2/18	64
Workshop/Training/Meeting	2018 OLCF User Meeting	5/15-17/2018	118
Workshop/Training/Meeting	Introduction to Summit	6/1/18	46

Event Type	Event Title	Date	Participants
Workshop/Training/Meeting	INCITE Webinar #2	6/7/18	98
Workshop/Training/Meeting	Data Science with R and pbdR at ORNL: From the CADES Cloud to the OLCF	6/18/18	66
Workshop/Training/Meeting	Introduction to HPC	6/26-28/2018	103
Workshop/Training/Meeting	Techniques and Tools to Start with HPC	7/13/18	66
Workshop/Training/Meeting	Performance Portability with Kokkos	7/23-26/18	56
Workshop/Training/Meeting	OpenACC Standards Committees Face-to-Face Meeting	7/31 - 8/3/18	46
Workshop/Training/Meeting	Score P / Vampire Workshop	8/17/18	21
Workshop/Training/Meeting	Profiling & Debugging with ARM Tools	9/14/18	31
Workshop/Training/Meeting	IBM HPCXXL Meeting	9/17-21/18	67
Workshop/Training/Meeting	Programming Methods for Summit's Multi-GPU Nodes	11/5/18	100
Workshop/Training/Meeting	Summit Workshop	12/3-6/18	180
Workshop/Training/Meeting	OLCF GPU Hackathon, TU Dresden	3/5-9/18	55
Workshop/Training/Meeting	OLCF GPU Hackathon, Pawsey Supercomputing Centre Hackathon	4/16-20/18	35
Workshop/Training/Meeting	OLCF GPU Hackathon, University of Colorado Boulder	6/4-8/18	59
Workshop/Training/Meeting	OLCF GPU Hackathon, National Center for Supercomputing Applications	9/10-14/18	48
Workshop/Training/Meeting	OLCF GPU Hackathon, Brookhaven National Laboratory	9/17-21/18	67
Workshop/Training/Meeting	OLCF GPU Hackathon, Swiss National Supercomputer Centre	10/1-5/18	55
Workshop/Training/Meeting	OLCF GPU Hackathon, Oak Ridge National Laboratory	10/22-26/18	83
Seminar Series	Towards an Efficient Framework for Bayesian Experimental Design, Inference, and Uncertainty Propagation	1/22/18	N/A
Seminar Series	Interpretable Deep Learning for Understanding Gene Regulation	1/29/18	N/A
Seminar Series	Charliecloud - A Lightweight Linux Container for User-Defined Software Stacks in HPC	2/7/18	N/A
Seminar Series	Transforming machine learning heuristics into provable algorithms: classical, stochastic, and neural	2/12/18	N/A
Seminar Series	Modelling Two-dimensional infrared spectra of proteins: Are we still bound to size?	2/26/18	N/A
Seminar Series	The Impact of the Power Architecture on the Summit Supercomputer	3/1/18	N/A
Seminar Series	On the quantification and efficient propagation of imprecise probabilities using Monte Carlo methods	3/14/18	N/A
Seminar Series	Deterministic and stochastic acceleration techniques for Richardson-type iterations	3/16/18	N/A

Event Type	Event Title	Date	Participants
Seminar Series	A Deep Dive into ZFS on Linux	3/23/18	N/A
Seminar Series	Machine Learning for Physics Discovery	3/28/18	N/A
Seminar Series	PathForward: DOE-Sponsored Vendor Research Projects	3/29/18	N/A
Seminar Series	Leadership AI at the OLCF	4/20/18	N/A
Seminar Series	PapyrusKV: A High-Performance Parallel Key-Value Store for Distributed NVM Architectures	5/24/18	N/A
Seminar Series	A Robust Fault-tolerant and Scalable Deduplication for Storage Systems	7/3/18	N/A
Seminar Series	System-level optimizations for high performance despite increasing heterogeneity of HPC architectures and applications	7/11/18	N/A
Seminar Series	OpenACC: Directives for GPU Computing are Ten Years Old	7/30/18	N/A
Seminar Series	Using OpenACC with the GNU Compiler Collection	7/30/18	N/A
Seminar Series	Parallel Machine Learning Software for the Prediction of Large-Scale Geostatistics Simulations using Maximum Likelihood Estimation	8/1/18	N/A
Seminar Series	Experiences Building Deductive Code Synthesis Tools for Python	8/6/18	N/A
Seminar Series	Scalable, Global Namespaces with Programmable Storage	8/7/18	N/A
Seminar Series	A Progressive Study for the Automatic Analysis of Internal Layers from Polar Radar Imagery	8/8/18	N/A
Seminar Series	Software Support for Dynamic Adaptive Neural Network Arrays	8/13/18	N/A
Seminar Series	Adaptive Parallelism Mapping in Dynamic Environments using Machine Learning	8/14/18	N/A
Seminar Series	Scaling Configuration Management for Maintaining A Compliant Environment	8/27/18	N/A
Seminar Series	Conditional Image Synthesis by Generative Adversarial Modeling	10/26/18	N/A
Seminar Series	Challenges in bidirectional porting of computer codes in an integrated cyberinfrastructure	10/8/18	N/A
Seminar Series	Capturing the complexity of renewable energy in Europe	11/1/18	N/A
Seminar Series	Application of Concurrent Programming to Numerical Computation	11/2/18	N/A
Seminar Series	Scaling Data-Driven Scientific Discovery in the Big-Data - A Data Lifecycle View	12/6/18	N/A

APPENDIX C. OUTREACH PRODUCTS

Date	Type of Product	Title
1/16/18	Highlight	A Shortcut to Modeling Sickle Cell Disease
1/16/18	Highlight	Optimizing Mini-apps for Better Portability
1/16/18	Highlight	Faces of Summit: Putting the System to the Test
1/16/18	PPT Slide	A Shortcut to Modeling Sickle Cell Disease
1/16/18	PPT Slide	Optimizing Mini-apps for Better Portability
1/16/18	PPT Slide	Faces of Summit: Putting the System to the Test
1/16/18	Quad Chart	A Shortcut to Modeling Sickle Cell Disease
1/30/18	Poster	ECP Meeting
2/1/18	Graphic	Summit graphic for Barb Helland
2/1/18	Graphic	Frontier web elements
2/7/18	Highlight	OLCF Launches New Website
2/7/18	Highlight	Particle Interactions Calculated on Titan Support the Search for New Physics Discoveries
2/7/18	Highlight	New Discoveries Within 'SIGHT'
2/7/18	Highlight	GM Revs up Diesel Combustion Modeling on Titan Supercomputer
2/7/18	PPT Slide	OLCF Launches New Website
2/7/18	PPT Slide	OLCF Launches New Website
2/7/18	PPT Slide	Particle Interactions Calculated on Titan Support the Search for New Physics Discoveries
2/7/18	PPT Slide	New Discoveries Within 'SIGHT'
2/7/18	PPT Slide	GM Revs up Diesel Combustion Modeling on Titan Supercomputer
2/7/18	Quad Chart	New Discoveries Within 'SIGHT'
2/7/18	Quad Chart	GM Revs up Diesel Combustion Modeling on Titan Supercomputer
2/7/18	Video	GM Accelerates Combustion Modeling on Titan
2/12/18	Highlight	Summit at GTC 2018
2/13/18	Highlight	Frontier: OLCF'S Exascale Future
2/27/18	Highlight	Titan Powers Data Dive into Tropical Soil Microbes
2/27/18	Highlight	A Novel Method for Comparing Plant Genes
2/27/18	Highlight	McNally Appointed Education Chair for iMasons
2/27/18	Highlight	Faces of Summit: Building a Better Summit
2/27/18	PPT Slide	Titan Powers Data Dive into Tropical Soil Microbes
2/27/18	PPT Slide	A Novel Method for Comparing Plant Genes
2/27/18	PPT Slide	McNally Appointed Education Chair for iMasons
2/27/18	PPT Slide	Faces of Summit: Building a Better Summit
2/27/18	Quad Chart	Titan Powers Data Dive into Tropical Soil Microbes
2/27/18	Quad Chart	A Novel Method for Comparing Plant Genes
3/1/18	Article	Stepping up to Summit
3/1/18	Report	Contributions for Operational Assessment
3/19/18	Poster	DAC Meeting -- OLCF Overview

Date	Type of Product	Title
3/19/18	Poster	DAC Meeting -- Summit Early Science
3/26/18	Graphic	Infographic: Summit will be the world's smartest supercomputer for open science
3/26/18	Highlight	Summit Installation Sparks a Safety Innovation
3/26/18	Highlight	Faces of Summit: Serving up Software
3/26/18	Highlight	Zeroing in on Text Understanding for Automated Cancer Report Classification
3/26/18	Highlight	Infographic: Summit will be the world's smartest supercomputer for open science
3/26/18	PPT Slide	Summit Installation Sparks a Safety Innovation
3/26/18	PPT Slide	Faces of Summit: Serving up Software
3/26/18	PPT Slide	Zeroing in on Text Understanding for Automated Cancer Report Classification
3/26/18	PPT Slide	Infographic: Summit will be the world's smartest supercomputer for open science
3/26/18	PPT Slide	With Supercomputing Power and an Unconventional Strategy, Scientists Solve a Next-Generation Physics Problem
3/26/18	Quad Chart	Zeroing in on Text Understanding for Automated Cancer Report Classification
3/26/18	Quad Chart	With Supercomputing Power and an Unconventional Strategy, Scientists Solve a Next-Generation Physics Problem
4/1/18	Graphic	Social banners for Summit launch (introducing Summit)
4/16/18	Fact Sheet	Summit White Paper one-pager for Thomas Zacharia
4/16/18	Highlight	Former OLCF User Brings Supercomputing-backed Flood Maps to FEMA
4/16/18	Highlight	Faces of Summit: Programming with Purpose
4/16/18	Highlight	A Problem with Polymer Theory
4/16/18	Highlight	Assessing the Cost of Air Pollution
4/16/18	PPT Slide	Former OLCF User Brings Supercomputing-backed Flood Maps to FEMA
4/16/18	PPT Slide	Faces of Summit: Programming with Purpose
4/16/18	PPT Slide	A Problem with Polymer Theory
4/16/18	PPT Slide	Assessing the Cost of Air Pollution
4/21/18	Graphic	Executive Board Certificates for User Meeting
4/25/18	Graphic	OLCF/NCCS door tags template (Summit branding)
5/1/18	Highlight	Faces of Summit: Tackling Storage
5/1/18	Highlight	Nuclear Physicists Wield HPC to Uncover Magic Isotopes
5/1/18	Highlight	Faces of Summit: Preparing to Launch
5/1/18	PPT Slide	Faces of Summit: Tackling Storage
5/1/18	PPT Slide	Nuclear Physicists Wield HPC to Uncover Magic Isotopes
5/1/18	PPT Slide	Faces of Summit: Preparing to Launch
5/23/18	Highlight	Another First for Quantum
5/23/18	PPT Slide	Another First for Quantum
5/29/18	Highlight	OLCF Hosts 14th Annual User Meeting amid Summit Buzz
5/29/18	Highlight	ADIOS and BigData Express Offer New Data Streaming Capabilities
5/29/18	Highlight	Faces of Summit: Leading a Systems Expedition
5/29/18	Highlight	Boosting Industry with OpenACC
5/29/18	PPT Slide	OLCF Hosts 14th Annual User Meeting amid Summit Buzz
5/29/18	PPT Slide	ADIOS and BigData Express Offer New Data Streaming Capabilities

Date	Type of Product	Title
5/29/18	PPT Slide	Faces of Summit: Leading a Systems Expedition
5/29/18	PPT Slide	Boosting Industry with OpenACC
5/30/18	Highlight	With Supercomputing Power and an Unconventional Strategy, Scientists Solve a Next-Generation Physics Problem
6/1/18	Fact Sheet	RSA Tokens one-pager
6/1/18	Video	Summit Time Lapse
6/1/18	Video	Summit Early Science
6/1/18	Website	Summit Launch materials dropbox for media
6/1/18	Website	Summit Launch
6/5/18	Fact Sheet	Lab Day tour talking points
6/8/18	Fact Sheet	Summit fact sheet edits
6/8/18	Fact Sheet	Exaop Story pdf for Jacobson
6/8/18	Graphic	Infographic: Summit by the Numbers
6/8/18	Graphic	Summit wall wrap installation
6/8/18	Graphic	Summit green backdrop
6/8/18	Graphic	Summit social tiles (x10)
6/8/18	Graphic	Summit branding elements for website
6/8/18	Graphic	Summit early science team photos for website
6/8/18	Graphic	Summit overlook screen graphic
6/8/18	Graphic	Summit Exaops pull-up banner stand
6/8/18	Graphic	Backdrop banner for launch event
6/8/18	Graphic	Summit presentation for media briefing
6/8/18	Highlight	Faces of Summit: Succeeding by Leading
6/8/18	Highlight	Genomics Code Exceeds Exaops on Summit Supercomputer
6/8/18	Highlight	Infographic: Summit by the Numbers
6/22/18	Highlight	Energy Secretary Rick Perry Attends ORNL Debut of Summit Supercomputer
6/25/18	Highlight	ORNL's Summit Supercomputer Named World's Fastest
6/25/18	Poster	OLCF's Top500 No. 1s
6/25/18	PPT Slide	Faces of Summit: Succeeding by Leading
6/25/18	PPT Slide	Genomics Code Exceeds Exaops on Summit Supercomputer
6/25/18	PPT Slide	Energy Secretary Rick Perry Attends ORNL Debut of Summit Supercomputer
6/25/18	PPT Slide	OLCF Readies Users for Summit
6/26/18	Highlight	OLCF Readies Users for Summit
6/26/18	Highlight	Buddy Bland Presented with Secretary's Appreciation Award
6/29/18	Poster	CSGF Poster
7/1/18	Article	Uncovering Magic Isotopes with the Power of HPC
7/10/18	Fact Sheet	OLCF Accomplishments for S1
7/12/18	Graphic	INCITE rebranding concept
7/15/18	Fact Sheet	Awards Night one-pagers
7/17/18	Highlight	International Teams Optimize Codes at Australia's First OLCF GPU Hackathon
7/17/18	Highlight	OLCF Workshop Introduces HPC Concepts

Date	Type of Product	Title
7/17/18	Highlight	Introduce Your Daughter to AI Event Sees OLCF Participation
7/17/18	Highlight	US Air Force and ORNL Partner in High Performance Computing and Weather Modeling System
7/17/18	Highlight	Titan Helps Scientists Fine-Tune Laser Interactions to Advance Cancer Treatments
7/17/18	PPT Slide	International Teams Optimize Codes at Australia's First OLCF GPU Hackathon
7/17/18	PPT Slide	OLCF Workshop Introduces HPC Concepts
7/17/18	PPT Slide	Introduce Your Daughter to AI Event Sees OLCF Participation
7/17/18	PPT Slide	US Air Force and ORNL Partner in High Performance Computing and Weather Modeling System
7/17/18	PPT Slide	Titan Helps Scientists Fine-Tune Laser Interactions to Advance Cancer Treatments
7/25/18	Fact Sheet	Responses to questions from FAZ one-pager
8/1/18	Poster	Frontier IPR
8/5/18	Graphic	Skins for HPSS
8/7/18	Highlight	OLCF Welcomes Wombat Test System
8/7/18	Highlight	ALCC Program Awards 14 Projects a Combined 729.5 Million Core Hours at the OLCF
8/7/18	Highlight	Transforming Gas into Fuels with Better Alloys
8/7/18	Highlight	Summer Interns Gain HPC Skills, Professional Development at the OLCF
8/7/18	PPT Slide	OLCF Welcomes Wombat Test System
8/7/18	PPT Slide	ALCC Program Awards 14 Projects a Combined 729.5 Million Core Hours at the OLCF
8/7/18	PPT Slide	Transforming Gas into Fuels with Better Alloys
8/7/18	PPT Slide	Summer Interns Gain HPC Skills, Professional Development at the OLCF
8/10/18	Graphic	Grace Hopper t-shirt design
8/29/18	Highlight	DOE Graduate Fellows Train on Titan for a Future in HPC
8/29/18	Highlight	Faces of Summit: Creating a Green Summit
8/29/18	Highlight	Preaching pbdR
8/29/18	Highlight	SmartTruck Steps Up Simulations for Certification by Computation
8/29/18	PPT Slide	DOE Graduate Fellows Train on Titan for a Future in HPC
8/29/18	PPT Slide	Faces of Summit: Creating a Green Summit
8/29/18	PPT Slide	Preaching pbdR
8/29/18	PPT Slide	SmartTruck Steps Up Simulations for Certification by Computation
9/4/18	Video	NVIDIA Video project -- 5 videos
9/17/18	Highlight	Uncharted Territory
9/17/18	Highlight	Faces of Summit: Getting Acclimated
9/17/18	Highlight	Summit Speeds Calculations in the Search for Exotic Particles
9/17/18	Highlight	Titan Speeds Inquiry into Engine Alloys That Can Take the Heat
9/17/18	PPT Slide	Uncharted Territory
9/17/18	PPT Slide	Faces of Summit: Getting Acclimated
9/17/18	PPT Slide	Summit Speeds Calculations in the Search for Exotic Particles
9/17/18	PPT Slide	Titan Speeds Inquiry into Engine Alloys That Can Take the Heat

Date	Type of Product	Title
9/17/18	Video	Joe Citeno Interview
9/25/18	Article	Mixed Precision: A Strategy for New Science Opportunities
10/1/18	Video	75-second History of Supercomputing
10/8/18	Graphic	Social media tile - Jacobson
10/9/18	Highlight	Mixed Precision: A Strategy for New Science Opportunities
10/9/18	Highlight	Simple Summit
10/9/18	Highlight	Annual OpenACC Meeting Promotes Programming Innovation
10/9/18	PPT Slide	Mixed Precision: A Strategy for New Science Opportunities
10/9/18	PPT Slide	Simple Summit
10/9/18	PPT Slide	Annual OpenACC Meeting Promotes Programming Innovation
10/9/18	Video	Simple Summit
10/10/18	Highlight	Research Team Breaks Exaop Barrier With Deep Learning Application
10/15/18	Highlight	Computing Genes to Support Living Clean
10/20/18	Highlight	Surprise finding: Discovering a previously unknown role for a source of magnetic fields
10/23/18	Highlight	Microscopy Images Put Deep Learning Code to the Test
10/29/18	Highlight	Gordon Bell Prize Finalist Code Adopts AI to Model Monster Earthquakes
11/1/18	Website	INCITE 2018
11/5/18	Highlight	OLCF Staff Participate in Merged Conferences for New Perspectives
11/5/18	Highlight	Augmented Reality Technology Shows Promise for Summit Surveillance
11/5/18	Highlight	Award Finalists Demonstrate Improved QCD Code for Supercomputing
11/5/18	PPT Slide	Research Team Breaks Exaop Barrier With Deep Learning Application
11/5/18	PPT Slide	Computing Genes to Support Living Clean
11/5/18	PPT Slide	Microscopy Images Put Deep Learning Code to the Test
11/5/18	PPT Slide	Gordon Bell Prize Finalist Code Adopts AI to Model Monster Earthquakes
11/5/18	PPT Slide	OLCF Staff Participate in Merged Conferences for New Perspectives
11/5/18	PPT Slide	Augmented Reality Technology Shows Promise for Summit Surveillance
11/5/18	PPT Slide	Award Finalists Demonstrate Improved QCD Code for Supercomputing
11/12/18	Fact Sheet	INCITE Fact Sheets
11/12/18	Graphic	Summit Early Science web tile
11/12/18	Web Article	SC18 Event Details
11/13/18	Highlight	Two DOE Supercomputers Top List of World's Fastest
11/20/18	Graphic	Conference Center banners (computing and Summit updates_
11/20/18	Highlight	Oak Ridge, Lawrence Berkeley National Labs Share 2018 ACM Gordon Bell Prize
11/27/18	Fact Sheet	Updated Summit Fact Sheet
11/27/18	Fact Sheet	Updated OLCF Fact Sheet
11/28/18	Graphic	Slides for Jeff Nichols CIO presentation
11/29/18	Highlight	Summit Early Science Program Starting Soon
11/30/18	Graphic	User survey web tile
12/3/18	Highlight	Teaching and Learning
12/3/18	Highlight	GPU Hackathon Places Summit within Reach for Visiting Teams

Date	Type of Product	Title
12/3/18	Highlight	New Storage System Boosts Efficiency to Better Preserve OLCF User Data
12/3/18	Highlight	Summit Facility Design Team Snags Tennessee Energy Award
12/3/18	Highlight	Pinnacle Engines Develops Efficient, Low-Emission Gasoline Engine Using Supercomputing
12/3/18	PPT Slide	Teaching and Learning
12/3/18	PPT Slide	GPU Hackathon Places Summit within Reach for Visiting Teams
12/3/18	PPT Slide	New Storage System Boosts Efficiency to Better Preserve OLCF User Data
12/3/18	PPT Slide	Summit Facility Design Team Snags Tennessee Energy Award
12/3/18	PPT Slide	Pinnacle Engines Develops Efficient, Low-Emission Gasoline Engine Using Supercomputing
12/10/18	Fact Sheet	Gordon Bell Winners -- combined one-pagers for Jeff Nichols
12/13/18	Graphic	Summit early science images for Director's holiday video
12/13/18	Poster	INCITE 2018 Projects Poster
12/20/18	Highlight	Breaching the Biomass Problem
12/20/18	Highlight	Ready for Science: Summit Completes System Acceptance
12/20/18	Highlight	UT Students Get Bite-Sized Bits of Big Data Centers in ORNL-Led Course
12/20/18	PPT Slide	Breaching the Biomass Problem
12/20/18	PPT Slide	Ready for Science: Summit Completes System Acceptance
12/20/18	PPT Slide	UT Students Get Bite-Sized Bits of Big Data Centers in ORNL-Led Course
12/21/18	Highlight	Titan Takes Fluoride from Taps and Toothpaste to Batteries
12/21/18	PPT Slide	Titan Takes Fluoride from Taps and Toothpaste to Batteries
12/31/18	Report	OLCF 2018 Annual Report draft

APPENDIX D. BUSINESS RESULTS FORMULAS

2018 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.

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

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.

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

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).

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

Mean Time to Failure

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

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

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).

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

APPENDIX E. DIRECTOR'S DISCRETIONARY PROJECTS ENABLED (AT ANY POINT) IN CY 2018

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Ji Qiang	LBNL	30000	2	Particle Accelerator Modeling on Exascale Computers
Jacob King	Tech-X Corporation	200000	24742	GPUelegant Performance
Shailendra Kaushik	GM	4000000	0	Multi-Disciplinary Optimization of Automobile Aerodynamic Performance, Thermal Performance, and Aero-Acoustics
Nathan See	SmartTruck	0	0	Investigating matching Unsteady CFD to Real World testing of Class 8 trucks
Eric Nielsen	NASA Langley Research Center	2000000	173762	FUN3D GPU Development
Marc-Olivier Delchini	ORNL	0	0	HPC4Mmfg - High Performance Computational Modeling of Synthetic Jet Actuators for Increased Freight Efficiency in the Transportation Industry
Allan Grosvenor	Microsurgeonbot	4000000	3899189	NextGen Low Cost Propulsion: Dual-Expander Cycle LOX/Methane Rocket Engine Supporting Research
John Schaefer	Boeing	0	0	Numerical Prediction of Nose Landing Gear Noise
Nathan See	SmartTruck	3000000	2013201	Investigation of major drivers of the fully non-linear unsteady CFD and its application for numerical certification for the wide combination of trucks and trailers integrated system using K20 GPU‚Äôs on Titan ‚Äì Phase1
Paul Shapiro	UT Austin	2000000	0	Simulating Reionization of the Local Universe: Witnessing our own Cosmic Dawn
Gonzalo Merino	U. of Wisconsin	1000000	494547	Management and Operation of the IceCube Neutrino Observatory (ICNO)
Alexander Tchekhovskoy	Northwestern University	0	0	Parallel Scaling of GPU-accelerated Simulations of Compact Binary Merger Remnant Disks, Outflows, and Nucleosynthesis
Bronson Messer	ORNL	0	0	Analysis and Visualization for Turbulence Studies in Core-collapse Supernovae
Brian O'Shea	Michigan State University	1000000	2562425	Scaling and performance enhancement of an astrophysical plasma code
Lorenzo Sironi	Columbia University	2000000	4430793	The interplay of relativistic reconnection and turbulence in astrophysical plasmas
Daniel Eisenstein	Harvard University	0	0	Cosmological N-body Simulations with Abacus

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Christian Cardall	ORNL	2000000	311	Towards 3D Core-collapse Supernova Ensemble Studies with GenASiS
Rod Linn	LANL	2000000	0	Scaling Wildland and Urban Fire Modeling Capabilities using HIGRAD/FIRETEC
Jessie Carman	NOAA-GFDL	2000000	764315	Air-Ocean-Land-Ice Global Coupled Prediction on Emerging Computational Architectures: A Framework for ESPC Coupled Models
Inanc Senocak	University of Pittsburgh	1000000	210	GEM3D: Open-Source Cartesian Adaptive Complex Terrain Atmospheric Flow Solver for GPU Clusters
Gur Pines	U. of Colorado	80000	4202	Overlapping Genes to Limit Mutational Escape
Daniel Jacobson	ORNL	20000000	0	Center For Bioenergy Innovation
Daniel Jacobson	ORNL	0	0	Scaling Up Parallelized Ortholog Detection Algorithms for Comparative Genomics of Bacterial Genomes
Lawrence Shapiro	Columbia University	4000000	0	Mechanisms of Functioning of Somatic Hypermutation on Antibody Affinity Maturation
Manjunath Gorentla Venkata	Mellanox	200000	0	Understanding and Constructing Fine-scale Genetic Map of Rhesus Macaque
Jeremy Smith	University of Tennessee	1000000	452381	Solvent Disruption of Biomass for the Production of Biofuels and Bioproducts
Kjiersten Fagnan	LBNL	0	0	JGI Data Archive
Victor Padilla-Sanchez	The Catholic University of America	0	0	Molecular Modelling and Visualization Viruses
Chongle Pan	University of Oklahoma	5000000	0	Large-Scale Integrated Omics Analyses of Complex Microbial Communities in Plant-Microbe Interfaces and Tropical Rainforests
Harel Weinstein	Cornell	5000000	7934951	Structural elements in the evolution of mechanisms in substance transporters
Joseph Curtis	NIST	500000	0	Atomistic Modeling of Small-Angle Scattering Experimental Data on HPC Resources via SASSIE-web
James Gumbart	Georgia Tech	4000000	5952095	Determining the role of AcrA in the bacterial multidrug efflux pump
Jan Michael Carrillo	ORNL, University of Tennessee	2000000	2861814	Large Scale Coarse-Grained Molecular Dynamics Simulations of Lipid Bilayers
Ron Dror	Stanford	3000000	21894388	Enabling the Design of Drugs that Achieve Good Effects Without Bad Ones
Ivaylo Ivanov	Georgia State University	0	0	Integrative Modeling of Transcription Initiation Assemblies in Gene Regulation

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Loukas Petridis	ORNL	4000000	3841345	Integrating Neutron Scattering with Molecular Simulation to Determine Structural Ensembles of Flexible Biosystems
Edward Lyman	U. of Delaware	2000000	398227	Hydrodynamics of complex membranes on subcellular scales
Aleksei Aksimentiev	UIUC	3000000	914589	All-atom structure and ionic conductivity of the nuclear pore complex
Arnold Tharrington	Department of Energy	3000000	2427588	HPC AIMD of Large Biomolecular Systems
Philip Kurian	Howard University	650000	0	Computing many-body van der Waals dispersion effects in biomacromolecules
Christopher Bahl	Institute for Protein Innovation	0	0	Design of GPCR Allosteric Modulators
Martin Karplus	Harvard University	8000000	3111829	Computational design of HIV vaccination schedule
Teresa Head-Gordon	UC Berkeley	1000000	1160149	Enhanced Sampling Methods for Biomolecular Simulation
Alison Sweeney	University of Pennsylvania	2500000	761512	Self-Assembly of Reflectin Proteins at Membrane Interfaces
Diwakar Shukla	UIUC	100000	8110	Computational investigation of nitrogen-uptake rate in plants
Robert Wilson	ORNL	500000	0	Biomimetic Vector Sensor Towed Arrays (BVSTA)
Sumanta Acharya	Illinois Tech	500000	0	Indirect Dry Cooling for Power Plants Using Encapsulated Phase Change Materials
Yue Ling	Baylor	2000000	393986	Direct Numerical Simulation of Spray Formation in a Two-phase Mixing Layer Between Two Parallel Gas and Liquid Streams
Jesse Ault	ORNL	2000000	726870	Using Parallel Computation to Improve Blood Cell Simulations in Cardiovascular Flows
David Gutzwiller	Numeca	2000000	679859	Porting and Demonstration of the NUMECA Omnis Environment for Multidisciplinary Simulation and Optimization.
Stephan Priebe	GE	3500000	2023392	Genesis LES Studies
David Buchta	UIUC	0	0	Effects of floating marine debris on nonlinear wave dynamics
Marc Olivier Delchini	ORNL	0	0	Nuclear Energy University Program (NEUP) - Computational Fluid Dynamics Benchmark Analysis of an Historical Sodium Loop Experiment
Charlotte Barbier	ORNL	0	0	Cavitation Bubbles near Solid Surfaces

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Vyacheslav Bryantsev	ORNL	500000	0	Rational Design of Flotation Agents for Rare Earth Ore Minerals
Nike Dattani	Oxford	1000000	0	Competing with Quantum Computers for Finding the Ground State Energy of the Iron-molybdenum Complex
John Parkhill	Notre Dame	2000000	1523	Accelerating Computational Chemistry a Million Times Over by Feeding Data-hungry Algorithms
Remco Havenith	U. of Groningen	3000000	4113994	GRONOR
Byeongjin Baek	SABIC	2000000	17	catalyst screening for hydrogenation processes
Srikanth Allu	ORNL	1000000	2483	Consortium of Advanced Battery Simulation
Kwangho Nam	University of Texas at Arlington	500000	12804	Multiscale Modeling of Kinase Conformational Change
Rio Yokota	Tokyo Institute of Technology	0	0	ADAC: Accelerated Data Analytics and Computing
Stephan Irle	ORNL	2000000	21615	GPU-accelerated quantum chemical ensemble-parallel nanomaterials growth simulations on microsecond timescales
Wibe de Jong	LBNL	2000000	0	The chemical universe through the eyes of deep learning neural networks
Olexandr Isayev	University of North Carolina at Chapel Hill	3500000	819769	Learning Quantum Mechanics (QM) with Neural Networks
Shiwei Zhang	William & Mary	2000000	5357832	Quantum Simulations of Photosystem II and Cuprate Superconductivity
Arun Yethiraj	U. of Wisconsin	2000000	582163	Polymers in deep eutectic solvents
Jason Byrd	ENSCO	2000000	0	Scaling and performance of Aces4 when applied to very large environmental chemical pollutants
Vyacheslav Bryantsev	ORNL	1000000	59	Atomic-Scale Simulations of Bulk and Interfacial Processes in Molten Salt Environment
Filip Pawlowski	Auburn University	4000000	37916	Coupled-cluster singles-and-doubles quality excitation energies for large molecular systems
Deborah Penchoff	University of Tennessee	1500000	8716000	Selective Binding of Rare Earth Elements and Actinides for Nuclear Forensics Applications
Ada Sedova	ORNL	4000000	4128298	Creation of a Tested HPC-based Workflow to Calculate VISION INS Spectra From DFT-based Frequency Calculations
Allison Talley	Proctor & Gamble Company	2000000	1680959	Investigation of anomalous NMR cross relaxation signals of surfactant aggregates via molecular dynamics simulations
Philippe Sautet	UCLA	2000000	0	Modelling catalysis at bimetallic surfaces

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Moetasim Ashfaq	ORNL	8000000	461137	The Computational Climate Science Integrated Allocation
William Gustafson	PNNL	0	0	The DOE Atmospheric Radiation Measurement Programs LES ARM Symbiotic Simulation and Observation (LASSO) Initialization, Forcing and Multiscale Data Assimilation Program
Dali Wang	ORNL	500000	15	Intelligent Knowledge Management Framework for Earth Science Data Initiative
Valentine Anatharaj	ORNL	100000	0	Provisioning of Climate Data
Clayton Naber	Pinnacle Engines	0	0	Development of Opposed-Piston Variable Compression Ratio Engine for Automotive Applications using High-Performance Computational Combustion Methodologies
Gary Cai	Fiat Chrysler Automobiles	0	0	A Fundamental Study of the Factors Affecting Adverse Autoignition (knock) in Internal Combustion Engines
Ramanan Sankaran	ORNL	0	0	HPC4mfg - Level-set Modeling Simulations of Chemical Vapor Infiltration for Ceramic Matrix Composites Manufacturing
Seung Hyun Kim	Ohio State University Research Foundation	1000000	23008	Development of a Physics-Based Combustion Model for Engine Knock Prediction: Phase II
Xinyu Zhao	University of Connecticut	900000	9490	Turbulence and flame resolved fire simulations
Venkat Raman	U. of Michigan	2000000	2177574	High-fidelity Simulations of Rotating Detonation Engines
Dario Alfe	University College London	4000000	2127878	New Frontiers for Material Modeling via Machine Learning Techniques with Quantum
Panchapakesan Ganesh	ORNL	2000000	43151	Data Driven Discovery by Design of Energy Materials
Andreas Glatz	ANL	2000000	1785460	OSCon
Anton Kozhevnikov	CSCS	0	0	ADAC-CSCS
Gabriel Kotliar	Rutgers	2000000	572296	application of GPU-accelerated quantum Monte-Carlo impurity solver to plutonium compounds
Prineha Narang	Harvard University	500000	0	Ab initio Exploration of Hydrodynamic Transport in Weyl and Dirac Semimetals
Jamison Daniel	ORNL	1000000	0	Cluster-based Visualization of Tera- and Peta-Scale Datasets
Neena Imam	ORNL	9000000	189497	Durmstrang

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Sergey Panitkin	BNL	1000	22885991	Next Generation Workload Management System
Bronson Messer	ORNL	1000000	0	CORAL Benchmarking
Norbert Podhorszki	ORNL	30000000	5627645	ADIOS
Robert Patton	ORNL, Department of Energy	3000000	196467	Scalable Deep Learning Systems for Exascale Data Analysis
Joshua New	ORNL	1000000	280066	Big Data Mining for Building Analytics
Manuel Arenaz	Appentra	50000	0	Porting Parallware Tool to Large HPC Installations Including Titan
Catherine Schuman	ORNL	2500000	407069	Scalable Evolutionary Optimization for Designing Networks
Ramakrishnan Kannan	ORNL	1000000	356599	Mini-Apps for Big Data
Jeffrey Young	Georgia Tech	750000	0	Evaluation and Porting of ORNL Mini-apps to Future Directive-Based Languages and Runtimes
Cory Hauck	ORNL	800000	624441	Diagnostics for Data Compression at Scale
Terry Jones	ORNL	3000000	471	UNITY: Unified Memory & Storage Space
Abhinav Vishnu	PNNL	3500000	0	Machine Learning on Extreme Scale GPU systems
Audris Mockus	University of Tennessee	1500000	4057443	Fingerprinting Online Users Using Doc2Vec Model and Bayesian Networks
John Cavazos	U. of Delaware	1000000	0	Large-Scale Distributed and Deep Learning of Structured Graph Data for Real-Time Program Analysis and Characterization
Jeremy Johnston	ORNL	3500000	0	Surrogate Based Modeling for Deep Learning Hyper-parameter Optimization
Peter Zaspel	U. of Basel	3000000	2606133	Extreme-scale Many-core Solvers for Data Analysis and Uncertainty Quantification
Jakub Kurzak	University of Tennessee	500000	0	SLATE
Catherine Schuman	ORNL	2000000	1779738	Scalable Neuromorphic Simulators: High and Low Level
Ian Foster	ANL	300000	20587	CODAR
Sunita Chandrasekaran	U. of Delaware	20000	754	Migrating Legacy Code to Novel Directive-based Programming Models
Mark Berrill	ORNL	0	0	Summit Acceptance Apps
Robert Brook	University of Tennessee	0	0	JICS/AACE Emerging Architectures Activities
John Turner	ORNL	0	0	ALExa ECP
George Bosilca	University of Tennessee	500000	3303	Open MPI x

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Hartwig Anzt	University of Tennessee	0	0	ECP-PEEKS
Allan Grosvenor	Microsurgeonbot	1500000	41551	Intelligent Middleware, Making CFD Accessible Through Leadership-Scale Deep Learning Research
Terry Jones	ORNL	1000000	0	ECP: Simplified Interface for Complex Memories
Judith Hill	ORNL	500000	11863	Computational Science Graduate Fellowship Program
Lin-Wang Wang	LBNL	4000000	3980543	HPC4mfg: Making Semiconductor Devices Cool Through HPC Ab-initio Simulations
Lonnie Crosby	University of Tennessee	3484800	0	UTK Benchmarking and Performance Projections
Alexander McCaskey	ORNL	300000	0	xacc-tnqvm-ibm-quantum-computing
Tjerk Straatsma	ORNL	0	0	Summit Benchmarking for Early Science
Tjerk Straatsma	ORNL	0	0	CANDLE Benchmarking for Early Science
Ramakrishnan Kannan	ORNL	0	0	Graph500 for OLCF
Ramakrishnan Kannan	ORNL	1000000	0	Scalable Non-linear Unmixing
Tjerk Straatsma	ORNL	0	0	Summit Benchmarking for Early Science - Edgar
Thomas Karnowski	ORNL	1000000	0	HPC4mfg - Reinforcement Learning-based Traffic Control to Optimize Energy Usage and Throughput
Tjerk Straatsma	ORNL	0	0	Summit Benchmarking for Early Science - Stancil
Seung-Hwan Lim	ORNL	3000000	0	A GPU-accelerated high dimensional data management for machine learning workloads in large scale GPU-CPU environments.
Helia Zandi	ORNL	2500000	0	Smart Grid Transactive Scalable Hybrid Management System
Jean-Roch Vlimant	Caltech	2000000	990483	large scale deep learning training & optimization
Peter Zaspel	U. of Basel	2000000	179	Extreme-scale linear solvers in high-dimensional approximation, machine learning (ML) and beyond
Karan Vahi	USC	100000	0	Pegasus Project
Robert Patton	ORNL, Department of Energy	2000000	0	Scalable Machine Learning for Connected Autonomous Vehicles
William Tang	Princeton	1000000	0	Extreme Scale PIC Research on Advanced Architectures

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
William Tang	Princeton	6000000	3	Big Data Machine Learning for Fusion Energy Applications
Igor Igumenshchev	Rochester	2000000	3070762	Development of Simulation Tools for Laser Fusion Experiments
David Green	ORNL	100000	0	Sparse Grid Solver for Fusion
George Fann	ORNL	100000	0	Optimizing High-order Time-stepping for Cgyro on SUMMIT
Stuart Loch	Auburn University	2000000	105413	High performance computation in support of neutron star mergers and high-Z plasma facing components on fusion energy experiments
Kyle Withers	US Geological Survey	1000000	0	Improving numerical simulations of earthquake ground motions using high-resolution 3-D seismic velocity and improved rupture models
Dalton Lunga	ORNL	200000	0	FMOW
Yidong Xia	Idaho National Laboratory	2000000	29533	Dissipative Particle Dynamics Modeling of Multiphase Flow and Transport in Nanoporous Shale Pore-Networks
Tsuyoshi Ichimura	University of Tokyo	0	0	Fast & Scalable Finite Element Method with Low-order Unstructured Elements for Earthquake Simulation on Summit
John McNelis	Oak Ridge Associated Universities (ORAU)	0	0	ORNL DAAC Data Visualization at EVEREST 2018
Katrin Heitmann	ANL	1000000	1362667	Sky Surveys
Jean-Roch Vlimant	Caltech	2000000	0	HEP Deep Learning
Boram Yoon	LANL	2000000	767639	Artificial Intelligence for Collider Physics
Malachi Schram	PNNL	1000000	0	Belle II
Eric Church	PNNL	0	0	Algorithms for Innovative Data Analysis
Brian Nord	Fermilab	2500000	0	Machine Learning for Astrophysics and Cosmology
Andre Walker-Loud	LBNL	2500000	6410924	The proton's structure and the search for new physics
Richard Brower	BU	0	0	ECP Lattice Field Theory Solvers
Liqun Zhang	TN Tech	4000000	8325831	Molecular Dynamics Simulations to Investigate the Structure, Dynamics and Functional Properties of Human Beta Defensin Type 3
Vipin Sachdeva	Silicon Therapeutics	1500000	128907	Large-scale Validation of Drug Discovery Pipeline Protocols
Darrin York	Rutgers	0	0	Computational Tools for High-throughput Lead Optimization

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Byung Park	ORNL	5000000	11958885	Accelerating Materials Modeling Loop of Leadership Computing and Spallation Neutron Source
Dongwon Shin	ORNL	4000000	118964	High Performance Cast Aluminum Alloys for Next Generation Passenger Vehicle Engines
Yangyang Wang	ORNL	3000000	9800099	Elucidating the Influence of Reversible Non-Covalent Interactions on Dynamic Properties for Rational Design of Soft Materials
Rick Archibald	ORNL	2000000	0	CADES/OLCF Computational Workflows for Materials Science
Marcus M√ºller	U. of G√∂ttingen	750000	398237	SOMA: Soft Coarse-Grained Monte-carlo Acceleration
Peter Cummings	Vanderbilt	4000000	48798	Materials Genome Screening of Soft Materials
Adolfo Eguiluz	University of Tennessee	2400000	59940	GW Many-body Solver for Ab-initio Multi-orbital Hamiltonians for Correlated Materials
Emmanuel Vallejo	UAEH	100000	0	Theoretical and Experimental Study on Polymeric Molecular Self-assembly on Metallic Substrates
Jacek Jakowski	ORNL	3000000	0	Electronic Structure Modeling of Impurity States
Bobby Sumpter	ORNL	10500000	15829339	Center for Nanophase Materials Sciences (CNMS)
Lin Lin	UC Berkeley	1000000	159	ESMATH
Bo Kong	Ames Laboratory	0	0	Multi-scale Simulations of Gas Atomization Process for Metal Alloy Powder Production
Axiel Birenbaum	ORNL	3000000	2063937	Reduced Symmetry in Oxide Heterostructures: Defects and Interfaces
Haixuan Xu	University of Tennessee	2000000	952396	Dislocation Interaction with Radiation-Induced Defect in Structrual Materials
Mordechai Kornbluth	Robert Bosch	500000	70883	Conductivity Mechanisms in Glasses
Miguel Fuentes-Cabrera	ORNL	1000000	45607	Molecular Dynamics Simulations of Liquid Metal Assembly at the Nanoscale
Trung Nguyen	Northwestern University	1000000	187668	Development of Models and Accelerated Sampling Techniques for Massively Parallel Molecular Dynamics Codes
Nouamane Laanait	ORNL	0	0	Deep Learning Applications in Microscopy at Mixed-Precision
James Morris	ORNL	2000000	72183	HPC4mfg - High-throughput calculations of light-weight alloys

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Zhiting Tian	VA Tech	1000000	742	Thermal Transport Properties of Two-Dimensional Polymers
Mohsen Asle Zaeem	U. of Missouri	3000000	354000	PREDICTING NANO/MICROSTRUCTURES AND PROPERTIES IN SOLIDIFICATION OF METALS
Roberto Longo	Texas A&M	1500000	0	First Principles Analysis and Design of Materials for Catalysis and Energy
Rajiv Kalia	USC	1000000	58881	Benchmarking Billion-to-Trillion Atom Molecular Dynamics Simulations of Turbulence
Prineha Narang	Harvard University	250000	5175	Ab initio Predictions of Spatially-resolved Nonequilibrium Transport Phenomena in Quantum Systems
Jiook Cha	Columbia University	2000000	75945	High-throughput brain imaging analysis in Alzheimer's disease related dementia.
Hong-Jun Yoon	ORNL	1000000	242	LungXNet: Scalable AI for High Throughput Multi-Labeled Medical Image Analysis
Predrag Krstic	SUNY Stony Brook	3000000	0	Gas-Liquid-Solid Interfaces for Energy Applications
Bamin Khomami	University of Tennessee,	3000000	11458013	Multiscale Modeling
Massoud Kaviany	U. of Michigan	1500000	985949	Human Thermosensation: Response of TRPV1 Domains to Temperature
Brendan McLaughlin	QUB	2000000	24	T-Iron
Phillip Stancil	UGA	1200000	4	Quantum Dynamics of Molecule-Molecule Collisions in Full-Dimensionality
Xinlian Liu	Hood College	0	0	Rapid Molecular Machine Learning through Dimensionality Reduction
Rachel Slaybaugh	UC Berkeley	4000000	1334848	Ex-core Dosimetry Calculations for a Pressurized Water Reactor and Comparison to Operational Data
Askin Guler Yigitoglu	ORNL	3000000	0	Nuclear-Renewable Hybrid Energy Systems
Igor Sfiligoi	General Atomics	10000	128236	BFG3D
Rajan Gupta	LANL	100000	0	Probing Novel Physics using Nucleon Matrix Elements
Aurel Bulgac	U. of Washington	4000000	2340202	Nuclear Fission: From More Phenomenology and Adjusted Parameters to More Fundamental Theory and Increased Predictive Power
Leah Broussard	ORNL	100000	222999	nEDM at SNS
Leah Broussard	ORNL	100000	0	nEDM at SNS
Kaushik Datta	USC	0	0	PRISMA

PI	Institution	Most Recent Titan Allocation	Usage	Project Name
Fabien Delalondre	EPFL	3000000	0	Brain Simulation Modeling
Benjamin Cumming	CSCS	0	0	Arbor
Josh McDermott	MIT	0	0	Next-Generation Hearing Aids Via Neural Network Models of the Auditory System
Brian Edwards	University of Tennessee	1000000	0	Nanoscale Ion Transport Through Carbon Nanotubes
Jerzy Bernholc	NC State	2000000	66735	Theoretical Investigations of Nanostructures
Prineha Narang	Harvard University	3000000	986547	Atom-by-Atom Quantum Defect Dynamics from First Principles
Jason Pries	ORNL	1500000	2960	High Fidelity Electric Motor/Generator Modeling and Optimization
Tameem Albash	USC	8000000	26729565	True Scaling of Time to Solution on a Quantum Annealer
Brajesh Gupt	Xanadu	3000000	1685997	Demonstrating quantum supremacy with Titan
Edison Liang	Rice U.	1500000	117148	Magnetized Jet Creation Using a Ring of Laser Beams
Itay Hen	USC	2000000	3774517	Exploring Error Mechanisms in Analog Quantum Computers
Andreas Herten	Forschungszentrum Juelich GmbH	0	0	GTC18 OpenPOWER OpenACC Tutorial
Chongle Pan	University of Oklahoma	2000000	877447	Proteogenomics Analysis of Tropical Soil Communities Under Long-term Nitrogen and Phosphorus Fertilization
John Schaefer	Boeing	0	0	Uncertainty Quantification and Sensitivity Analysis of Turbulence Model Coefficients for the Common Research Model
Joseph Smith	Missouri University of Science and Technology	1000000	1378	Numerical Investigation of Flow Instabilities Inside a Westinghouse SMR
Raul Payri	UPV	1500000	1544393	DNS Atomization
Pino Martin	U. of Maryland	2000000	1119905	Scaling the CRoCCo Code
Srikanth Allu	ORNL	3000000	0	hpc4MFG
Duane Rosenberg	Colorado State University	100000	0	GPU Performance Analysis for Turbulence Codes