ORNL/SPR-2015/133

US Department of Energy, Office of Science

High Performance Computing Facility Operational Assessment 2014 Oak Ridge Leadership Computing Facility

February 2015

Ashley D. Barker
David E. Bernholdt
Arthur S. Bland
James J. Hack
James H. Rogers
Brian Smith
T.P. Straatsma
Kevin G. Thach
Sudharshan S. Vazhkudai
Jack C. Wells
Julia C. White

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

ORNL/SPR-XXXX/XXX

US Department of Energy, Office of Science

HIGH PERFORMANCE COMPUTING FACILITY OPERATIONAL ASSESSMENT 2014
OAK RIDGE LEADERSHIP COMPUTING FACILITY

Ashley D. Barker
David E. Bernholdt
Arthur S. Bland
James J. Hack
James H. Rogers
Brian Smith
T.P. Straatsma
Kevin G. Thach
Sudharshan S. Vazhkudai
Jack C. Wells
Julia C. White

Date Published: February 2015

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

CONTENTS

Page

LIST OF FIGURES

Figure **Page**

LIST OF TABLES

Table **Page**

ABBREVIATED TERMS

3D	three-dimensional
ACCEL	Accelerating Competitiveness through Computational ExceLlence
ACME	Accelerated Climate Model for Energy
ACTS	Assessment and Commitment Tracking System
ADIOS	Adaptable IO System
ADP	automated discovery process
ALCC	ASCR Leadership Computing Challenge
ALCF	Argonne Leadership Computing Facility
ALICE	A Large Ion Collider Experiment
ALS	Advanced Light Source
ALTD	Automatic Library Tracking Database
ANL	Argonne National Laboratory
ASCAC	ASCR Advisory Council
ASCR	Advanced Scientific Computing Research
ASDF	Adaptable and Modern Seismic Data Format
ATO	authority to operate
AU	aprun usage
C&A	certification and accreditation
CAAR	Center for Accelerated Application Readiness
CADES	Compute and Data Environment for Science
CAM-SE	Community Atmosphere Model—Spectral Element
CASL	Center for Advanced Simulation of Light Water Reactors
CCI	Common Communication Interface
CCS	carbon capture and sequestration
CCSN	core-collapse supernova
CEP	Central Energy Plant
CESM	Community Earth System Model
CNMS	Center for Nanophase Materials Science
CPU	central processing unit
CSCS	Swiss National Computing Center
CSEEN	Computational Scientists for Energy, the Environment and National Security
CY	calendar year
DART	Days Away, Restricted, or Transferred
DD	Director's Discretionary
DESI	Dark Energy Spectroscopic Instrument
DFT	density functional theory
DIMM	dual in-line memory module
DME	dimethyl ether
DNE	Distributed Namespace
DOE	Department of Energy
DOE-OSO	DOE ORNL Site Office
DRAM	dynamic random access memory
EECS	Electrical Engineering and Computer Science
ESNet	Energy Sciences network
EVEREST	Exploratory Visualization Environment for Research in Science and Technology
FEMA	Federal Emergency Management Agency
FP	functional partitioning
FY	fiscal year

Gb	gigabit
GB	gigabyte
GPU	graphics processing unit
HACC	Hardware/Hybrid Accelerated Cosmology Code
HiPSTAR	High-Performance Solver for Turbulence and Aeroacoustics Research
HOMME	High-Order Methods Modeling Environment
HPC	high-performance computing
HPSS	High Performance Storage System
INCITE	Innovative and Novel Computational Impact on Theory and Experiment
IOSI	I/O Signature Identifier
ISI	Institute for Scientific Information
ISM	integrated safety management
IT	information technology
LBNL	Lawrence Berkeley National Laboratory
LDRD	Laboratory Directed R&D
LHC	Large Hadron Collider
libPIO	I/O placement strategy
LSST	Large Synoptic Survey Telescope
MB	megabyte
MiniApps	reduced, proxy applications
MPI	Message Passing Interface
MTTF	mean time to failure
MTTI	mean time to interrupt
NCCS	National Center for Computational Sciences
NERSC	National Energy Research Scientific Computing Center
NICS	National Institute of Computational Sciences
OA	operational assessment or overall availability
OAR	Operational Assessment Report
OLCF	Leadership Computing Facility
OPV	organic photovoltaics
ORAU	Oak Ridge Associated Universities
ORISE	Oak Ridge Institute for Science and Education
ORNL	Oak Ridge National Laboratory
OSO	ORNL Site Office
OUG	OLCF Users Group
PEMP	Performance Evaluation and Measurement Plan
PETSc	Portable, Extensible Toolkit for Scientific Computing
PFS	parallel file systems
PIC	particle-in-cell
PM	preventative maintenance
PTS	Publication Tracking System
QMC	quantum Monte Carlo
R&D	research and development
RATS	Resource Allocation Tracking System or Resource and Tracking System
RMP	Risk Management Plan
RSS	research safety summary
RUR	Resource Utilization Reporting
S3D	scientific application
SA	scheduled availability
SBMS	Standards-Based Management System
SciDAC	Scientific Discovery through Advanced Computing

SM	streaming multiprocessor
SNS	Spallation Neutron Source
SR	self-reports
STCI	Scalable runTime Component Infrastructure
SUPER	Institute for Sustained Performance, Energy, and Resilience
TB	terabyte
TRC	Total Recordable Cases
UAO	User Assistance and Outreach
UT-B	UT-Battelle

Executive Summary

HIGH PERFORMANCE COMPUTING FACILITY
2014 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2015

EXECUTIVE SUMMARY

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

The OLCF upgraded the 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, where each existing AMD Opteron connects to the NVIDIA Kepler as a CPU-GPU pair. The completed XK7 system, with more than 27 PF of peak computational capacity, is named Titan.

The reporting year, CY 2014, is one of superlatives: the highest-yet rating from users on overall satisfaction, the greatest number of core-hours delivered to research projects, and the largest percentage of capability usage since the OLCF began tracking the metric in 2009. This, coupled with the extremely high utilization rate, represents the fulfillment of the promise of Titan: maximum use by maximum size simulations.

The impact of all of this and more are reflected in the science accomplishments of OLCF users, with publications this year in notable journals *Science*, *Nature*, *Nature Physics*, *Nature Communications*, *Nature Scientific Reports*, the *Proceedings of the National Academy of Sciences*, *Geophysical Review Letters*, *Physical Review X*, and *Physical Review Letters*. The science achievements included in the *2014 OLCF Operational Assessment Report* reflect first-ever or largest simulations in their communities; for example, the first ab initio simulation of a cuprate, one of the class of high-temperature superconductors that have a major future role to play in our energy economy.

The Titan system provides the largest extant heterogeneous architecture for computing and computational science. Usage is high, delivering on the promise of a system well-suited for capability simulations for science. This success is due in part to innovations in tracking and reporting the activity on the compute nodes, and using this information to further enable and optimize applications, extending and balancing workload across the entire node. The OLCF continues to invest in innovative processes, tools, and resources necessary to meet continuing user demand. The facility's leadership in data analysis and workflow was featured at the Department of Energy (DOE) booth at SC14, highlighting work with researchers and institutions in the US and abroad to couple its unique computational and data resources with experimental and observational data across a broad range of scientific domains.

Effective operations of the OLCF play a key role in the scientific missions and accomplishments of its users. This Operational Assessment Report (OAR) delineates the policies, procedures, and innovations implemented by the OLCF to continue delivering a multi-petaflop resource for cutting-edge research. This report covers 2014, which, unless otherwise specified, denotes January 1, 2014, through December 31, 2014.

COMMUNICATIONS WITH KEY STAKEHOLDERS

Communication with the Program Office

The OLCF regularly communicates with the Advanced Scientific Computing Research (ASCR) Program Office through a series of established events. These include weekly Integrated Project Team calls with the local DOE ORNL Site Office (DOE-OSO) and the Program Office, monthly highlight reports, quarterly reports, the annual Operational Assessment, an annual Budget Deep Dive, 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 sponsoring agencies, potential users, and the public.

Communication with the User Community

The OLCF's communications with users take a wide variety of forms and are tailored to the objective, ranging from relating science results to the larger community or instructing users on more efficiently and effectively using OLCF systems. The OLCF offers many training and educational opportunities throughout the year for both current facility users and the next generation of high-performance computing (HPC) users (see Section 1.4.6).

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 4.5 in 2014. Ninety-six percent of respondents (272) rated their overall satisfaction with communications from the OLCF as "satisfied" or "very satisfied, which was an increase from eighty-five percent the previous year. Ninety-eight percent of users responded that they feel adequately informed of OLCF changes. The OLCF uses a variety of methods to communicate with users, including the following:

- weekly e-mail message
- welcome packet
- general e-mail announcements
- opt-in e-mail notification lists
- OLCF website
- conference calls
- OLCF User Council and Executive Board
- one-on-one interactions through liaisons and analysts
- social networking vehicles

SUMMARY OF 2014 METRICS

In consultation with the DOE program sponsor and proposed in the 2013 OAR, a series of metrics and targets were identified to assess the operational performance of the OLCF in CY 2014. The 2014 metrics, target values, and actual results as of December 31, 2014, are noted throughout this report and summarized in Section 8. The OLCF exceeded all of the metric targets.

RESPONSES TO RECOMMENDATIONS FROM THE 2013 OPERATIONAL ASSESSMENT REVIEW

In February 2014 the OLCF presented the 2013 operational activities of the center to our sponsor, the DOE Office of Advanced Scientific Computing Research (ASCR). The complete list of questions posed to reviewers and their responses, ORNL actions, and DOE ASCR comments and actions are given in Appendix A. The reviewers recommended tracking and reporting the use of GPUs. This *2014 OLCF Operational Assessment Report* discusses the additional functionality within the resource accounting and tracking system that was developed and introduced, beginning in 2013 and continuing through the current reporting period. Highlights of GPU-related activities and reports can be found in the following sections.

- Section 2.12, GPU Activity, has results from the ALTD tracking mechanism and the new Resource Utilization Reporting (RUR) tool.
- Section 2.12 contains case studies showing GPU activity of select OLCF user projects and collaborations to further increase GPU usage.
- Section 4.1 defines and describes how the RUR tracking mechanism was implemented.
- Section 4.1.2 highlights the first detailed study of GPU reliability on a large scale. The results describe and quantify very low failure rates for the Kepler GPU, which contributed to the high system availability.

Many of the science achievements by users of Titan would only have been possible through the employment of GPUs. Section 3, Strategic Results, highlights several scientific studies incorporating GPUs in full production simulations. Section 1, User Support, illustrates support and collaboration between OLCF staff and users to promote and facilitate full, effective use of the Titan hybrid architecture.

User Results

HIGH PERFORMANCE COMPUTING FACILITY
2014 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2015

1. USER RESULTS

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

OLCF RESPONSE: Yes. In 2014, the Oak Ridge Leadership Computing Facility (OLCF) supported 1,071 users on more than 250 projects. The OLCF has established a user support model for effectively supporting users that is 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 was 4.6, which was an increase from 4.4 in 2013 and 4.2 in 2012. Overall ratings for the OLCF were positive; 96% of users reported being "satisfied" or "very satisfied." The survey also asks users to rate their overall satisfaction with OLCF Support Services; the mean rating in 2014 was 4.6, which was an increase from 4.4 in 2013.

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 user support 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 2014 with 90% of tickets being resolved within 3 business days. The OLCF continues to enhance its user support, collaboration, training, outreach, and communication. The center also engages in activities to promote high-performance computing (HPC) to the next generation of researchers.

1.1 USER RESULTS SUMMARY

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

To promote continual improvement at the OLCF, users are sent a survey soliciting their feedback regarding support services and their experience as users of the facility. The 2014 survey was launched on October 2, 2014, and remained open for participation through November 24, 2014. The survey was sent electronically to individuals with enabled accounts on Innovative and Novel Computational Impact on Theory and Experiment (INCITE), Advanced Scientific Computing Research (ASCR) Leadership Computing Challenge (ALCC), and/or Director's Discretionary (DD) projects. Of 967 possible respondents, 312 users completed the survey, an overall response rate of 32.3%.

Information was collected about the various users, user experience with the OLCF, and OLCF support capabilities. Attitudes and opinions regarding performance, availability, and possible improvements for the OLCF and its staff were also solicited. Data collected from the user survey was analyzed by the Oak Ridge Institute for Science and Education (ORISE) using both quantitative and qualitative methods. The two fundamental goals that drove the collection and subsequent analysis were to catalog the types of users and to understand their needs. Analysis included basic descriptive statistics and qualitative coding of responses to open-ended questions. Responses to specific survey items were used to cross-check the responses to other, directly related items to ensure that all responses were valid (e.g., only people who selected that they had used a particular machine could rate their satisfaction with various aspects of that machine). The results of the 2014 survey can be found on the OLCF website.[1]

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

Table 1.1. 2014 User result metrics summary

Metric description	2013 target	2013 actual	2014 target	2014 actual
Overall OLCF Satisfaction score on the user survey.	3.5/5.0	4.4/5.0	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 ½ of questions that scored below satisfactory (3.5) in the previous period.	No question scored below satisfactory (3.5/5.0) on the 2013 survey.	Results will show improvement in at least ½ of questions that scored below satisfactory (3.5) in the previous period.	No question scored below satisfactory (3.5/5.0) on the 2014 survey.
OLCF survey results related to problem resolution.	3.5/5.0	4.4/5.0	3.5/5.0	4.6/5.0
Percentage of user problems addressed within 3 business days.	80%	92.3%	80%	90.0%
Average of all user support services ratings.	3.5/5.0	4.4/5.0	3.5/5.0	4.6/5.0

1.2 USER SUPPORT METRICS

The OLCF exceeded all of the User Support metrics for 2014. The OLCF metric targets and calendar year (CY) actual results 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 2013		CY 2014	
	Target	Actual	Target	Actual
Overall OLCF Satisfaction rating	3.5/5.0	4.4/5.0	3.5/5.0	4.6/5.0
Average of all user support services ratings	3.5/5.0	4.4/5.0	3.5/5.0	4.6/5.0

[1] http://www.olcf.ornl.gov/media-center/center-reports/2014-outreach-survey/

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 agreed upon by the DOE OLCF program manager define 3.5/5.0 to be satisfactory.

Overall ratings for the OLCF were positive with 96% of users responding that they are satisfied or very satisfied with the OLCF overall. With regard to the degree of overall satisfaction with the center, the percentage of satisfied and very satisfied respondents has steadily increased from 2007 (86%) to the highest percentage to date of 96% in 2014.

Key indicators from the survey, including overall satisfaction, are shown in Table 1.3. They are summarized and broken out by program. The data shows that the satisfaction among all three allocation programs is fairly similar for these four key satisfaction indicators.

Table 1.3. Satisfaction rates by program type for key indicators

Indicator	Mean	Program		
		INCITE	ALCC	DD
Overall satisfaction with the OLCF	4.6/5.0	4.6/5.0	4.7/5.0	4.6/5.0
Overall satisfaction with support services	4.6/5.0	4.5/5.0	4.7/5.0	4.6/5.0
Overall satisfaction with compute resources	4.6/5.0	4.6/5.0	4.7/5.0	4.6/5.0
Overall satisfaction with data resources	4.4/5.0	4.3/5.0	4.3/5.0	4.4/5.0

1.2.2 Average Rating across All User Support Questions

The response to the specific question on overall satisfaction with user support services provided by the OLCF was high, with an average response of 4.6/5.0. The calculated mean of all answers to all user support services questions on the 2014 survey was 4.6/5.0, indicating that the OLCF exceeded the 2014 user support metric target and that users have a high degree of satisfaction with user support services. In response to an open-ended question about the best qualities of the OLCF, user assistance was listed as the top choice by 50% of the survey respondents. The following comments are samples from the survey:

> "I think the best qualities of OLCF have remained the same: world class computing capabilities, backed up by excellent user liaison service and technical support. The center really wants to see us succeed and will often go beyond expectations to enable this."

> "Besides the great high-performance computing facilities, the best thing about the OLCF is the great user support which enables fast and efficient service."

> "Love the fact that you keep the compilers (especially C++) up to date with the latest releases!"

> "Support staff is very helpful, compute systems are fast and efficient, and technical documentation and training materials are also very helpful."

1.2.3 Improvement on Past Year Unsatisfactory Ratings

Each year the OLCF works to show improvement in 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

2013 and 2014 surveys. However, based on feedback received from the 2013 User Survey in conjunction with other feedback channels, the OLCF took the following actions in 2014 to enhance the user experience at the OLCF.

- Implemented cross-system job submission
- Purchased necessary equipment to increase the quotas on the NetApp in 2015
- Deployed additional data transfer nodes
- Implemented a supported Globus endpoint
- Developed additional training tutorials for the OLCF website
- Assembled a data management guide for the OLCF website
- Made enhancements to the "MyOLCF" user portal
- Began development of an enhanced User Dashboard expected to launch in CY 2015

1.2.4 Assessing the Effectiveness of the OLCF User Survey

Before sending the survey, OLCF staff met with the ORISE evaluation specialist to review the content of the survey questions to ensure that they accurately addressed the concerns of the OLCF and that all technical terminology was appropriately used. The evaluator specifically reviewed the response options for each of the selection items and discussed how variations in question type could influence the meaning and utility of the data they would generate.

Several targeted notifications were sent to those eligible to participate in the survey. The initial survey invitation from James Hack (center director) was sent on October 2, 2014, and subsequent follow-up reminders were sent by Ashley Barker (User Assistance and Outreach [UAO] group lead), Jack Wells (director of science for the National Center for Computational Sciences [NCCS]), ORISE, the OLCF User Council, and individual members of the OLCF. The survey was advertised on the OLCF website and was mentioned in the weekly communications e-mail sent to all users. Survey responses were tracked on a daily basis to assess the effectiveness of the various communication methods. The notifications from center management and the OLCF User Council were the most effective, but the results show that other efforts, such as including the notice in the weekly communication, also contributed to the survey response rate.

The OLCF has a relatively equally balanced distribution of new users and users for 1 to 2 years. The number of users in the more than 2 years category increased in 2014 (Table 1.4).

Table 1.4. User survey participation

	2013 survey	2014 survey
Total number of respondents (Total percentage responding to survey)	367 (30%)	312 (32%)
New users (OLCF user <1 year)	43%	26%
OLCF user 1–2 years	26%	27%
OLCF user >2 years	31%	48%

Survey respondents were asked to classify the program types with which they were affiliated. Table 1.5 is a summary of responses according to respondent affiliation and program type.

Table 1.5. User survey responses by affiliation and program type

Category	Response rate (%)[a]
Affiliation[(b)]	
University	40
DOE/Lab/government	34
Other	7
Industry	6
Foreign	13
Program type[(b)]	
INCITE	53
DD	51
ALCC	32
Other	0.3

[a] Total is greater than 100% because survey respondents can be associated with more than one type of project.
[b] Percentage of the original survey list.

1.2.4.1 Statistical Analysis of the Results

Statistical analysis of key survey areas is shown in Table 1.6. The results reflect overall satisfaction with the facility, services, and computational resources.

Table 1.6. Statistical analysis of key results

Survey topic	# of survey respondents	# of survey responses to this question	Mean	Variance	SD
Overall satisfaction with the OLCF	312	307	4.64	0.40	0.63
Overall satisfaction with Titan	312	257	4.46	0.34	0.58
Overall satisfaction with data resources	312	283	4.37	0.71	0.84
Overall satisfaction with support services	312	292	4.56	0.58	0.76
Overall satisfaction with communications	312	272	4.54	0.34	0.58
Overall satisfaction with training	312	229	4.39	0.42	0.65
Overall satisfaction with the website	312	275	4.35	0.46	0.68

The OLCF examined the variance and standard deviation for several key questions and found them to be within acceptable parameters. Most of the responses were within one standard deviation from the mean. No question was above two deviations from the mean.

1.3 PROBLEM RESOLUTION METRICS

The operational assessment (OA) metrics for problem resolution are the following:

- 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 (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 can resolve a reported problem directly, including identifying and executing the necessary corrective actions so that the problem is resolved from the user's perspective. Occasionally, the center receives problem reports for which it has limited ability to resolve the root cause because of factors beyond its control. In such a scenario, "addressing the problem" requires that OLCF staff identify and carry out all corrective actions at its 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 ticket with the product vendor; provide the vendor necessary information about the issue; and then provide a workaround to the user, if possible.

The OLCF uses request tracker software to track queries (i.e., tickets) and ensure that response goals are met or exceeded. Users may submit queries via e-mail, the online request form, or phone. E-mail is the predominant source of query submittals. The software collates statistics on tickets issued, turnaround times, and so on, to produce weekly reports. These statistics allow the OLCF staff to track patterns and address anomalous behaviors before they have an adverse effect on the work of other users. The OLCF issued 2,482 tickets in response to user queries for CY 2014 (Figure 1.1). The center exceeded the problem-resolution metric and responded to 90% of the queries within 3 business days (Table 1.7).

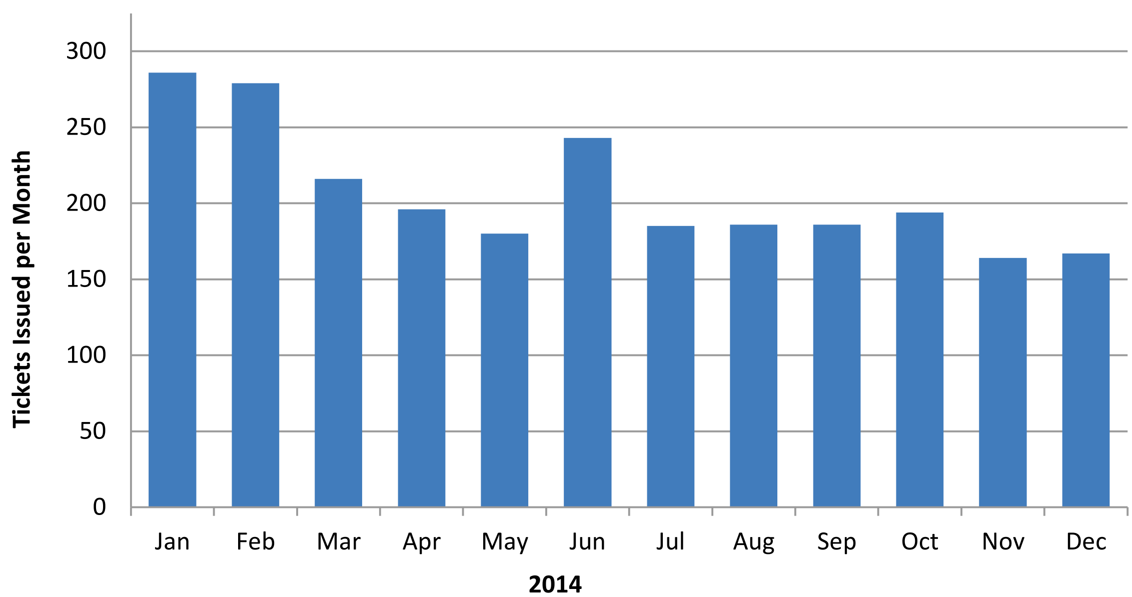

Figure 1.1. Number of helpdesk tickets issued per month.

Table 1.7. Problem resolution metric summary

Survey Area	CY 2013		CY 2014	
	Target	Actual	Target	Actual
Percentage of problems addressed in 3 business days	80%	92.3%	80%	90.0%
Average of problem resolution ratings	3.5/5.0	4.4/5.0	3.5/5.0	4.6/5.0

Each ticket is assigned to one user assistant or account analyst, who establishes customer contact and tracks the query from first report to final resolution, providing not just fast service but also service tailored to each customer's needs. Although they are dedicated to addressing queries promptly, user assistants and account analysts consistently strive to reach the "right" or best solution rather than merely achieve a quick turnaround. Tickets are categorized by the most common types. The top three reported categories in 2014 were account queries, login issues, and batch queues.

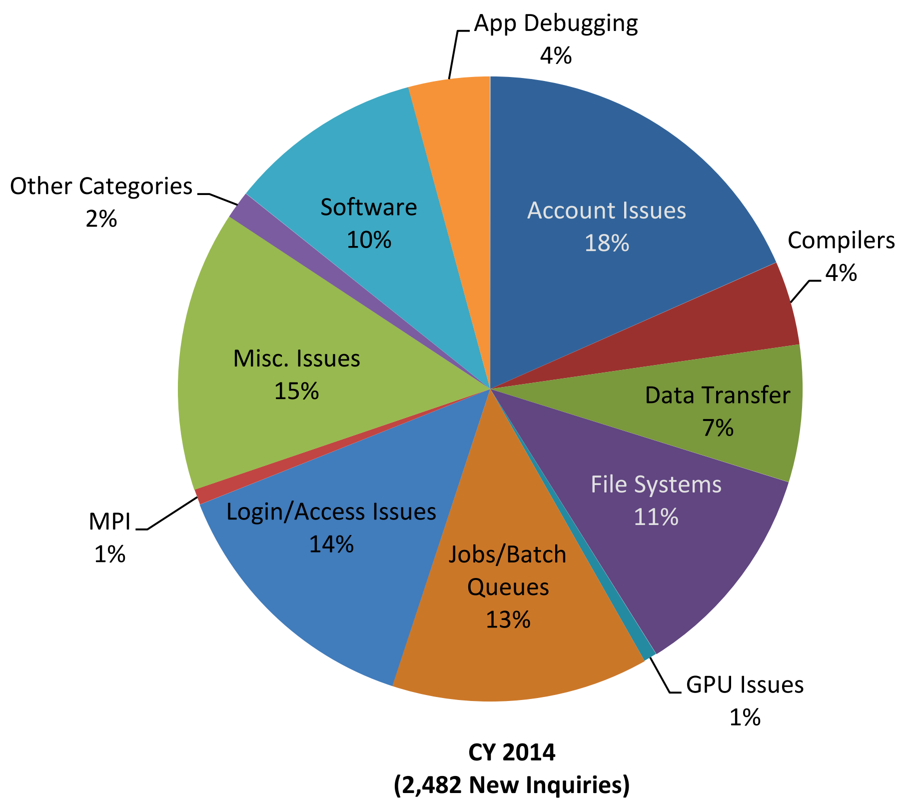

Figure 1.2. Categorization of helpdesk tickets.

1.4 USER SUPPORT AND OUTREACH

The OA data for user support and outreach includes the following:

- Examples of in-depth collaboration between facility staff and the user community
- A summary of the training and outreach events conducted during this period

The following sections discuss key activities and contributions in the areas that the OLCF recognizes as the pillars of user support and outreach.

1.4.1 User Support

The OLCF recognizes that users of HPC facilities have a range of needs requiring a range of 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 center provides complementary OLCF user support vehicles that include user assistance and outreach staff as well as scientific, data, and visualization liaisons.

1.4.1.1 User Assistance and Outreach Analysts

UAO analysts are responsible for addressing user queries. Some of the most common UAO activities include the following:

- enabling access to OLCF resources
- helping users compile and debug large science and engineering applications
- identifying and resolving system-level bugs in conjunction with other technical staff and vendors
- installing third-party applications and providing documentation for usage
- engaging other OLCF staff to ensure that users have up-to-date information about OLCF resources and to solicit feedback
- researching, developing, and maintaining reference and training materials for users
- communicating with users
- developing and delivering training
- acting as user advocates

1.4.1.2 Scientific Liaisons

Liaisons are a unique OLCF response to HPC problems faced by users. Several examples of liaison collaboration with and support of users are provided in Sections 1.4.1.2 and 1.4.1.3.

Scalability and physical fidelity of XGC

OLCF’s scientific liaison Ed D’Azevedo collaborated with Choong-Seock Chang's INCITE Project "High-Fidelity Simulation of Tokamak Edge Plasma Transport” to improve the scalability and physical fidelity of the XGC advanced fusion simulation suite, which is developed in the Center for Edge Physics Simulation project in the DOE Scientific Discovery through Advanced Computing (SciDAC) program.

XGC1 is a first-principles kinetic particle-in-cell (PIC) code for understanding the turbulent plasma transport near the edge region of a tokamak fusion device. The success of ITER is critically dependent upon effective confinement of high-temperature plasma, especially near the edge region.

As the project’s INCITE liaison, D’Azevedo has contributed to ensure the scalability and performance portability of the simulation code. XGC1 has been scaled to more than 16,000 nodes on Titan and takes advantage of both the GPUs and 16-core CPUs. XGC1 has also been ported to the Piz Daint Cray XC30 at the Swiss National Computing Center (CSCS) and the Mira IBM Blue Gene at Argonne Leadership Computing Facility (ALCF). The porting exercise uncovered various incompatibilities between compilers and libraries. For example, OpenMP “critical” directives with variables might not be accepted, or some compilers might not support nested OpenMP. As the code was scaled to a large number of MPI tasks, other code options for performing collectives and broadcasts using the MPI library were used to conserve memory or to obtain better performance.

D’Azevedo has also ported the key kernel computations such as the particle push routines to execute on the GPU version using PGI CUDA Fortran, and there is ongoing investigation of the effectiveness of OpenACC compiler directives. OpenACC version 2.0 has a new capability to call device functions within computational kernels; however, the compiler technology is still in development since accuracy, robustness, and performance issues still must be resolved before migration into the production code.

Different ways of expressing the same computation may yield very different performance levels on the GPUs.

As the physics models are refined and enhanced, the performance kernels of the XGC code will evolve as well. For example, activation of a new Fokker-Planck-Landau collision physics on highly non-Maxwellian edge plasma nearly doubled the overall computation time. D'Azevedo worked with Oak Ridge National Laboratory's (ORNL's) Patrick Worley in the SciDAC Institute for Sustained Performance, Energy, and Resilience (SUPER) to perform performance profiling to find opportunities for optimization. They identified a load imbalance and insufficient work per grid cell to fully exploit OpenMP parallelism. The module was restructured to use OpenMP processing across multiple grid cells. D'Azevedo worked with Barry Smith, developer of PETSc (Portable, Extensible Toolkit for Scientific Computing) in the SciDAC FastMath Institute, to obtain a special thread-safe version. He also developed a thread-safe option using LAPACK band solver. The NTCC pspline cubic spline library was also modified to make pspline thread-safe. The resulting collision module was accelerated by nearly 5× and overall time in XGC1 reduced by 1.7×.

Scalability of LS-DALTON and CPU/GPU utilization

Dmitry I. Lyakh is the scientific liaison collaborating with the INCITE project "Large-Scale Coupled-Cluster Calculations of Supramolecular Wires" led by Poul Jørgensen, director of the qLEAP center and professor at Aarhus University in Denmark. The primary goal of this INCITE project is to extend the applicability of coupled-cluster theory, one of the most accurate yet practical formalisms of electronic structure theory, to large chemical systems composed of hundreds to thousands of atoms. The successful accomplishment of this goal could revolutionize the field of electronic structure theory, leading to a much better predictive ability delivered by ab initio quantum-chemical simulations.

During the course of the past year, Lyakh has been improving the scalability of the LS-DALTON code and increasing the overall efficiency of its CPU/GPU utilization on heterogeneous HPC nodes. The LS-DALTON code is a part of the DALTON quantum chemistry software suite and is developed primarily at the qLEAP center. Lyakh removed the replicated memory bottleneck in the linear-scaling coupled-cluster module of the LS-DALTON code, thus improving the scalability of the most computationally expensive part of the code, the tensor contractions. In the original parallelization scheme, entire distributed arrays had to be gathered in local buffers of a node to perform a BLAS call. Lyakh's new approach performs distributed tensor contractions with a constant amount of additional (buffer) memory allocated on each HPC node, which does not depend on the number of MPI processes in the MPI communicator. This removes the major obstacle for achieving scalability in the linear-scaling coupled-cluster module.

Contrary to the original implementation, the module created by Lyakh decomposes tensor contractions (computationally expensive operations) into tasks distributed among MPI processes. To maximize the computational efficiency on each node, task pipelining has been implemented so that the MPI data transfers can be overlapped with matrix multiplications. To reduce the overhead of converting tensor contractions into matrix multiplications, Lyakh designed and implemented an efficient tensor transpose algorithm, which will further increase the efficiency of CPU/GPU utilization on each HPC node. The general GPU porting is still in progress, although several largest tensor contractions have already been ported on GPUs by the LS-DALTON developers using OpenACC. To date, Lyakh has designed a unified scheme for simultaneous CPU/GPU utilization in the coupled-cluster module of LS-DALTON, based on the task decomposition introduced in his CPU tensor contraction module. Furthermore, he has written a general-purpose library for performing tensor contractions on NVIDIA GPUs, which can be used as a backend. These contributions have been directly committed to the LS-DALTON github repository and will be available in forthcoming public releases.

In LS-DALTON, the MPI communication pattern in the coupled-cluster module is complex, making the use of MPI-3 one-sided (passive) communications necessary to minimize synchronization overhead and leverage the availability of RDMA engines on the Cray Gemini network. Unfortunately, multiple middleware problems have been observed in the Cray DMAPP library, on which the Cray one-sided MPI is based. Lyakh has been working with Cray staff in identifying those issues, describing the essence of the problem and potential reasons, preparing bug reports, and so on. Once these problems are resolved, it is anticipated that Lyakh's contributions to improving the overall scalability and efficiency of LS-DALTON will be demonstrated at scale.

GPU port of CAME-SE tracer transport routines

OLCF science liaison Matt Norman collaborated with Mark Taylor, Sandia National Laboratories, on his INCITE project "High Resolution Simulation for Climate Means, Variability and Extremes" to improve the quality of a GPU port of the Community Atmosphere Model—Spectral Element (CAM-SE) tracer transport routines.

Norman did significant work to optimize the existing implementation of CAM-SE's tracer transport routines on the GPUs. For example, he rewrote the pack and unpack routines for element coupling to make them more readable and enable them to better use the GPUs' DRAM bus, achieving a 20% performance improvement of these routines. Also, he placed additional variables into local per-streaming multiprocessor (SM) shared memory for the hyperviscosity routines, resulting in a 4× improvement in runtime for those particular kernels. Finally, he restructured each routine so that a fixed number of vertical levels are simulated within a CUDA block. This both improves efficiency and allows any number of vertical levels to be specified without overflowing the per-SM shared memory resources, which would cause the model to fail. These CUDA enhancements will allow future problems to be run on GPUs, previously untenable because of vertical level limits. They also improved the speedup of the transport routines by approximately 50%. Taylor and his team expect to use these improvements in their production INCITE runs in 2015.

Compilation of the Accelerated Climate Modeling for Energy (ACME) simulation code took approximately 90 minutes at the beginning of 2014, resulting in a significant bottleneck for debugging and tweaking parameters for future tuning purposes. Norman and Richard Archibald (of ORNL) altered the scripts for the cmake build system of the dynamical core of the atmospheric model, HOMME (High-Order Methods Modeling Environment), which reduced the compilation time of the atmospheric portion of the model by roughly 20 minutes. After profiling experiments, this entailed reducing optimizations on a select few files, which are not highly important to runtime. Also, Norman performed numerous profiling experiments for CAM-SE with Cray, Intel, and PGI compilers. The conclusion reported to the project was that the Cray compiler outperforms the others for typical production problem sets, largely because of its improved vectorization. For the Cray compiler, in particular, Norman reproduced, reported, and worked around numerous bugs to enable the compiler to succeed for HOMME. There were also additional numerous module problems with parallel-netcdf, all of the cray-netcdf modules, and the asyncpe/5.23 module, all of which resulted in failed compilation and were identified and fixed by Norman so that the project could continue.

QMCPACK debugging and GPU assistance

In late March, the OLCF received a GPU assistance request from two INCITE projects led by Paul Kent and David Ceperley, respectively. Both projects actively develop and use QMCPACK in their production quantum Monte Carlo (QMC) calculations. The teams had an unidentified GPU memory allocation problem on Titan, which occurred only when the particle count went beyond 1,500 electrons. It thus placed a severe limitation on the physical systems that the GPU code could handle.

INCITE scientific liaison Ying Wai Li of the OLCF identified two separate issues after inspection of the error behaviors. First, she determined that the memory allocation failure was due to improper linking or mismatched CUDA compiler/runtime/libraries versions on Titan. She mended the problem by rerouting those paths and making relevant changes to ensure proper compilation and linking of the code. Second, she found that the particle count limitation was an implementation flaw that exhausted the GPU thread count and shared memory during matrix inversions. She eliminated this problem by completely rewriting this part of the GPU code, introducing the use of cuBLAS library.

Enabling large particle count calculations is a major improvement that benefits advancements not only on Titan but on all GPU platforms. The November 10, 2014, public release of QMCPACK contained Li's modified code and featured this contribution as one of the release highlights.

1.4.1.3 Data Liaisons

I/O bottleneck in SPECFEM3D_GLOBE

SPECFEM3D_GLOBE is a code that simulates global and regional seismic wave propagation. Jeroen Tromp of Princeton University and his team are using Titan to study the structure of Earth using adjoint tomography simulations. Adjoint tomography involves monitoring the interaction of a forward wavefield in which the waves travel from the source to the receivers, and an "adjoint" wavefield in which the waves travel inversely from the receivers to the source. The major bottleneck to running the code at scale on Titan involved the enormous amount of I/O required to process the input files and to output the results. To address these bottlenecks, the team started using the Adaptable IO System (ADIOS) middleware library for both input and output of data. ADIOS capabilities allowed them to reduce their dependence on I/O and focus on the scientific aspects of the simulation. To analyze and understand the scientific results, a parallel reader was developed that directly reads ADIOS output data into the VisIt visualization tool (Figure 1.3). This reader has been optimized and tuned to provide specific functionality to the scientific team as well. An addition outcome of this work is the creation of a new seismic data format to improve parallel I/O performance and workflow. This new Adaptable and Modern Seismic Data Format (ASDF) has been published on github, and is available for use and further development by the solid earth sciences community.[2]

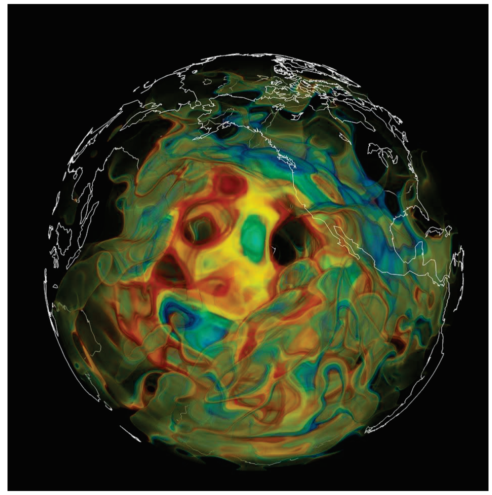

Figure 1.3. Volume rendering of shear wave perturbations of ADIOS data using VisIt.

PIConGPU and scaling I/O

The PIConGPU team had trouble in scaling their application I/O during runs at OLCF for the 2013 Gordon Bell award submittal (for which they were named finalists). OLCF staff instructed and trained the team on incorporating ADIOS, and they began using it. They tested the I/O performance on the Atlas file system in June 2014 and experienced much better performance using ADIOS than using their previous I/O solution based on parallel HDF5 (Figure 1.4; note the data are reported based on bits, not bytes). The team, led by PI Michael Bussman, has a 2015 INCITE for the project "Targeting Cancer with High Power Lasers," during which additional collaboration to enhance the performance will be carried out.

[2] http://asdf.readthedocs.org/en/latest/; https://github.com/krischer/ASDF

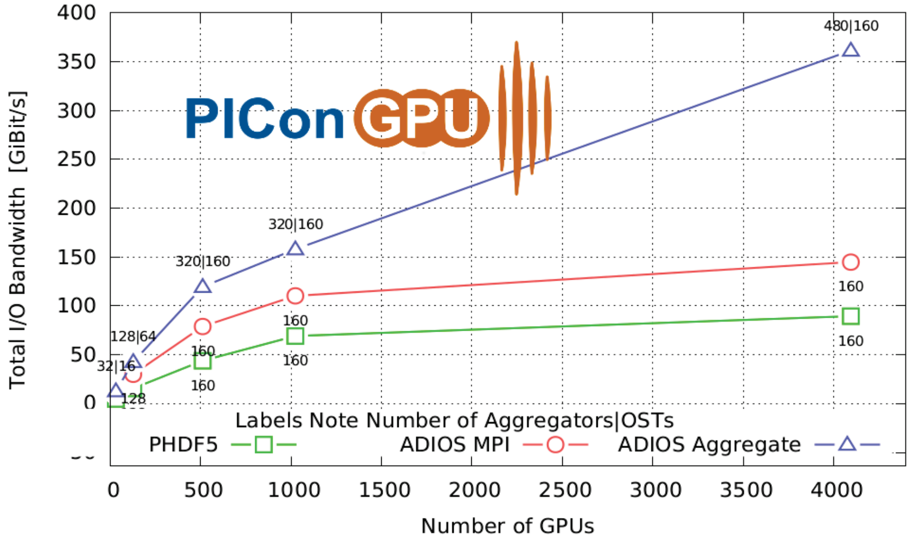

Figure 1.4. Comparison of total I/O bandwidth using different I/O solutions.

ADIOS added to QMCPACK to manage analysis

QMCPACK is a framework for ab initio many-body electronic structure simulations. A problem of the quantum MC approach is that the total size of the analysis output is very large. In practice, the data has been so large and has been generated so frequently that only simple averages have been written out for later analysis. For example, in a specific configuration to determine the properties of a graphite system consisting of 64 atoms and 256 electrons, using 4,096 nodes of Titan (i.e., 64K cores per hour of allocation), QMCPACK generates about 250 GB of data every minute; and the total can be higher for simulations of other materials. Therefore, routines would have to be added to the QMCPACK code directly to generate more output data than the existing simplified statistics.

The data liaisons have been working to add ADIOS into QMCPACK and to enable writing out of all data for later analysis. When QMPACK is converted from double precision to single precision, only half as much data is written. Additionally, the metadata collection part of the ADIOS output step was turned off to skip the step involving global communication among the processes. The metadata is generated after the run on a login node. The resulting I/O overhead is around 4%, which is acceptable.

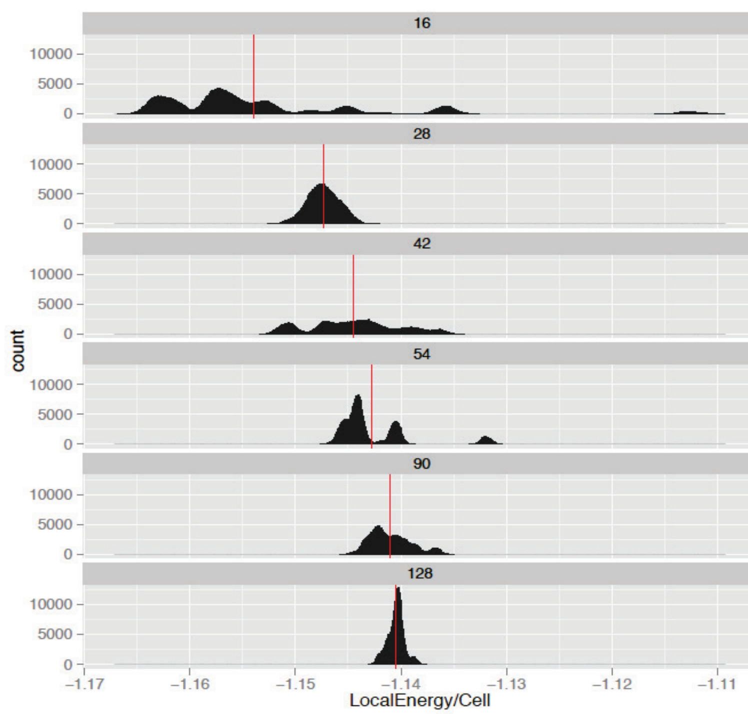

Figure 1.5. Preliminary study of convergence and distributional properties of QMCPACK.

The liaisons are also working to develop deeper statistical analysis of QMCPACK data on two fronts. One front is the analysis of the new detailed data output with ADIOS. The liaisons developed a new ADIOS reader for pbdR (a system for scalable statistical analysis, see pbdr.org) and are exploring new statistical analysis of the detailed data. The goal is to have a combination of in situ analysis and postprocessing analysis with advanced statistical methods.

The second analysis front is to use statistical design of experiments to reduce the number of required QMCPACK runs while meeting precise error bound requirements. A preliminary result is a new look at convergence and distributional properties of QMCPACK. Figure 1.5 shows that mean (red) convergence over cell size (top to bottom) is proceeding as expected, but distributions are more skewed and multimodal than expected.

1.4.1.4 Visualization Liaisons

Liquid crystals

Liquid crystals are a diverse family of macromolecules composed of rigid and flexible segments. In recent years, liquid crystals have fueled the widespread adoption of thin flat-panel electronic displays for computers and consumer televisions; in the future, they could be employed in nanoscale coatings, improved optical and photovoltaic devices, or advanced biosensors. However, these new applications require the use of thin film liquid crystals, which are susceptible to rupture and tearing. Understanding the driving forces of these ruptures is necessary to achieve better control over liquid crystal thin films.

A Titan Early Science project was the first to study liquid crystal thin films at experimental length and time scales and relate the rupture, or dewetting, process to the molecular-level driving forces. Led by Scientific Computing staff member Michael Brown, the Early Science project collaborated with OLCF visualization staff member and computational scientist Michael Matheson, who developed the necessary software to visualize these large-scale simulations. Software developed to aid in the visualization and analysis provided invaluable insight into the complex mechanisms in these huge systems of liquid crystals.

Current software tools were unable to process and analyze models at this scale. Even the interactive work with smaller models was very slow. Rendering software was developed that could deal with large, complex molecular models and manipulate them in real time. Advanced lighting was used in rendering to create images that give valuable cues to the scientist, such as determining spatial orientation and position among the molecules.

This work appeared as the cover story in the March 2014 print edition of *Nanoscale*, a high-impact journal of the Royal Society of Chemistry,[3] with a cover image created by Matheson (Figure 1.6).

Figure 1.6. Visualization of Titan Early Science project generated by Michael Matheson, OLCF.

[3] T. D. Nguyen, J.-M. Y. Carrillo, M. A. Matheson, and W. M. Brown. Rupture mechanism of liquid crystal thin films realized by large-scale molecular simulations. *Nanoscale*. **6** (6), 2014.

Core-collapse supernova

A core-collapse supernova (CCSN) is the explosion of a star, powered by the gravitational collapse of the star's core from a radius that can be thousands of kilometers to a radius of tens of kilometers. Scientifically, understanding supernovae is important to insight into how elements are formed. Michael Matheson of the Scientific Computing visualization staff collaborated with the INCITE project "Three-Dimensional Simulations of Core-Collapse Supernovae with Chimera" to visualize its CCSN simulation to gain understanding into the physical and chemical processes of supernovae and identify quantitative measures of these physical processes. The project is led by Eric J. Lentz, a computational stellar astrophysics with appointments in the Department of Physics and Astronomy at the University of Tennessee and in the ORNL Physics Division,

Recent numerical work has clearly demonstrated that the CCSN problem is thoroughly asymmetric and requires the coupling of detailed neutrino transport and nuclear physics with self-gravitating fluid dynamics. The complexity of these simulations required a collaborative effort between Matheson and the Lentz team to develop visualizations of the full-physics three-dimensional (3D) CCSN simulations that isolate features of the behavior. These provide insight into the physics of the CCSN mechanism and generate imagery that can communicate that insight to other CCSN specialists, scientists, and the general public. The workflow necessary to create these visualizations was continuously modified and optimized. Specifically, the visualization tool VisIt was run in parallel across hundreds of nodes of the visualization and analysis clusters. Each node was able to process a subset of data with parameters chosen by exploring the simulation results in an iterative fashion. As a result, animations could be viewed in the OLCF's visualization laboratory in 3D stereo almost as soon as the simulation had completed a time step. The processing method also allowed changes to the analysis and visualization techniques, as new animations could be generated in less than an hour. The ability to generate simulations quickly was extremely valuable for CCSN scientists, as it provides a method to gain insight from the spatially and temporally complex results of the simulation in near real time.

Plasma blobs

XGC is a gyrokinetic PIC code used to simulate plasma inside a magnetically confined tokamak. Studying the plasma along the edges is critical to understanding the plasma as a whole. "What happens at the edge is what determines the steady fusion performance at the core," says Princeton Plasma Physics Laboratory researcher, C-S. Chang. But when it comes to studying the edge, "the effort hasn't been very successful because of its complexity," he added. To help the scientific team understand and analyze their simulation output, a data reader was developed for VisIt that directly reads the ADIOS time-varying output files. Once read in, VisIt can perform a variety of visualization and analysis operations to help scientists understand the science being simulated. The data reader was used to visualize the blobs that occur along the edge of the plasma. Understanding the blobs is critical to understanding the behavior of the plasma, and visualizations like those produced by VisIt help scientists in their analysis. For more details of Chang's science achievements, see Section 3.2.5.

1.4.2 Outreach and Education

It is not feasible to discuss all of the support activities conducted by the OLCF staff in 2014. Sample activities include:

- The survey once again showed that users prefer online training to any other form of training. The OLCF continued to add new tutorials to the OLCF website, including "Running ParaView on Titan," "Concurrent Kernels," and "Concurrent Kernels II." Feedback on the tutorials has been very positive and users continue to tell us that they appreciate the additional training tools. The number of page views for the tutorials increased from 10,715 in 2013 to 17,570 in 2014. Users rated the tutorials on the OLCF website at 4.4/5.0 on the User Survey.

- Outreach staff created a data management user guide to help co-locate all of the user support documentation related to data management to make it easier for users to find the information. Since its release in the fall, the guide has had more than 550 page views.
- Staff helped facilitate the implementation of a supported Globus endpoint by advocating for the need, helping with the testing, and providing the user documentation for the service.
- Staff continued enhancements to the OLCF Resource Allocation Tracking System (RATS) to address additional reporting requirements, including GPU statistics. See Section 4.4.3 for details.
- Staff continued updating the "MyOLCF" site with additional functionality and began work on a replacement site with new and enhanced features to debut in 2015. Here is a typical comment:
    - "If I log into MyOLCF, the interface has greatly improved, and has a nice display for queued jobs and account usage."

1.4.2.1 OLCF User Group and Executive Board

Elections for the OUG executive board were held for the first time at the 2014 OLCF user meeting. OUG members nominated a slate of candidates and elected three to fill 3-year terms on the OUG Executive Board. Seven additional members from the previous OLCF User Council and other user committees were appointed to shorter terms on the board. Three members will be elected each year at the annual user meeting to fill vacancies left as the terms of three board members expire.

The chair of the 2014 Executive Board is Balint Joo and the vice-chair is user Mike Zingale. The following are the current board members.

Appointed to 1 year terms:

- Balint Joo, Executive Board chair, Thomas Jefferson National Accelerator Facility
- C. S. Chang, Princeton Plasma Physics Laboratory
- Stephane Ethier, Princeton Plasma Physics Laboratory

Appointed to 2 year terms:

- Brian Wirth, University of Tennessee
- Katrin Heitmann, Argonne National Laboratory
- David Dixon, University of Alabama
- Joe Oefelein, Sandia National Laboratories

Elected to 3 year terms:

- John Turner, ORNL
- Mike Zingale, vice-chair, State University of New York–Stony Brook
- Rangan Sukumar, ORNL

OUG members met once a month by conference call and webinar to discuss OLCF news, resources, policies, and timely HPC tutorials and techniques. Typically, between 25 and 40 members participate in the monthly calls; the highest number was 85 for a conference call featuring a Summit overview. The OUG Executive Board meets shortly before or after the monthly call to provide the OLCF with in-depth feedback and guidance on topics such as training, facility resources, and policies. The board has also organized working groups to supply user feedback on specific topics, such as the annual user meeting. The OUG meets in person at least once a year, usually at the annual onsite OLCF user meeting.

1.4.2.2 Training, Education, and Workshops

Workshops, user conference calls, training events, and seminars are integral components of both user assistance and outreach. While training can obviate difficulties with conducting science on such large-scale systems, training events can also serve to engage both the public and the user community. In addition to training users how to use the resources available at the OLCF, the training program focused on two additional areas: staff development and collaboration with other HPC centers.

See Appendix B for a summary of training events. Select notable 2014 events are highlighted below.

Getting Started Workshops

Entering 2014 and a new round of INCITE and ALCC projects, OLCF offered the “Getting Started” workshop geared at relaying essentials to new users. The 3 hour training event was held twice, in January and July. The first Getting Started workshop was also delivered as a webinar for users unable to visit ORNL. The workshops were intended to give users a quick start at using the center and a good overview of common pitfalls new users face when using leadership computing resources. User assistance staff presented best practices for staying informed about OLCF computing resources via email lists, online status indicators, and the @OLCFStatus Twitter feed; installing and updating software; resolving common runtime errors; using data management resources for storing and transferring data; and downloading and running debugging tools like Allinea’s Distributed Debugging Tool and Cray’s performance analysis tool CrayPat-lite.

First OpenACC Hackathon

The OLCF invited scientists and GPU developers to spend an intense week together. The inaugural Hackathon convened from October 27–31, 2014, in Knoxville, Tennessee. Six teams of 4–12 users were each paired with two expert mentors from Cray, NVIDIA, the Portland Group, the CSCS, and the OLCF Scientific Computing and UAO groups. The event hosted 45 people over 5 days. “Without our mentors, it would have taken us several months to do what we were able to achieve in a week,” said research scientist Rangan Sukumar of ORNL’s Computational Data Analytics Group. “Having somebody that really knows the best practices showing us how the compilers come together and explaining aspects of how to write high-performance code with OpenACC was invaluable.” Since the Hackathon, vendor partners who participated have gathered compiler bugs observed during the event that might not have been picked up or submitted by users through traditional means, and three teams have requested projects with OLCF to continue their application development on Titan.

Joint Facilities User Forum

In partnership with seven national laboratories, OLCF co-organized a 3 day workshop to discuss approaches to handling data and the future of data-driven scientific discovery. The Joint Facilities User Forum on Data-Intensive Computing brought together users and HPC center staff from June 16–18, 2014, in Oakland, California, to discuss advances in managing, analyzing, and visualizing data; successes, failures, best practices and lessons learned; and practical ways to work with data. Several OLCF staff members took part in the workshop: Fernanda Foertter and Ashley Barker served on the program committee and as session chairs during the first and second days’ meetings, respectively. OLCF’s Chris Fuson presented on transferring large data sets over the wide area network to laboratories and facilities, and Norbert Podhorszki gave a presentation on using ADIOS. ORNL’s Rangan Sukumar delivered a presentation on machine learning for data-driven discovery, and Stuart Campbell gave a presentation titled “Accelerating Scientific Discovery at the Spallation Neutron Source,” an ORNL neutron scattering facility that generates large amounts of data that can be more rapidly analyzed with leadership computing resources. Other ORNL participants included Clay England, Doug Fuller, and John Harney.

2014 OLCF User Meeting

The OLCF retooled the User Meeting in 2014. Previously, the user meeting was very training oriented and covered topics such as how to get on the system, basic debugging, and compiling code. In 2014 the OLCF included project PIs and software developers. Although users could still receive training and skills preparation, the event was much more research oriented. The basic format involved a keynote speaker in the morning followed by an afternoon session that focused on the scientific progress users had made on Titan. A total of 26 PIs and 104 total users attended this year's meeting.

SC14 Training Workshop

The OLCF's Fernanda Foertter and a group of colleagues at other HPC centers led the first "Best Practices in HPC Training" workshop at the SC14 supercomputing conference in New Orleans. About 50 people attended the workshop. The format included twelve 10 minute talks that discussed workshop development, training courses, survey options for polling users, and more.

1.4.2.3 Training and Outreach Activities for the Next Generation

The OLCF maintains a broad program of collaborations, internships, and fellowships for young researchers. Twenty-nine faculty members, student interns, and postdoctoral researchers were supported from January 1, 2013, through December 31, 2013.

Examples of user engagement and outreach to the next generation of HPC users include the following.

- Sean McDaniel, an intern during the summer of 2014, spent his time creating smaller, faster-running kernels that mimic real-world applications that can be used to test supercomputers and ensure the hardware and software are always robust. McDaniel took home the SC14 supercomputing conference Undergraduate Research Award for his poster "Comparing Decoupled I/O Kernels Versus Real Traces in the I/O Analysis of the HACC Scientific Applications on Large-Scale Systems." McDaniel worked under the direction of mentors Sarp Oral and Hai Ah Nam from the OLCF. He will be pursuing a PhD in computer science at the University of Delaware this spring, where he plans to specialize in HPC.

- Matthew Donovan, a senior in computer/electrical engineering at Tennessee Tech University, is a three-time veteran of the ORNL internship program. He spent the summer of 2014 working under the guidance of the OLCF's Adam Simpson to design and build a supplementary physics experiment for Tiny Titan. The experiment uses a 6 foot metal contraption that records a ball dropping and sends the information to Tiny Titan. Students then plug the coordinates into an algorithm to show the gravity constant. Donovan wrote the program from end to end and made it fully customizable. The rising college senior said that his internship confirmed that computing was really what he wanted to pursue. "The first year I did some business stuff. The second year I did all computer science work. Being my third year, I'm actually doing the stuff I *really* enjoy. It's developing very simple GUIs. I got to use my building skills. It's definitely made me realize that I enjoy the mixture between the hardware and the software."

- Chris Martin, a sophomore at the University of Tennessee, was able to combine his graphic design experience with his interest in computer science during his internship. Martin worked on a web front end for power and water usage effectiveness coupled with humidity and temperature.

- On August 19, 2014, Adam Simpson and Suzanne Parete-Koon represented OLCF at a career mentoring event, Bazinga, hosted by the University of Tennessee's Women in Electrical Engineering and Computer Science (EECS) chapter. The event was designed to encourage students to learn about EECS careers. Simpson and Parete-Koon demonstrated parallel programming with Tiny Titan and used the packed audience as a human supercomputer that estimated the value of pi using MC methods.

- The OLCF partnered with the National Institute of Computational Sciences (NICS) to deliver the annual Crash Course in Supercomputing on June 20, 2014. The Crash Course in Supercomputing is designed to give participants an overview of HPC concepts and techniques. The target audience is summer interns and those with no prior HPC exposure. A total of 120 people participated in the event.
- The OLCF outreach staff participated in the recruitment fair at the Grace Hopper Celebration of Women in Computing conference held October 8–10, 2014, in Phoenix and presented a poster on the OLCF User Facility to conference participants.
- For the past 5 years, the OLCF's Dustin Leverman has participated in the Student Cluster Competition at the annual supercomputing conference as a Supercomputing committee member. For the past 2 years, he has served as the committee chair. The goal of the competition is to expose students to the field of HPC. By participating in this event, the OLCF has the opportunity to engage with students interested in HPC. As a result of its participation, OLCF has hired two former participants in the competition into the facility, including Leverman himself.

The OLCF Petascale Initiatives Postdoctoral Research Program, funded by the American Recovery and Reinvestment Act of 2009, was completed on August 14, 2014. The purpose of the program was to foster innovation in computational science. During the project tenure from 2009 to 2014, the OLCF hired 16 postdoctoral researchers, 8 of whom are still currently working at ORNL in various capacities. In 2014, OLCF postdoctoral associates made the following contributions to both computational and domain sciences:

- Materials Science: Ying Wai Li studied the magnetic properties of magnetic materials at finite temperature using first principles calculations, an important topic for technological applications such as magnetic data storage and permanent magnets in electric motors and generators. (See also Section 1.4.1.2).
- Chemical Sciences: Mingyang Chen developed computationally efficient multilevel approaches to treat solvated molecular systems using density functional theory and higher-level electronic structure methods. Dmitryo Lyakh developed highly scalable implementations of electronic structure applications at the coupled-cluster level of theory for complex excited state and multi-reference state molecular systems. (See also Section 1.4.1.2).
- Computational Biology: Trung Nguyen improved the performance of molecular dynamics applications so that the time-to-solution for very large problems was significantly reduced. Jan-Michael Carrillo used coarse-grained molecular dynamics simulations to understand the behavior and interactions of polymers at surfaces and interfaces, which includes interactions of charged or neutral polymers with nanoparticles and biomaterials. Juanjuan Chai developed a novel automated bioinformatics tool to cluster proteins and identify the homologous groups, an important task for protein function annotation and reconstruction of evolutionary history.
- Nuclear Physics: Supada Laosooksathit collaborated with domain scientists from the CERN ALICE (A Large Ion Collider Experiment) project that measures properties of the quark gluon plasma, optimizing their computational simulation software for multithreaded and accelerated supercomputers such as the OLCF's Titan.

In 2015, the OLCF will initiate the OLCF Distinguished Postdoctoral Fellowship Program in support of the Computational Scientists for Energy, the Environment and National Security (CSEEN) initiative. These postdoctoral associates will be critical in the OLCF's efforts to prepare applications for the Summit architecture and ensure portability of applications between Summit and the National Energy Research Scientific Computing Center (NERSC) and ALCF architectures (Section 1.5).

Tiny Titan

The OLCF outreach team identified a way to improve how we answer one of the most commonly asked questions about how supercomputers work: What does it mean for computers to collaborate? To that end, members of the OLCF designed a public engagement tool called "Tiny Titan." Tiny Titan is a portable cluster of nine Raspberry Pi computers running an interactive fluid simulation. Each computer features a bright LED, which is color-coded to match particles in the simulation. Clearly identifying which computer powers each region of the simulation gives visitors a concrete grasp of how supercomputers divide problems into smaller components. Showcasing how particles change color when they pass between computers gives a strong visual example of network communication in action. Illuminating these concepts provides additional opportunities to engage the public, excite students, and bolster support for HPC among public officials (Figures 1.7 and 1.8).

Tiny Titan has represented the OLCF at regional schools, local and national news outlets, the US House of Representatives' 2014 National User Facilities Organization meeting, and DOE's National Laboratory Day. It has been incorporated into the overlook area of the OLCF where tours are conducted regularly. The tool has helped make Titan and the power of supercomputers more understandable to the numerous visitors to the facility each year.

Figure 1.7. Secretary of Energy Ernest Moniz (left) Texas Representative Eddie Bernice Johnson (D, center), and Alaska Senator Lisa Murkowski (R, right) visit Tiny Titan during the 2014 National Laboratory Day event.

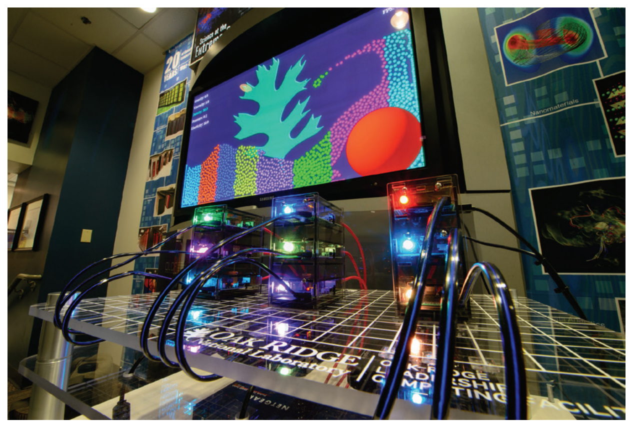

Figure 1.8. Tiny Titan on display in the OLCF overlook area.

1.4.2.4 Communications

The OLCF continues its rapid and widespread dissemination of the scientific impact of user research. For example, in 2014 the OLCF Outreach staff worked closely with the authors timing the release of a highlight with the publication of a *Nature Physics* paper[4] resulted in a great deal of media interest. The highlight, as well as the publication itself, was picked up by a number of media outlets including: Sciencedaily.com, Scienceseeker.org, Biomedox.com and Einnews.com.

Relationship-building with media outlet is actively promoted. In April, Erika Engelhaupt, an editor and blogger at *Science News,* toured EVEREST and met with Jack Wells, director of science. *Science News* is a bi-weekly print- and online- magazine that reports on basic science research. In late July the center hosted HPCwire editor, Nicole Hemsoth, for a comprehensive day of touring, meetings with key

[4] Yin, Z., Haule, K., Kotliar, G. (2014). Spin dynamics and orbital-antiphase pairing symmetry in iron-based superconductors. *Nature Physics*, 10(11) 845-850. doi: 10.1038/NPHYS3116.

OLCF officials and conversation with ORNL communications. In September the OLCF hosted a visit by Janet Ivey, host of the PBS children's science show, Janet's Planet.

The OLCF Outreach team played a key part in preparing the communications and public relations materials for the Summit announcement. The Outreach team initiated and hosted biweekly video conference calls with integrated team members from OLCF, ORNL, IBM, NVIDIA, Mellanox, and Lawrence Livermore National Laboratory to coordinate communications and plan media relations efforts for the Summit launch.

Further, the OLCF team worked to produce several communication assets, including an institutional press release, the Summit website, a Summit video detailing the research possibilities of the new system, a CORAL fact sheet, a Summit poster for SC14 supercomputing conference, and a Summit fact sheet. The OLCF team also worked with other integrated team members to consult on vendor press releases and vendor-specific videos. The official announcement was made on November 14, 2014, and the OLCF team launched the Summit website and video on the same day. Within 2 weeks, the Summit website received 3,595 page views and the Summit announcement was featured in 99 different news articles.

Several OLCF staff members gave talks on Summit at SC14, the ASCR Advisory Council (ASCAC), and other meetings. The OLCF Outreach team works to engage new and next-generation users and showcases OLCF research through strategic communication activities such as tours, highlights, fact sheets, posters, snapshots, the OLCF website, and center publications (see Appendix C). The Outreach team was responsible for the creation of 54 highlights and for more than 130 total outreach products in 2014. Throughout the year, the OLCF provides tours to groups of visitors who range from middle-school students through senior-level government officials. The center gave tours for 283 groups in 2014.

UAO is responsible for the OLCF website, which received 341,598 page views in 2014. The top two visited pages on the OLCF website in 2014 included the Titan resource page and the Titan User Guide. The OLCF User Guides were the highest-rated aspect of the OLCF website with a mean rating of 4.5/5.0, and 94% of users indicated they were satisfied or highly satisfied with the User Guides. UAO continued to add to the content to the Titan User Guide, which received 13,843 total page views and 11,191 unique page views in 2014. The OLCF Getting Started Guide, the OLCF Support landing page, and the Titan computer resource page rounded out the top five most visited places on the OLCF website. The tutorials were once again popular in 2014 with two of the tutorials being among the ten most visited pages on the OLCF website with over 10,000 combined total page views.

1.5 LOOKING FORWARD

1.5.1 Application Portability

The portability of scientific and engineering applications is increasingly important to the users of ASCR computing facilities. Application developers target a wide range of architectures for their user base. Moreover, the lifespans of applications are much longer than that of any computer architecture, which further illustrates the need for applications to be developed for changing architectures. Further, many of OLCF's PIs have allocations at multiple facilities, and their science campaigns are greatly facilitated by having portable applications.

Recognizing the responsibility to contribute to making applications both architecturally and performance portable, at least between the architectures in the ASCR computing facilities, Tjerk Straatsma, leader for the OLCF Scientific Computing group; Katie Antypas, head of the Services Department at NERSC; and Timothy Williams, ALCF Principal Project Specialist, are taking the initiative to coordinate application readiness activities, develop a strategy to provide guidance and tools encouraging application development that is portable across different architectures, and put mechanisms and resource allocations in place that enable Early Science teams to test and run applications on different architectures. This collaboration does not simply address the next-generation systems that will be coming to ALCF, NERSC, and OLCF. Using appropriate abstractions to get portability and performance on these pre-exascale systems provides a path to be continued toward exascale.

1.5.2 Center for Accelerated Application Readiness

Simultaneously with the announcement of OLCF's next computer system, Summit, the OLCF Center for Accelerated Application Readiness (CAAR) issued a call for proposals to establish eight new partnership projects. These will be to prepare computational science or engineering applications for highly effective use on Summit, which will become available to users in 2018. Each of these projects will consist of a 3 year Application Readiness phase in which the code refactoring and porting work will take place and an Early Science phase for tuning the code to the Summit architecture and demonstrating the application through a scientific grand-challenge project. The teams will be partnerships between the core developers of the applications and OLCF staff assigned to the projects; and they will receive technical support from the IBM/NVIDIA Center of Excellence, as well as access to computational resources that include Titan, early delivery systems, and Summit as they become available.

In addition to delivery of scalable applications optimized for Summit's architecture, performance portability to other architectures is an important consideration. The CAAR readiness efforts will be carried out, particularly for shared applications, in collaboration with the ALCF and NERSC to enhance application portability across their respective architectures.

The partnership teams will develop and execute a technical application porting and performance improvement plan with reviewable milestones for the Application Readiness phase of the project. Key components of this plan include code optimization, refactoring, testing, and profiling of the applications. This plan will be developed according to a management plan that outlines the responsibilities of the core application developers and of the Scientific Computing staff member assigned by the OLCF and the OLCF postdoctoral fellow.

Business Results

HIGH PERFORMANCE COMPUTING FACILITY
2014 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2015

2. BUSINESS RESULTS

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

OLCF RESPONSE: Yes. The OLCF provides a series of highly capable and reliable systems for the user community. The 2014 reporting period includes full calendar year production periods for the HPC resources Cray XK7 (Titan), the Cray XC30 (Eos), and the Lustre file systems (Atlas). The effectiveness with which these resources were delivered is demonstrated by the business result metrics, which were met or exceeded in all cases. The OLCF team managed policies and job-scheduling priorities that maximized access to these production systems. Again in 2014, the OLCF delivered all of the compute hours committed to the three major allocation programs: INCITE, ALCC, and DD. OLCF leadership computational resources support scientific research through production simulation across many scientific domains, providing the key computing and data resources that are critical to their success.

2.1 BUSINESS RESULTS SUMMARY

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

2.2 CRAY XK7 (TITAN) RESOURCE SUMMARY

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

2.3 CRAY XC30 (EOS) RESOURCE SUMMARY

Eos is a four-cabinet Cray XC30. The system, with 744 Intel Xeon E5-2670 compute nodes and 47.6 terabybtes (TB) of memory, provides the OLCF user community with a substantive large-memory-per-node computing platform. The Eos nodes are connected by Cray's Aries interconnect in a network topology called "Dragonfly." All INCITE users are automatically granted access to the XC30.

2.4 LUSTRE FILE SYSTEMS (ATLAS) RESOURCE SUMMARY

In September 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 petabytes (PB) and block-level performance of more than 1.3 TB/second. The /atlas file systems are the default high-performance file systems for all compute systems.

The previous generation Lustre file systems, collectively the four /widow file systems, were decommissioned during the prior reporting period (2013).

2.5 DATA ANALYSIS AND VISUALIZATION CLUSTER (RHEA) RESOURCE SUMMARY

Rhea is a 512-node commodity-type Linux cluster. The primary purpose of Rhea is to provide a conduit for large-scale scientific discovery via pre- and post-processing of simulation data generated on Titan. Users with accounts on INCITE- or ALCC-supported projects are automatically given an account on Rhea. Director’s Discretion (DD) projects may request access to Rhea. Each of Rhea’s nodes contains two 8-core 2.0 GHz Intel Xeon processors with hyper-threading and 64 GB of main memory. Rhea is connected to the OLCF’s 32 PB high performance Lustre filesystem, Atlas.

2.6 HIGH PERFORMANCE STORAGE SYSTEM RESOURCE SUMMARY

The OLCF provides a long-term storage archive system based on the High Performance Storage System (HPSS) software product co-developed by IBM, Los Alamos National Laboratory, Sandia, Lawrence Livermore National Laboratory, Lawrence Berkeley National Laboratory, and ORNL. The ORNL HPSS instance is currently over 40 PB in size providing up to 40 gigabits (Gb) per second of read and write performance. The archive has ingested over 150 TB in a single day several times in the last year; the previous daily maximum was just over 50 TB.

The archive is built from hardware from Dell, Hewlett Packard, Brocade, NetApp, DataDirect Networks, and Oracle. There is currently over 2 PB of disk cache for bursting data into the archive at a rate of 20 gigabytes (GB) per second, and over 26 GB/s of read/write bandwidth to 154 Oracle T10K tape drives. There are 6 Oracle SL8500 tape libraries for tape archival storage that each have 10,100 slots; the archive’s maximum size is over 500 PB using these libraries.

2.7 VISUALIZATION RESOURCE SUMMARY

The Exploratory Visualization Environment for Research in Science and Technology (EVEREST) has three computing systems and two separate state-of-the-art visualization display walls. The primary display wall spans 30.5 × 8.5 feet and consists of 18 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 display provides 37 million pixels; the stereoscopic design creates depth in an image for a 3D effect.

2.8 OLCF COMPUTATIONAL RESOURCE SUMMARY

The OLCF provided the Titan and Eos computational resources in 2014 (see Table 2.1).

2.8.1 OLCF HPC Resource Production Schedule

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

Table 2.1. OLCF production computer systems, 2014

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 + 14 streaming multiprocessor (SM) 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	744	2 × 8-core SMP	64 GB DDR3-1600 per node; 47,600 GB DDR3 aggregate	Aries (Dragonfly)

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

System	Type	Production date[a]	Performance end date[b]	Notes
Atlas	Lustre parallel file system	October 3, 2013	Null	Delivered as two separate file systems, /atlas1 and /atlas2. 32 PB capacity
Eos	Cray XC30	October 3, 2013	Null	Production with 744 Intel E5, 2,670 nodes.
Titan	Cray XK7	May 31, 2013	Null	Production with 18,688 hybrid CPU-GPU nodes (AMD 6274/NVIDIA K20X)
JaguarPF	Cray XK6	September 18, 2012	October 7, 2012	Production at 240,000 cores until September 18, when partition size was reduced to 120,000 AMD Opteron cores. Additional Kepler installation. TitanDev access terminated
JaguarPF	Cray XK6	February 13, 2012	September 12, 2012	Full production until September 12, when partition size was reduced to 240,000 AMD Opteron cores. Beginning of Kepler installation
JaguarPF	Cray XK6	February 2, 2012	February 13, 2012	Stability test. Restricted user access. 299,008 AMD Opteron 6274 cores. Includes 960-node Fermi-equipped partition
JaguarPF	Cray XK6	January 5, 2012	February 1, 2012	Acceptance. No general access 299,008 AMD Opteron cores
JaguarPF	Cray XK6	December 12, 2011	January 4, 2012	142,848 AMD Opteron cores
JaguarPF	Cray XT5	October 17, 2011	December 11, 2011	117,120 AMD Opteron cores
JaguarPF	Cray XT5	October 10, 2011	October 16, 2011	162,240 AMD Opteron cores
JaguarPF	Cray XT5	September 25, 2009	October 9, 2011	224,256 AMD Opteron cores
JaguarPF	Cray XT5	August 19, 2008	July 28, 2009	151,000 AMD Opteron cores

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

2.8.2 Business Results Snapshot

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

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 90% and the overall availability (OA) target is 80%. For year 2, the SA target for an HPC compute resource remains 90% and the OA target increases to 85%. For year 3 through the end of life for the associated compute resource, the SA target for an HPC compute resource remains 90% and the OA target increases to 90%. SA targets are thus described as 90% for all production years. OA targets are 80%/85%/90% for the first, second, and subsequent years.

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

The discussion of SA targets as a single value (not increasing), regardless of system year, is new for 2014, and will continue in 2015.

The Atlas, Titan, and Eos systems all celebrated their 1 year production anniversaries in 2014. The reported results for each system measure the full 2014 calendar year and intentionally do not reflect the partial results to their respective production anniversaries. In all cases, the OLCF result exceeded the 2014 target for the accompanying metric.

Because an outage that may define the SA, OA, mean time to interrupt (MTTI), or mean time to failure (MTTF), may occur outside the reporting period, the data reflected here artificially assume calculation boundaries of January 1, 2014 0:00 and January 1, 2015 0:00.

Table 2.3. OLCF business results summary for HPC systems

	Measurement	2013 target	2013 actual	2014 target	2014 actual
Cray XK7 (Titan)	Scheduled availability	85%	98.70%	90%	99.59%
	Overall availability	80%	93.82%	85%	95.80%
	MTTI (hours)	NAM	173.47	NAM	310.83
	MTTF (hours)	NAM	467.94	NAM	1,246.54
	Total usage	NAM	89.93%	NAM	89.63%
	Core-hours used*	NAM	2,640,915,296	NAM	4,217,292,935
	Core-hours available	NAM	2,936,516,529	NAM	4,705,171,200
	Capability usage				
	INCITE projects	NAM	60.31%	NAM	69.44%
	All projects	30%	59.38%	35%	62.58%

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

Table 2.4. OLCF business results summary for Eos

	Measurement	2013 target	2013 actual	2014 target	2014 actual
Cray XC30 (Eos)	Scheduled availability	NIP	NIP	NAM	99.78%
	Overall availability	NIP	NIP	NAM	97.06%
	MTTI (hours)	NIP	NIP	NAM	340.09
	MTTF (hours)	NIP	NIP	NAM	1,748.20

MTTF = Mean time to failure.
MTTI = Mean time to interrupt.

NAM = Not a metric. No defined metric or target exists for this system. Data provided as reference only.
NIP = Not in production. This system was not available as a production resource.

Table 2.5. OLCF business results summary for HPSS

	Measurement	2013 target	2013 actual	2014 target	2014 actual
HPSS	Scheduled availability	95%	99.99%	95%	99.83%
	Overall availability	90%	97.60%	90%	98.56%
	MTTI (hours)	NAM	450.00	NAM	297.71
	MTTF (hours)	NAM	2,919.78	NAM	546.56

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

Table 2.6. OLCF business results summary for the external Lustre file systems

	Measurement	2013 target	2013 actual	2014 target	2014 actual
/atlas1	Scheduled availability	85%	99.67%	90%	99.73%
	Overall availability	80%	97.02%	90%	99.15%
	MTTI (hours)	NAM	232.97	NAM	377.62
	MTTF (hours)	NAM	430.82	NAM	624.02
/atlas2	Scheduled availability	85%	98.34%	90%	99.50%
	Overall availability	80%	93.62%	90%	98.59%
	MTTI (hours)	NAM	224.79	NAM	411.26
	MTTF (hours)	NAM	425.35	NAM	792.41

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

2.9 RESOURCE AVAILABILITY

Details of the definitions and formulas describing scheduled availability, overall availability, MTTI, and MTTF are provided in Appendix D.

2.9.1 Scheduled Availability

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

As shown in Table 2.7, the OLCF has exceeded the scheduled availability targets for the facility's computational resources for 2013 and 2014.

Table 2.7. OLCF business results summary: Scheduled availability

	System	2013 target	2013 actual	2014 target	2014 actual
Scheduled Availability	Cray XK7	85%	98.70%	90%	99.59%
	Cray XC30	NIP	NIP	NAM	99.78%
	HPSS	95%	99.99%	95%	99.83%
	/atlas1	85%	99.67%	90%	99.73%
	/atlas2	85%	98.34%	90%	99.50%

NIP = Not in production. This system was not available as a production resource.
NAM = Not a metric. No defined metric or target exists for this system. Data provided as reference only.

2.9.1.1 Assessing Impacts to Scheduled Availability

The operational posture for the Cray XK7 system contains a regularly scheduled weekly preventative maintenance (PM) period. PM is exercised only with the concurrence of the Cray hardware and software teams and the HPC Operations team. Typical PM includes software updates, application of field notices, and hardware maintenance to replace failed components. Without concurrence, the systems are allowed to continue in operation.

In 2014, OLCF staff executed scheduled maintenance on the Cray XK7 a total of 20 times, associated with hardware maintenance, SMW and CLE firmware/software upgrades, field notices, Lustre software stack testing, and integration with the /atlas file systems. Seven unscheduled outages were reported in 2014, including external power quality events, XDP failures (a component of the XK7 cooling system), and a small number of hardware or software failures that could not be recovered. Similarly, OLCF performed scheduled maintenance on Eos 20 times in 2014, with five unscheduled outages.

2.9.2 Overall Availability

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

As shown in Table 2.8, the OLCF has exceeded the overall availability targets for the facility's computational resources for 2013 and 2014.

Table 2.8. OLCF business results summary: Overall availability

	System	2013 target	2013 actual	2014 target	2014 actual
Overall Availability	Cray XK7	80%	93.82%	85%	95.80%
	Cray XC30	NIP	NIP	NAM	97.06
	HPSS	90%	97.60%	90%	98.56%
	/atlas1	80%	97.02%	90%	99.15%
	/atlas2	80%	93.62%	90%	98.59%

NIP = Not in production. This system was not available as a production resource.
NAM = Not a metric. No defined metric or target exists for this system. Data provided as reference only.

2.9.3 Mean Time to Interrupt

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

The MTTI summary is shown in Table 2.9.

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

	System	2013 target	2013 actual	2014 target	2014 actual
MTTI (hours)	Cray XK7	NAM	173.47	NAM	310.83
	Cray XC30	NIP	NIP	NAM	340.09
	HPSS	NAM	450.00	NAM	297.71
	/atlas1	NAM	232.97	NAM	377.62
	/atlas2	NAM	224.79	NAM	411.26

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

NIP = Not in production. This system was not available as a production resource.

2.9.4 Mean Time to Failure

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

The MTTF summary is shown in Table 2.10.

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

	System	2013 target	2013 actual	2014 target	2014 actual
MTTF (hours)	Cray XK7	NAM	467.94	NAM	1,246.54
	Cray XC30	NIP	NIP	NAM	1,748.20
	HPSS	NAM	2,919.78	NAM	546.56
	/atlas1	NAM	430.82	NAM	624.02
	/atlas2	NAM	425.35	NAM	792.41

NAM = Not a metric. No defined metric nor target exists for this system. Data provided as reference only.
NIP = Not in production. This system was not available as a production resource.

2.10 RESOURCE UTILIZATION

2014 Operational Assessment Guidance

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

Observation: The numbers that are reported for the Cray XK7 resource are Titan core-hours, where a single Titan node-hour comprises 16 AMD Opteron core-hours and 14 NVIDIA Kepler SM-hours. The OLCF refers to the combination of these traditional core-hours and SM-hours as "Titan core-hours", denoting that they are the product of a hybrid node architecture. System production requires the use of node-hours, where all resources of both the CPU and GPU comprising a single node are aggregated. The use of node-hours impacts all scheduling and accounting activities. Users describe all job submission activity in node-hours as the smallest unit. In addition, allocation programs are moving to a node-hour basis, where the conversion from node-hours to Titan core-hours remains a straight-forward calculation for comparison purposes. Subsequent versions of this calculation are expected to continue to shift emphasis to node-hours, with the conversion to Titan core-hours available, to better reflect the specific systems at a particular facility.

2.10.1 Resource Utilization Snapshot

For the Cray XK7 for the operational assessment period January 1–December 31, 2014, 4,217,292,935 Titan core-hours were used from an available 4,705,171,200 Titan core-hours. The total system utilization for the Cray XK7 was 89.63%.

2.10.1.1 Resource Utilization Measurement Units

For the 2014 reporting period, system accounting measured the consumption of 1 Titan node-hour as 30 Titan core-hours (the combination of 16 Opteron core-hours and 14 NVIDIA Kepler SM-hours).

2014 INCITE allocations were provided to approved projects in terms of Titan core-hours; other 2014 programs were allocated in terms of Titan node-hours, where there remains a direct correlation from the

Titan node-hour to the Titan core-hour. This adjustment is consistent with the method required for allocating node-hours to applications. The job scheduler for the OLCF compute resources is Adaptive Computing's Moab, coupled to the Cray resource manager, Torque. Moab/Torque allocates resources at the granularity of a single node, not a compute core, regardless of the composition of that node.

2.10.2 Total System Utilization

2014 Operational Assessment Guidance

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

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

The Cray XK7 Titan system utilization is shown in Table 2.11. The measurement period is for 2014, irrespective of the prescribed allocation period of any single program. As an example, the INCITE allocation period follows a calendar year schedule. The ALCC program follows an allocation cycle that runs for 12 months beginning July 1 of each year. System utilization for 2014 is 89.63%.

Table 2.11. The 2014 Cray XK7 Titan utilization

Time period	CPU hours consumed	CPU hours available	Percent of available hours consumed
January	359,157,465	401,109,888	89.54%
February	292,876,541	371,470,720	78.84%
March	373,752,344	410,528,640	91.04%
April	341,863,865	366,191,360	93.36%
May	359,011,753	395,867,904	90.69%
June	330,876,336	373,797,376	88.52%
July	353,625,422	388,486,144	91.03%
August	365,888,574	408,519,680	89.56%
September	336,642,217	388,841,216	86.58%
October	368,141,896	405,155,840	90.86%
November	360,314,191	394,232,704	91.40%
December	375,142,331	400,969,728	93.56%
Total	4,217,292,935	4,705,171,200	89.63%

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 with high fidelity the consumption of Titan core-hours by program, project, user, and system. Figure 2.1 summarizes the Cray XK7 utilization by month and by program for all of 2014.

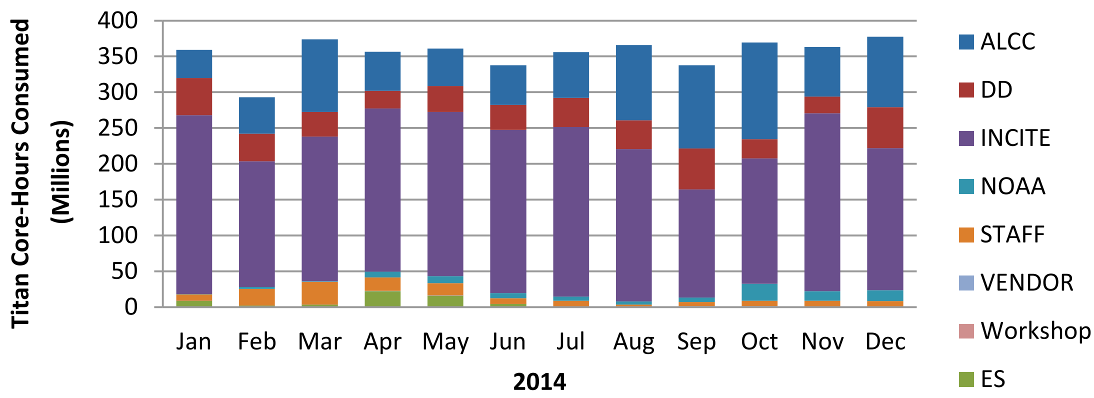

Figure 2.1. 2014 XK7 resource utilization—Titan core-hours by program.

2.10.2.1 Assessing Total System Utilization

The monthly utilization of available Titan core-hours is shown in Figure 2.2. Utilization of the available core-hours continues to be very high.

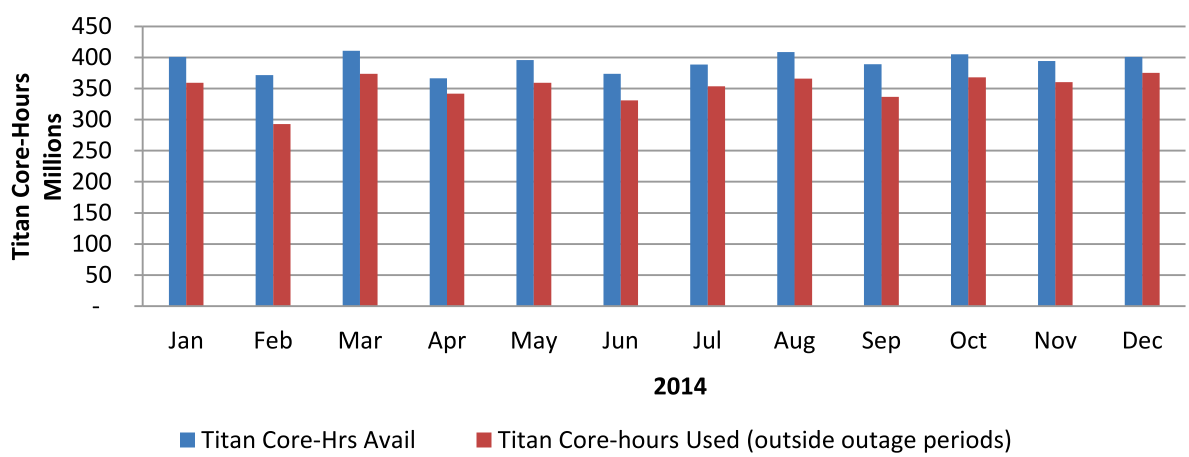

Figure 2.2. System Utilization on Titan in 2014.

2.10.2.2 Performance of the Allocated Programs

All allocation programs, including INCITE, ALCC, and DD, are aggressively monitored to ensure that projects within these allocation groups maintain appropriate consumption rates. The 2014 INCITE allocation program was the largest program in 2014 with a commitment for 2.25B Titan core-hours. The consumption of these allocation programs is shown in Table 2.12.

Non-renewed INCITE projects from 2013 continued running through January, 2014. The policy to permit an additional, final, month for completion was recognized as a best practice during a previous operational assessment review. It also serves to increase system utilization while new projects establish a more predictable consumption routine. 2014 non-renewing INCITE projects were allowed to continue to submit jobs in accordance with existing queue policy on both Titan and Eos through January 31, 2015. ALCC projects from the 2013 allocation period (ending June 30, 2014) were additionally granted extensions where appropriate.

Table 2.12. The 2014 Allocated program performance on the OLCF resources

Program	Allocation	Hours consumed	Percent of total
INCITE	2,250,000,000	2,562,284,100	62.12%
Titan		2,459,685,452	
Eos		102.598,648	
ALCC	Not applicable	941,437,489	22.83%
ALCC_2014[a]	380,600,000	397,398,975	
ALCC_2015[b]	1,153,400,000	544,038,514	
DD	Not applicable	464,653,171	11.27%
ES	Not applicable	57,743,654	1.40%
NOAA[5]	87,500,000	98,285,095	2.38%
Total		4,124,403,509	100.00%

[a]ALCC_2014 reflects the ALCC allocation period from July 1, 2013 to June 30, 2014.
[b]ALCC_2015 reflects the ALCC allocation period from July 1, 2014 to June 30, 2015.

2.11 CAPABILITY UTILIZATION

Capability usage defines the minimum number of nodes allocated to a particular job on the 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. As Titan reached production as a new system in 2013, the applicable metric was again 30%. This increased to 35% in subsequent production years, beginning May 31, 2014 (in the middle of the reporting period).

The OLCF Resource Utilization Council uses queue policy on the Cray systems to support delivery of this metric target, providing queues specifically for capability jobs with 24-hour wallclock times and increased priority.

The OLCF Capability Utilization Definition is summarized in Table 2.13.

Table 2.13. OLCF Capability Utilization Definition

System	Year 1		Subsequent years	
	Definition for capability	Capability metric	Definition for capability	Capability metric
Cray XK7 Titan	20% of available nodes on the largest system	30% of delivered hours	20% of available nodes on the largest system	35% of delivered hours

The OLCF continues to exceed expectations for capability usage of its HPC resources (Table 2.14). Keys to the growth of capability usage include the liaison role provided by the Scientific Computing Group members, who work hand-in-hand with users to port, tune, and scale code; and OLCF support of the application readiness efforts (CAAR), where staff 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.

5 NOAA = National Oceanic and Atmospheric Administration; period of performance reported: January 1, 2014–December 31, 2014.

Table 2.14. OLCF capability usage on the Cray XK7 system

	Leadership usage	CY 2013 target	CY 2013 actual	CY 2014 target	CY 2014 actual
Cray XK7	INCITE	NAM	60.31%	NAM	69.44%
	ALCC	NAM	43.54%	NAM	52.53%
	All Projects	30%	59.38%	30/35%	62.58%

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, 62.58%, was well above the 2014 target of 30/35%. This consumption varies modestly during the year, affected by factors including system availability and the progress by the various projects within their research. 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 that use >60% (11,250) or more 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. The distribution of the consumption of hours by capability jobs, by month, using stacked bars that show the contribution in both the 20–60% bin and the >60% bin, is shown in Figure 2.3.

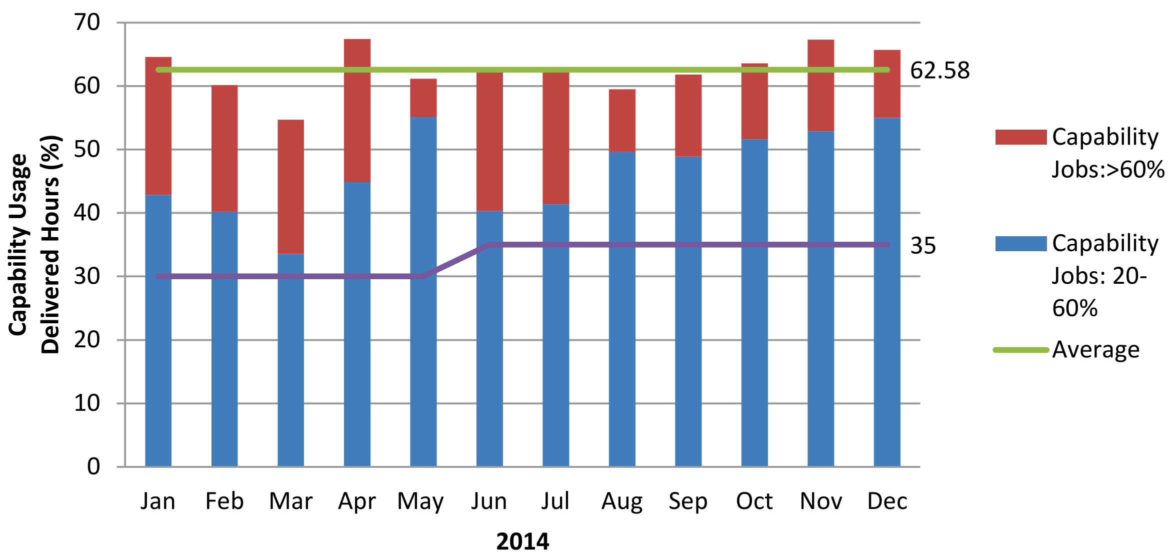

Figure 2.3. Effective scheduling policy enables capability usage.

Capability jobs are not restricted to the INCITE program. There are capability jobs across the Early Science, ALCC, and DD programs as well. The contribution to capability utilization by allocation program is shown in Figure 2.4, again using stacked bars that show the contribution in both the 20–60% bin and the >60% bin.

For each program type, the calculation describes the ratio of compute hours delivered by capability jobs to the compute hours delivered to all jobs.

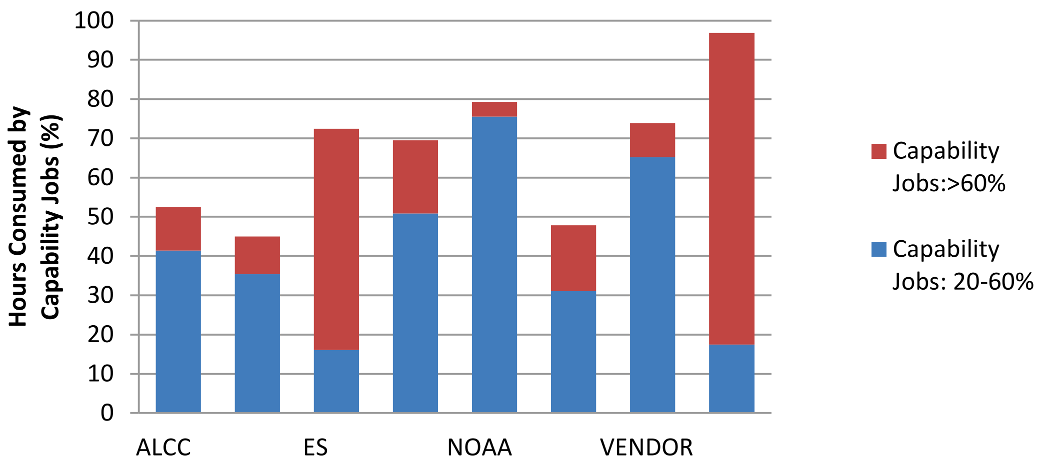

Figure 2.4. Capability usage by job size bins and project type.

2.12 GPU ACTIVITY

The heterogeneous architecture of Titan provides a key capability to users, allowing them to exploit a new hybrid compute node that contains both a CPU and the NVIDIA Kepler GPU. On any hybrid node, the GPU is an option for the user. There is no explicit requirement to use it.

In response to recommendations from previous Operational Assessment Reviews, the OLCF has worked closely with Cray and NVIDIA to develop techniques for estimating the usage of GPUs in Titan. The Facility uses two complementary techniques. Both methods have edge cases that may not record GPU usage. The two methods, ALTD—an ORNL developed tool that examines the libraries that are linked in an executable code, and Resource Utilization Reporting (RUR)—a application from Cray that retrieves hardware counter information from each GPU driver at the end of each job, are described below. In the technique descriptions and in the reports of utilization, the phrase "GPU enabled" refers to applications that use the node's GPU during execution.

2.12.1 Estimating GPU use via Automatic Library Tracking Database

The OLCF continues to use the Automatic Library Tracking Database (ALTD) method for tracking whether an application is employing the GPU. First deployed in 2012, ALTD actively monitors the compilation phase of individual applications and, at link time, creates a unique record for that application that contains a list of each of the libraries that were linked against that particular binary. When this application is executed via aprun, a new ALTD record is written to the database that contains the name of the executable, the batch job ID, and other supporting information. To determine whether a specific executable takes advantage of the GPU, OLCF examines whether the executed job, for which OLCF has all of the per-job scheduling information, was linked against an accelerator-specific library. For additional details about this method, see the *2013 OLCF Operational Assessment Report*.[6] Results for 2014 are shown in Figure 2.5. Due to a technical issue during a December 2014 database migration, three weeks of ALTD data were not gathered. This accounts for the low reported value of GPU use in December.

[6] http://info.ornl.gov/sites/publications/Files/Pub51404.pdf

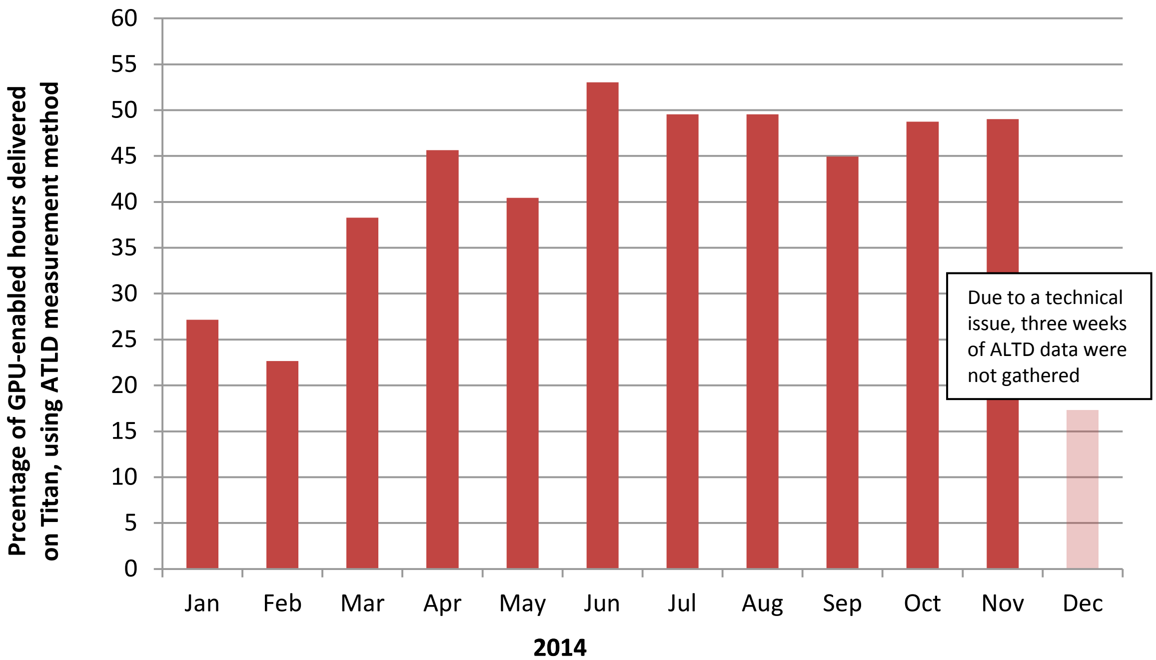

Figure 2.5. Tracking GPU usage on Titan using linked-library data from ALTD.

2.12.2 Estimating GPU activity via Cray's Resource Utilization Reporting

The method using ALTD makes a binary assumption: the application uses the GPU, or it does not. To understand the degree to which the GPU is being used requires more information directly from the GPU.

To meet this requirement, the OLCF actively engaged with both NVIDIA and Cray in 2012 and 2013 to determine what NVIDIA hardware counters could be exposed, and what system accounting features were available. The results of those discussions drove new features in the NVIDIA device driver and the accompanying API and library, and changes to the Cray RUR software, so that additional information about the GPU usage, on a per-job basis, is available at the conclusion of each normally terminated job. This initial development effort was completed in Spring 2014, with the release of an updated NVIDIA driver and with changes to the available Cray software stack, beginning with CLE4.2 UP02. The result provides an integrated database and toolset that combines accounting and GPU usage information in a single location. The data is available on a per-job basis. This information supports multiple efforts within the OLCF to identify and characterize jobs that can significantly leverage the power of the NVIDIA GPU. RUR results begin with data available in March 2014.

See Section 4.1.1 for details about how GPU activity is tracked and measured at the OLCF using RUR. The number of GPU-enabled compute jobs and the level of GPU activity for those jobs can vary widely across time, allocation program, and individual project. For calendar year 2014, approximately 40% of all delivered compute time on Titan for both INCITE and ALCC was by GPU-enabled applications. The DD program tends to trend slightly lower in the percentage of compute time that is GPU-enabled, since many DD projects are in the beginning phases of porting code to GPUs. In the second half of the measurement period, GPU-enabled INCITE applications were consistently responsible for about 50% of the delivered hours to those projects. Figure 2.6 shows the percentage of GPU-enabled compute time by month for the three primary allocations programs.

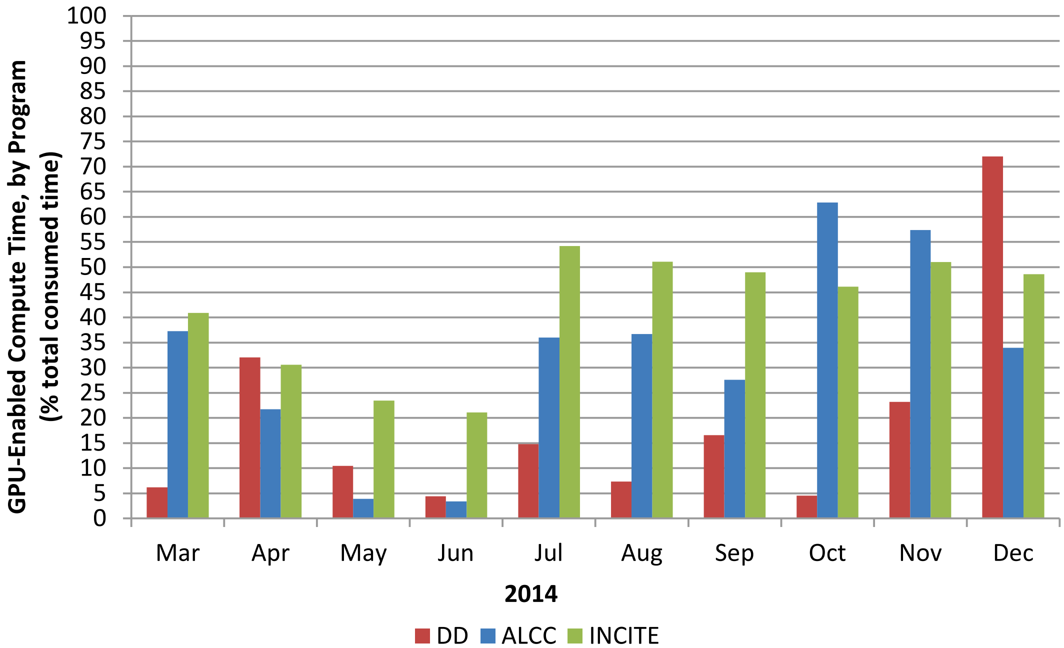

Figure 2.6. GPU-enabled jobs for the DD, ALCC, and INCITE user programs, RUR method.

For all GPU-enabled compute jobs on Titan in 2014, the reported average GPU activity of those applications was about 35% of the total runtime. This varies widely by application and by allocation program, ranging from as low as 9% for the DD program in December, to as high as 80% for the INCITE program in the same month. GPU activity varies based upon the algorithms used by individual projects. Figure 2.7 shows the percentage of GPU activity by month for exclusive-process GPU-enabled jobs for the major allocation programs. Data was compiled from the RUR tool data for the primary allocation programs.

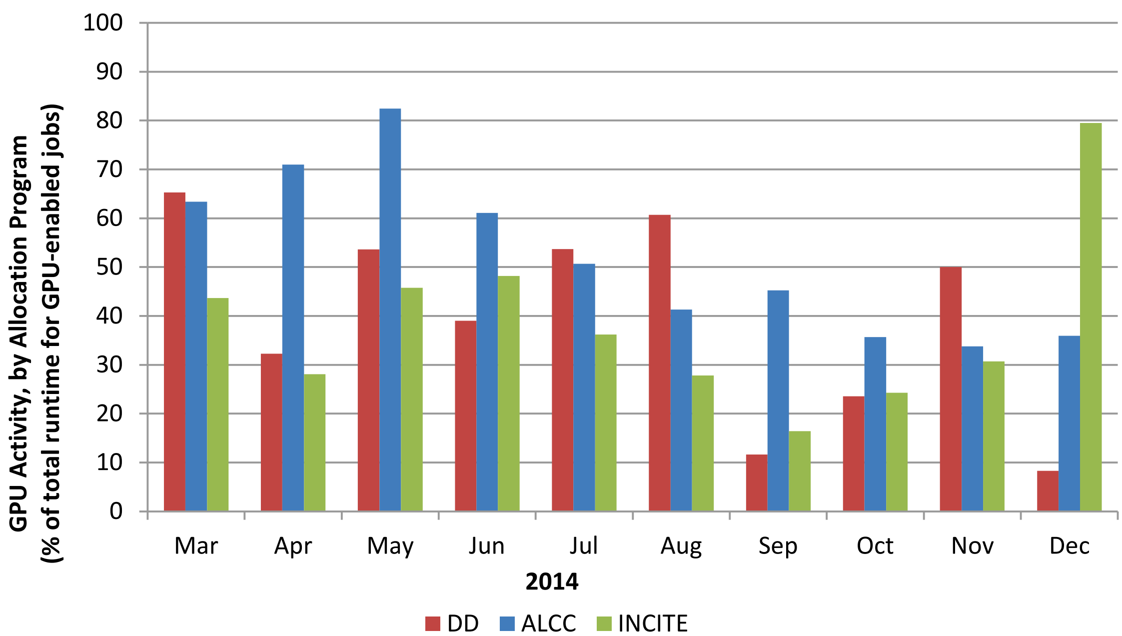

Figure 2.7. For jobs that employed the GPU, the percentage of total runtime when the GPU was active.

2.12.2.1 Case Study: Collaboration Leading to Improved GPU Use

In 2014, NCCS liaisons worked with INCITE project CPH103, "QMC Simulations Database for Predictive Modeling and Theory," and its PI David Ceperley with the goal of porting the quantum MC simulation code QMCPACK to Titan and optimizing the code to run on its many GPUs.

The result of this successful collaboration is evident in the GPU statistics. For most of 2014, the project ran no GPU-enabled jobs on Titan; in December, the reworked code was running many GPU-enabled jobs, and the GPU activity for those jobs was greater than 62% (not shown). Figure 2.8 shows the total percent of consumed time—42%—when GPUs were enabled. This data is from the subproject CPH103GEO, which employs QMCPACK. For comparison the total hours consumed is also shown.

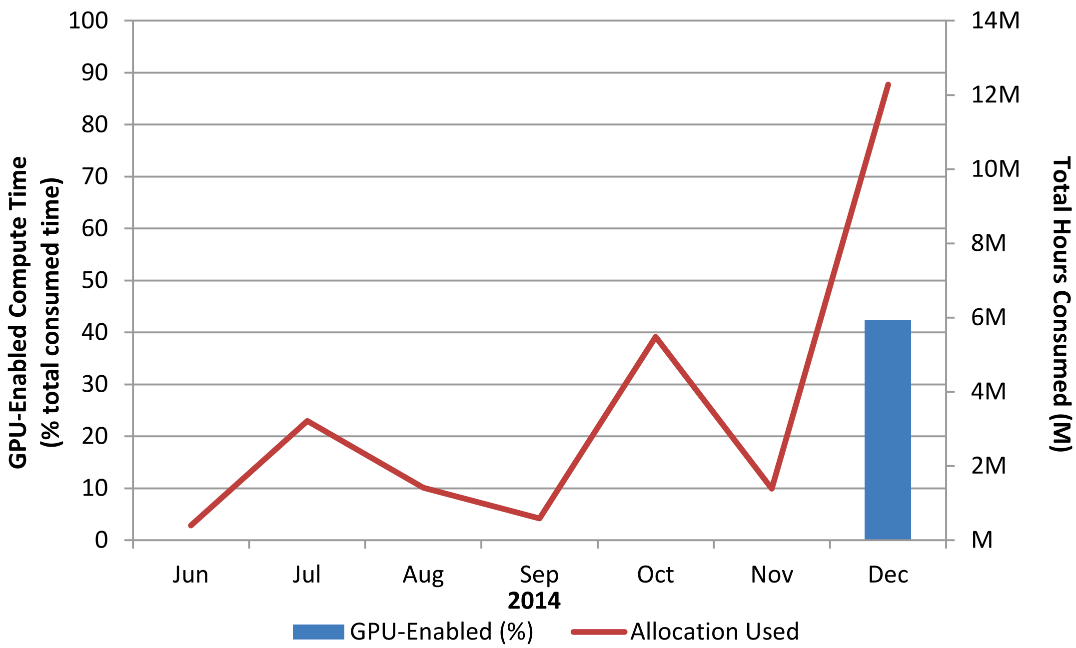

Figure 2.8. Percentage of GPU-enabled compute time (left scale) and total hours consumed (right scale) by month for CPH103GEO, which employs QMCPACK.

2.12.2.2 Case Study: Transitioning from Discretionary Program to INCITE

PI Michael Bussmann and his DD project APH005, "Laser-Wakefield Simulations Using PIConGPU," will continue in 2015 under an INCITE award after successfully demonstrating effective use of Titan's GPUs. Specifically, PIConGPU interleaves the data transfer between Titan's many distributed GPUs and the computation on each single GPU, so that the GPUs can execute the algorithmic steps continuously without interruption for communication.

After their initial porting efforts in March, subsequent jobs were primarily GPU-enabled. Nearly all of the compute jobs that ran in the latter half of the year were GPU-enabled (see Figure 2.9), and GPU activity (not shown) for those jobs was as high as 60%.

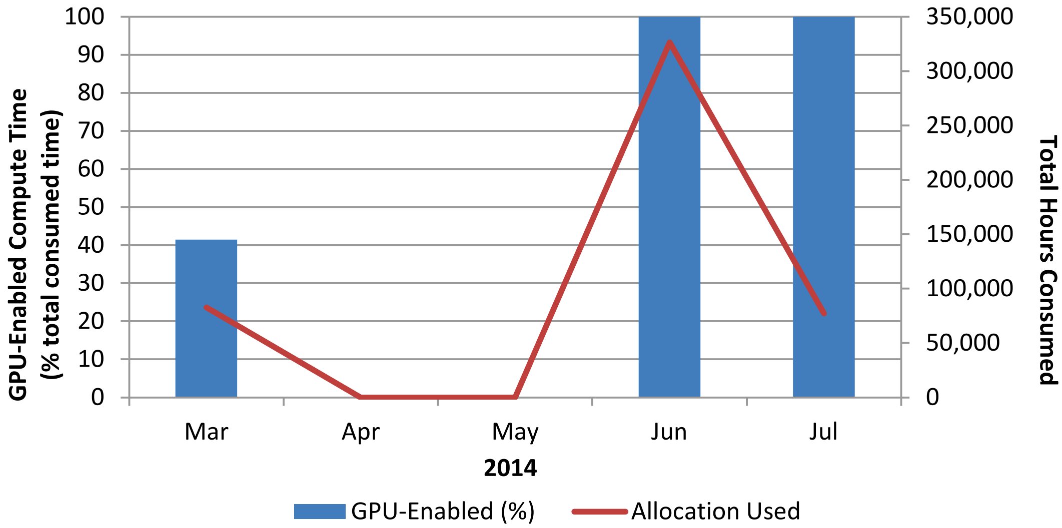

Figure 2.9. Percentage of GPU-enabled compute time (left scale) and total hours consumed (right scale) by month for APH005, which employs PIConGPU.

2.12.2.3 Case Study: Future Collaborations to Increase GPU Use

The 2014 GPU statistics efforts have enabled OLCF staff to more easily identify projects that can significantly benefit from application changes that can offload work to the GPU. OLCF staff have identified several projects of this nature and begun efforts to collaborate in 2015.

PI Vittorio Michelassi and his project ARD106, "HIPSTAR-G-01," runs HiPSTAR (High-Performance Solver for Turbulence and Aeracoustics Research), developed for direct numerical simulation studies on current computing architectures. The project will continue in 2015 under an INCITE award. Less than 5% of consumed compute time for ARD106 was GPU-enabled in 2014 (see Figure 2.10), a number that should improve in 2015 through collaborations of Michelassi's team with Scientific Computing Group liaison Michael Matheson, whose expertise is in engineering simulations.

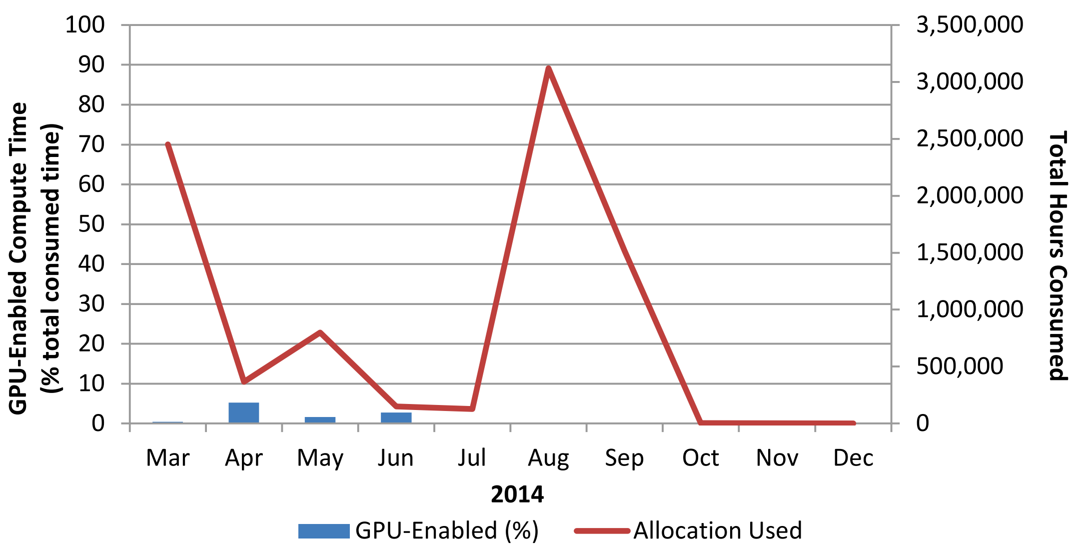

Figure 2.10. Percentage of GPU-enabled compute time (left scale) and total hours consumed (right scale) by month for ARD106, which employs HIPSTAR.

2.12.3 Comparing ALTD and RUR Methods

ALTD was used heavily to track GPU activity through 2012 and 2013, before being supplemented by the 2014 implementation of RUR. Both remain in use. Nine months of overlapping results allow an assessment of the effectiveness of each approach (see Figure 2.11). Neither method is infallible. OLCF has identified several edge cases of GPU-enabled applications that could escape detection using the ALTD method. It is also possible that ALTD can produce false-positives based on the link information provided. The RUR method is new for 2014. Changes to the NVIDIA device driver continue (bug fixes) but no additional features (exposed counters) are expected. In addition, OLCF identified a specific race condition that can cause under-reporting of GPU usage. In this example, jobs that are intentionally constructed to run until terminated by a wallclock exception may not be accurately reported. The volume of results affected by this limitation is estimated as 43% of all batch jobs. The OLCF continues to work with both Cray and NVIDIA to identify other potential edge cases in the RUR method and deal with them as effectively as the underlying GPU hardware and firmware will allow.

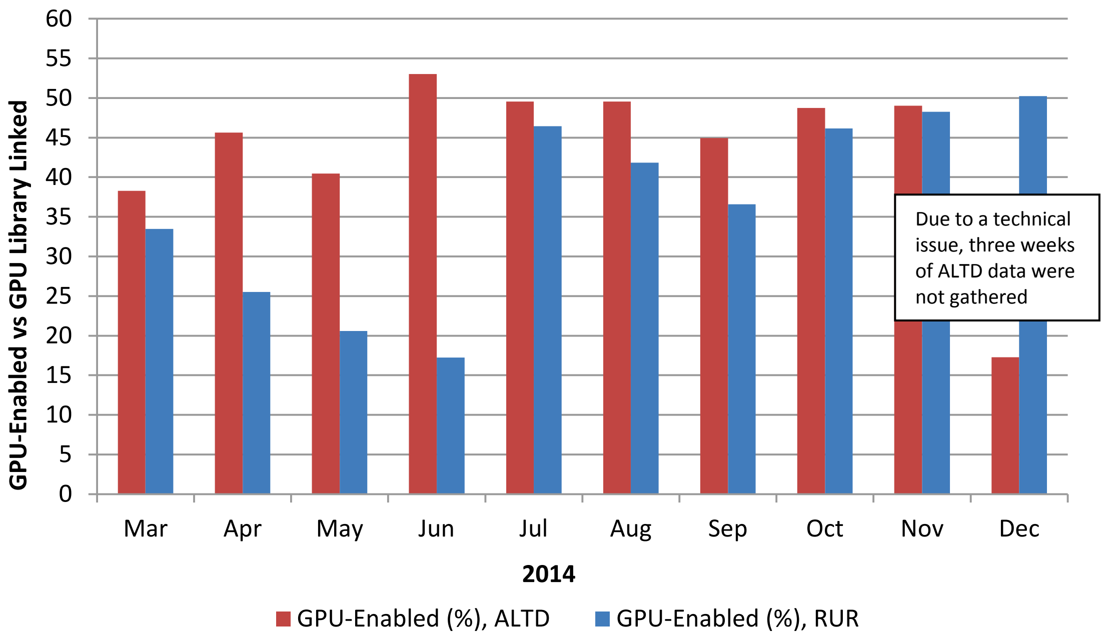

Figure 2.11. GPU Usage on Titan: ALTD and NVML/RUR.

2.13 THE APRUN ADVISORY TOOL

The aprun-Usage (AU) module, first presented as an Innovation in the 2013 OLCF Operational Assessment report, has significantly improved users' performance on Titan.

The CPU on each Titan node has sixteen compute cores, with eight AMD Bulldozer modules. Each Bulldozer has two integer units and one floating point unit. Applications that use more than eight cores to perform high volumes of floating-point operations on the CPUs can experience contention for the Bulldozer's floating-point units since the default ALPS policy is to schedule 16 ranks per node (default placement). To accommodate this, some users explicitly request 9–16 ranks per node (dense placement), which also incurs floating-point contention when using the CPUs. Others attempt to avoid this contention by launching only 2–8 ranks per node (sparse placement), but the default ALPS policy will schedule those ranks so that they share floating-point units. Users typically launch a single MPI rank (single placement) when using the GPU.

The AU module targets the sparse jobs that use 2–8 ranks on 16 or more nodes. It reviews the aprun options and detects sparse jobs. If the job does not set the proper flags, it warns the user by adding a message to the job output. It does not "fix" the flags for the user. Once users set the proper flags, ALPS will schedule the ranks such that they do not share floating point units, and they will typically experience a 1.4–1.7× increase in performance.

The module went live in January 2014. In Figure 2.12, the AU module targeted the green portion of each column. Early in the year, most jobs were sparse and were reviewed by the AU module. The sparse usage decreased in the spring, with a spike in September. The September spike is thought to have been caused by an influx of new users on Titan.

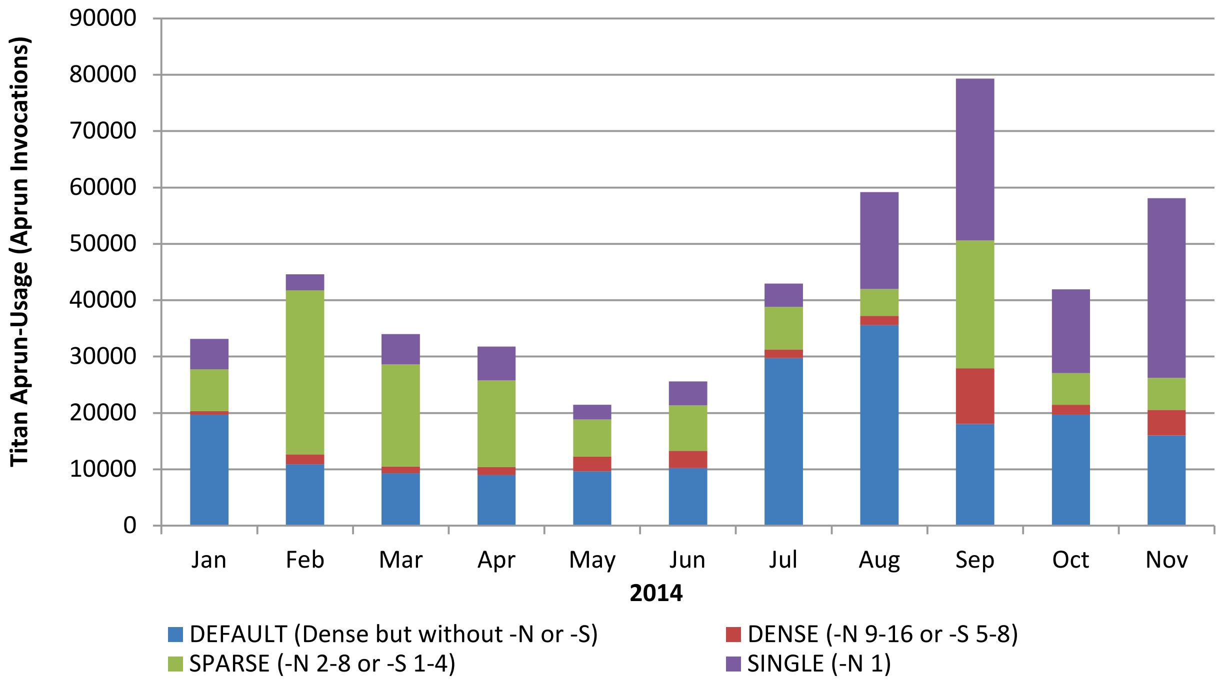

Figure 2.12. The aprun usage module targets sparsely placed MPI ranks.

Figure 2.13 shows that of the sparse invocations, the module clearly had a positive impact on users. In February, the first full month with the module, 83.4% of the 29,116 sparse jobs were notified. In March alone, only 8.6% of the 18,079 jobs were warned about sub-optimal aprun options. The September spike saw a large percentage of sparse jobs warned, which dropped significantly in October and November. This clearly demonstrated the impact of the AU module to the benefit of OLCF users.

As a side consequence of looking at the AU module, a dramatic increase in single-rank jobs was observed in the latter part of 2014. These are clearly GPU users.

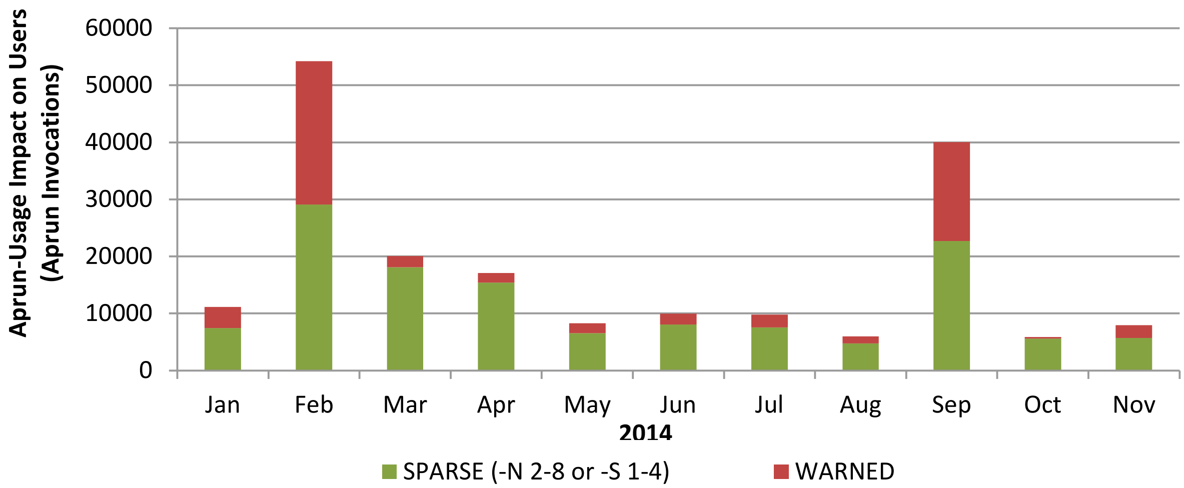

Figure 2.13. Statistics showing the impact of the aprun usage module on users.

2.14 SAFETY

The provision of a safe working environment and the demonstrated safety-conscious attitude of all subcontractors and employees remain important considerations. In the face of the very high volume of work required by the Cray XK7 maintenance activity, the ability to foster and promote a safe work environment remains paramount. See Section 6, Site Office Safety Metrics, for a description of how the facility incorporates DOE site office safety recommendations into its operations.

There were no reportable incidents occurring in this program in 2014.

Strategic Results

HIGH PERFORMANCE COMPUTING FACILITY
2014 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2015

3. STRATEGIC RESULTS

CHARGE QUESTION 3: Is the facility enabling scientific achievements consistent with DOE strategic goals?

OLCF RESPONSE: Yes. The center 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 projects and user programs operating within the OLCF will advance DOE's mission to ensure America's security and prosperity by addressing its energy, environmental, and nuclear challenges through transformative science and technology solutions. The select number of accomplishments that are described in this section serve to communicate how the OLCF is advancing all three strategic objectives of DOE's Strategic Plan Goal 1, "Science and Energy: Advance foundational science, innovate energy technologies, and inform data driven policies that enhance U.S. economic growth, job creation, energy security, and environmental quality, with emphasis on implementation of the President's Climate Action Plan to mitigate the risks of and enhance resilience against climate change," as stated in the *U.S. Department of Energy Strategic Plan: 2014–2018* (March 2014):

- Strategic Objective 1—Advance the goals and objectives in the President's Climate Action Plan by supporting prudent development, deployment, and efficient use of "all of the above" energy resources that also create new jobs and industries.
- Strategic Objective 2—Support a more economically competitive, environmentally responsible, secure and resilient US energy infrastructure.
- Strategic Objective 3—Deliver the scientific discoveries and major scientific tools that transform our understanding of nature and strengthen the connection between advances in fundamental science and technology innovation.

3.1 SCIENCE OUTPUT

2014 Operational Assessment Guidance

The facility tracks and reports the number of refereed publications written annually based on using the facility's resources. This number may include publications in press or accepted but not submitted or in preparation. This is a reported number, not a metric. In addition, the facility may report other publications where appropriate.

3.1.1 OLCF Publications Report

In 2014, 229 refereed publications resulting from the use of OLCF resources were published, as identified in the search and compilation completed on February 12, 2015.[7] Only publications appearing in print in 2014 were counted. Within this 2014 count, users and OLCF staff jointly authored 19 of these publications, users authored 189 without OLCF co-authorship, and OLCF staff published 21 without user co-authorship. In the 2013 OLCF Operational Assessment Report (OAR), 262 publications were reported.

The OLCF follows the recommendation from the 2007 report of the ASCR Advisory Committee Petascale Metrics Panel to report and track user products, including, for example, publications, project milestones (requested quarterly; also examined in the INCITE renewal process), and code improvement. The methodology used to identify publications in 2014 is summarized in Section 3.1.2.

Sponsor guidance allows the reporting of "publications in press or accepted." In previous years, the OLCF has followed this guidance inclusively, reporting manuscripts accepted for publication but not appearing in print. However, search tools may not automatically discover papers that are accepted for publication but not yet generally available to the public. Moreover, reporting papers that are not yet in print, e.g., in 2013, but that will be available in 2014 may complicate the tracking of publications from year to year. The OLCF considers this undesirable; therefore, the OLCF interpretation of the guidance is that only publications appearing in print in the year under review (e.g., 2014) are eligible for tabulation in the current report. The number of publications reported within previous OARs will be reevaluated in the light of this new interpretation, with updates being communicated within regular reporting opportunities.

3.1.2 Methodology for OLCF Publication Discovery and Reporting

The OLCF requires an effective and efficient process to discover, curate, and report publications that have used the facility's computational resources; and no single collection method, if implemented in isolation, is sufficient to the task (e.g., user-reported lists of publications are error prone and may be submitted post-deadline, and automated searches can be no more complete than is the database searched). However, tools developed at ORNL for automated discovery of documents from distributed databases have advanced the ability of the OLCF to reliably perform this task. Therefore, active database searches are integrated with the collection of user-reported publications to create a collection of publications to meet sponsor guidance for reporting.

Three data sets using different methodologies for collection were generated and used to obtain the results for the 2014 *OLCF Operational Assessment Report*:

- An automated discovery process (ADP), using a software tool named COBRA, applied to searching the Institute for Scientific Information (ISI) Web of Science publication database for papers published in final form in 2014 by OLCF users
- Self-reports (SR) of papers published in 2014 by OLCF users and ORNL staff supporting the OLCF program
- ORNL's Publication Tracking System (PTS) for papers published in 2014 from OLCF users and ORNL staff

In previous years, ORNL Research Library staff performed a manual search for papers published by OLCF users within the ISI Web of Science database in support of this publication discovery and reporting. This method was exceedingly time consuming and expensive and was considered to be no longer needed, given the operational improvements in publication discovery and reporting; therefore, the ORNL Research Library did not participate in this task for 2014.

---

[7] In this document, "year" refers to the calendar year unless it carries the prefix "FY" indicating the fiscal year.

Table 3.1 contains statistics from the results of these three approaches. The ADP method identified 167 OLCF publications. The SR method identified 30 OLCF publications not identified by ADP, and PTS identified 32 OLCF publications not identified by ADP. After all of the data sets were compiled into one, duplicates were removed, yielding 229 unique, confirmed publications that resulted from the use of OLCF resources.

Table 3.1. Summary of statistics from OLCF publication-discovery methodology

Process	Unique, confirmed OLCF publications
ADP	167
SR	30
PTS	32
Total	**229**

3.2 SCIENTIFIC ACCOMPLISHMENTS

The OLCF advances DOE's science and engineering enterprise through robust partnerships with its users. The following subsections provide brief summaries of selected scientific and engineering accomplishments, as well as resources for obtaining more information. While they cannot capture the full scope and scale of achievements enabled at the OLCF in 2014, these accomplishments advance the state of the art in science and engineering research and development (R&D) and are advancing DOE's science programs toward their targeted outcomes and mission goals. As an additional indication of OLCF achievements, OLCF users published many breakthrough publications in high-impact journals in 2014, including four in *Science*, one in *Nature*, one in *Nature Physics*, one in *Nature Communications*, one in *Nature Scientific Reports,* two in the *Proceedings of the National Academy of Sciences*, one in *Geophysical Review Letters,* one in *Physical Review X*, and ten in *Physical Review Letters.* Also of interest is the large breadth of journal and conference titles in which OLCF publications appear in 2014. More than 115 unique journal and conference titles are represented in this list.

3.2.1 Titan Takes on the Universe: Salman Habib, Argonne National Laboratory, INCITE

Objective: To perform a large-scale cosmology simulation of the evolution of the distribution of matter in the universe using half a trillion particles—each particle representing roughly a thousandth of the mass of a typical galaxy—and mining the data for "halos," dark matter–dominated clumps in which galaxy formation takes place.

Impact: PI Salman Habib's Argonne National Laboratory (ANL) team models structure formation in the universe using its Hardware/Hybrid Accelerated Cosmology Code (HACC). HACC is designed to exploit multiple leadership-class supercomputing architectures, including Titan's hybrid CPU–GPU system. HACC was introduced in 2009 as the first large-scale cosmology code capable of running on all modern architectures, and it earned Gordon Bell Prize nominations in 2012 and 2013 for outstanding achievements in high-performance computing. Over the next few years, the team will use this scalable, modular code to test the standard model of cosmology and explore the nature of dark energy and dark matter in a virtual universe on Titan and ANL's IBM Blue Gene/Q supercomputer, Mira.

The team will make results from HACC simulations available to astronomers and cosmologists to support efforts in seeking out distant galaxies and astrophysical sources and cataloging their properties and spatial distribution via large-scale sky surveys. By running large-scale cosmology simulations on DOE leadership-class computers and then making the results available to the scientific community, researchers can compare different cosmological models with observational data; produce synthetic (or virtual) galaxy catalogs to guide the design of sky surveys; and use the statistics generated from large-scale simulations to develop cosmological emulators, or precision prediction tools, that will help survey teams interpret large amounts of data. Several surveys have members collaborating closely with the HACC team, including the Baryon Oscillation Spectroscopic Survey, the Dark Energy Survey, the Dark Energy Spectroscopic Instrument (DESI), and the Large Synoptic Survey Telescope (LSST), the last two of which are now under construction.

Accomplishments: The recent HACC simulation on Titan, titled the "Q Continuum" simulation, generated two PB of data (well exceeding Titan's total system memory, see Figure 3.1). This data set is being processed to find halos and determine how they interact and merge. The resulting halo merger tree information is fed into modeling codes to generate a picture of how galaxies form and evolve. The analysis of the simulation will enable a large number of scientific projects, including high-quality sky maps for DESI and LSST and fundamental investigations of the formation of the large-scale structure of the universe. One of the computational challenges the team overcame was designing a scheme to balance workload evenly on Titan's GPU nodes by shifting work to less-busy nodes when highly dense regions were calculated, so that the code could run as efficiently as possible. By using the computational power of Titan's NVIDIA GPUs, the halo finder was able to identify more than 167 million halos (equaling 46 percent of all the particles in the simulation) in less than 36 minutes. The program calculates and stores a set of descriptors, such as mass, velocity, and radial profiles, of halos consisting of more than 40 particles.

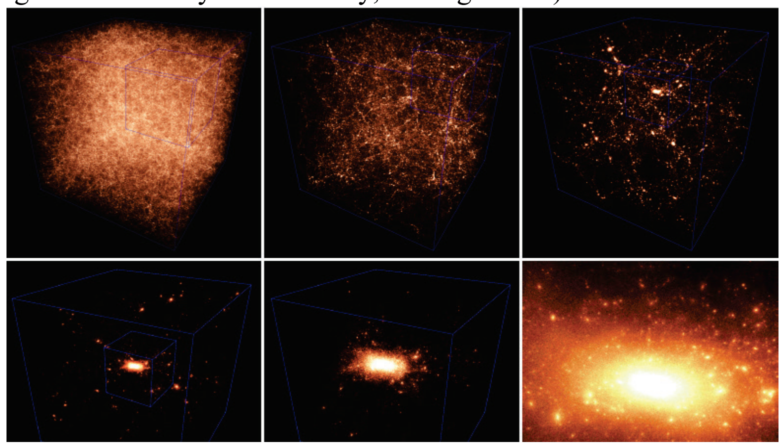

Figure 3.1. Halo particles (members of clumps of matter where galaxy formation takes place) in the Q Continuum simulation performed on Titan, one of the largest high-resolution simulations of the universe to date. The resolution of this simulation is designed to match modern-day sky surveys that find and catalog distant galaxies. The series of images (from upper left to bottom right) "zooms in" on one of the largest galaxy cluster-scale halos in the Q Continuum run, giving some impression of the enormous dynamic range of the simulation.

OLCF Contributions: Scientific liaison Bronson Messer diagnosed and resolved GPU software and hardware issues relevant to the Q Continuum run. HACC, with its particular heritage in the OpenCL standard and considerable scalability, was the only code to encounter this particular set of problems. These were severe upon presentation, and resolution was essential for this scientific achievement. The science-driven urgency with which OLCF engaged this user project resulted in the identification and correction of the software bugs. With these issues resolved, the simulation used almost all of Titan (90 percent of the nodes) with an efficient throughput for such a high fraction of machine utilization. OLCF is also in the process of creating a virtual machine in ORNL's Compute and Data Environment for Science (CADES) to provide a portal for access to reduced HACC data for researchers beyond the team's immediate collaborators. The team used 190 million Titan core hours, 99.8 percent at 60–100 percent of Titan's nodes.

Related Publications:

Amol Upadhye, Rahul Biswas, Adrian Pope, Katrin Heitmann, Salman Habib, Hal Finkel, Nicholas Frontiere. 2014. "Large-scale structure formation with massive neutrinos and dynamical dark energy," *Physical Review D* **89**(10): 103515. doi: 10.1103/PhysRevD.89.103515

Katrin Heitmann, Salman Habib, Hal Finkel, Nicholas Frontiere, Adrian Pope, Vitali Morozov, Steve Rangel, Eve Kovacs, Juliana Kwan, Nan Li. 2014. “Large-scale simulations of sky surveys,” *Computing in Science and Engineering* **16**(5): 14–23. doi: 10.1109/MCSE.2014.49

Katrin Heitmann, Nicholas Frontiere, Chris Sewell, Salman Habib, Adrian Pope, Hal Finkel, Silvio Rizzi, Joe Insley, and Suman Bhattacharya. 2014. “The Q Continuum simulation: Harnessing the power of GPU accelerated computers,” submitted to *Astrophysical Journal Supplement*; *arXiv.org*, http://arxiv.org/abs/1411.3396.

3.2.2 The Complexities of Combustion: Jacqueline Chen, Sandia National Laboratories, INCITE

Objective: To simulate a jet flame burning dimethyl ether (DME) using a direct numerical simulation code known as S3D. The results will serve as a benchmark for model development for both combustion and pollutant formation processes used by combustion researchers to develop predictive models for the design of next-generation diesel engines.

Impact: PI Jacqueline Chen’s team has simulated jet flames in the past, but the latest simulations on Titan were a breakthrough for two reasons: the inclusion, for the first time, of DME and the valuation of the highest Reynolds number ever achieved by the team, 13,050. The increased Reynolds value (a measure of the mixing intensity and dynamic range of turbulence) allows the team to resolve a wider range of turbulence scales in space and time, a major breakthrough in trying to match experimental conditions and evaluate turbulent mixing and combustion models.

The ultimate goal is a predictive modeling capability that will lead to shorter, cheaper engine design cycles for US industry. The low-temperature diesel engine is one of the most thermodynamically efficient internal combustion engines, yet one of the challenges during its design is ensuring low pollutant emissions. Pollutant emissions are sensitive to premixing, temperature variations, fuel compositions, and other factors that converge during combustion. Therefore, there is strong demand for models that can predict combustion, including pollutant formation, accurately.

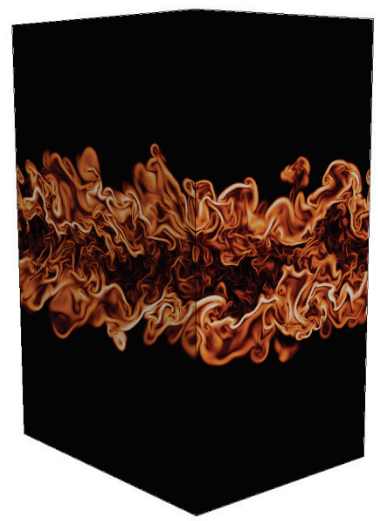

Figure 3.2. The logarithm of the scalar dissipation rate (that is, the local mixing rate), where white denotes high mixing rates and red, lower mixing rates.

Image credit: Ankit Bhagatwala, Sandia.

Accomplishments: Chen’s Sandia team member Ankit Bhagatwala employed S3D on Titan to simulate a jet flame burning DME in an attempt to match the conditions of a companion experiment at Ohio State University. The jet flame simulation is used to probe fundamental turbulent flame physics associated with local extinction, in which parts of the flame burn out, and reignition—a phenomenon that may occur in practical combustors like diesel engines.

Specifically, the team wanted to know the dependence of reignition on the local mixing rate, i.e., the rate of fuel and air mixing during the combustion process (Figure 3.2). They found that oxygenated fuels such as DME generate considerably more stable intermediates such as formaldehyde, rendering the flame more robust against local extinction than conventional hydrocarbons such as methane.

Furthermore, the simultaneous imaging of formaldehyde and hydroxyl radical, a neutral form of hydroxide ion, was evaluated to determine its effectiveness at measuring the peak heat release rate. The simulation data verified that this imaging method performed extremely well at predicting the maximum heat release rate in the flame. This method was subsequently applied to the experimental data as a diagnostic to measure the extent of local extinction and reignition.

If researchers can identify strategies to minimize flame extinction, they will greatly enhance efficiency and minimize undesired emissions in combustion devices such as engines. But even with computers as powerful as Titan, performing a direct numerical simulation of a diesel jet flame matching all of the aero-thermo-chemical conditions, and with all of its chemical reactions, is out of the question. Researchers therefore simulate a jet flame configuration at a lower Reynolds number that matches the critical thermo-chemical conditions associated with combustion, using around 30 chemical molecules to model the combustion of DME. Executing runs using this large number of molecules is possible only on systems like Titan, and it advances the team's goals toward simulating the behavior of more realistic fuels, including biofuels.

Xinyu Zhao, a member of Chen's group from the Combustion Energy Frontier Research Center, is using this data to evaluate a newly proposed turbulent mixing model. Such models may ultimately be used in engineering-scale computational fluid dynamics simulations that run on desktops and computer clusters, and they could optimize the designs of combustion devices using diverse fuels. Because industrial researchers must conduct thousands of calculations around a single parameter to optimize a part design, individual calculations need to be inexpensive.

Titan has proved the perfect platform for Chen's research. S3D is six times faster on Titan than on the OLCF's previous CPU-only system Jaguar, in part because of OLCF contributions in porting S3D to Titan's hybrid architecture.

OLCF Contributions: OLCF's Ramanan Sankaran, a computational combustion scientist, worked closely with the team to prepare S3D to run effectively on Titan. Sankaran, Ray Grout of the National Renewable Energy Laboratory, and John Levesque of Cray, hybridized S3D to extract multiple levels of parallelism to achieve good performance and portability. S3D was one of the projects of OLCF's Center for Accelerated Application Readiness Early Science that used the OpenACC community standard programming environment. The success of the S3D simulations reinforce Titan's promise of performance portability, or the idea that improvements made to applications so that they can run on Titan will translate to other systems. The team used 111 million Titan core hours; 77 percent of the time used 20–60 percent of Titan's nodes and 2 percent of the time used 60–100 percent of Titan's nodes.

Related Publications:

Ankit Bhagatwala, Zhaoyu Luo, Han Shen, Jeffrey A. Sutton, Tianfeng Lu, and Jacqueline H. Chen. 2014. "Numerical and experimental investigation of turbulent DME jet flames," *Proceedings of the Combustion Institute* **35**(2): 1157–1166. doi:10.1016/j.proci.2014.05.147.

Stephen B. Pope. 2013. "A model for turbulent mixing based on shadow-position conditioning," *Physics of Fluids* **25**: 110803. doi:10.1063/1.4818981.

John M. Levesque, Ramanan Sankaran, and Ray Grout. 2012. "Hybridizing S3D into an exascale application using OpenACC: An approach for moving to multi-petaflops and beyond," presented at the SC12 supercomputing conference. doi: 10.1109/SC.2012.69.

Online Story:

OLCF Staff Writer, "The Complexities of Combustion", *OLCF News* (November 11, 2014).

3.2.3 Researchers Get Warmer in Understanding High-Temperature Superconductors: Paul Kent, ORNL, INCITE

Objective: To apply ab initio many-body QMC simulations to materials problems, particularly those for which standard electronic structure approaches such as density functional theory (DFT) fail or can be used only empirically. The team focuses on transition metals and transition-metal oxides that have a wide range of technological applications.

Impact: DOE considers achieving a quantitative understanding of electronic phenomena in novel materials a grand challenge. More specifically, by being able to better predict superconductors' behaviors, scientists can extend this predictive capability into a wide range of applications, including energy storage, catalysis, energy production, and metals that can be used as structural materials.

Until recently, most computational research into superconducting materials was performed using DFT or other heuristic models. While effective at studying electronic interactions for many classes of materials, DFT is incapable of solving ab initio computational models for cuprates— in this study, copper-oxide–based high-temperature superconductors—because of difficulties in describing the strongly coupled electron interactions. By validating the QMC method for simulating cuprates, researcher Paul Kent's team moved closer to understanding the remarkable properties of high-temperature superconductors. As supercomputing power increases, this QMC method can be applied to heavier, more complex transition metals and metal oxides.

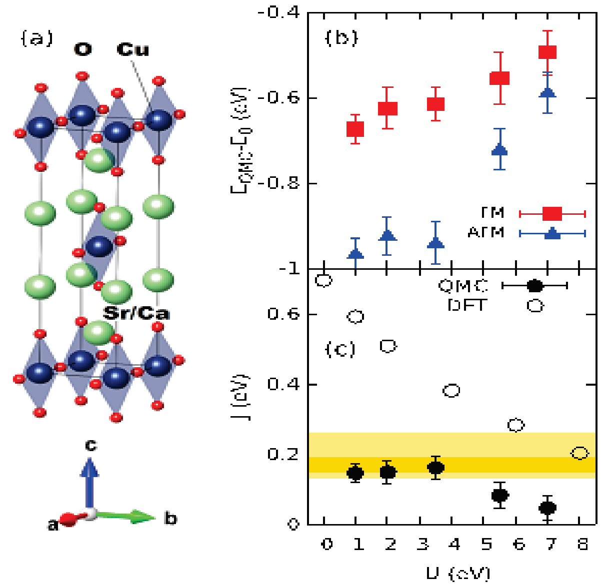

Figure 3.3. (a) The crystal structure of (Ca/Sr)$_2$CuO$_3$. (b) The calculated quantum Monte Carlo (QMC) energies of different magnetic states. Lower QMC energies correspond to more accurate predictions. (c) The calculated magnetic couplings from different methods, with the experimental range highlighted in yellow. The best QMC results from (b) give the best agreement with experiment, indicating that the QMC predictions are robust. Image credit: Foyevtsova et al.

Accomplishments: The Kent team was able to obtain spin superexchange coupling[8], or a type of magnetic coupling between atoms within a material, through simulation in quantitative agreement with actual experiment. Unlike simulations using DFT, which allow researchers to tune predefined parameters, the team used the QMC method, a more computationally expensive, but more extensive, calculation. By using QMC, the team performed the first ab initio simulation of a cuprate.

The team was able to simulate Ca_2CuO_3, one of the lighter metal oxide superconductors. Because of computational constraints relating to memory, the team cannot yet simulate the full range of cuprates, but a simulation of Ca_2CuO_3 serves as strong proof of the QMC method (see Figure 3.3).

Initial simulations were computationally expensive, but through the end of 2014, the Kent team was able to reduce memory usage by a factor of 8, allowing for more computation on GPUs and for some computation to be done on CPUs that previously required too much memory.

OLCF Contributions: OLCF computational scientist Ying Wai Li helped researchers use GPUs to run large supercell simulations. Using an earlier CUDA-based implementation of certain compute-intensive routines, the researchers were limited to around 730 electrons in a simulation. Li's work helped the team simulate more than 1,000 electrons in a simulation. Li also created a new matrix inversion, allowing the team to delete hundreds of line of source code, thereby streamlining the computational demands of running QMCPACK. The team used 55 million Titan core hours, 13 percent of which required less than 20 percent of Titan, 78 percent of which used 20–60 percent of Titan, and 9 percent of which used more

[8] Superexchange is the strong (usually) antiferromagnetic coupling between two next-to-nearest neighbor cations through a non-magnetic anion.

than 60 percent of Titan. See sections 1.4.1.2 and 1.4.1.3 for a more detailed description of the OLCF liaison contributions to this INCITE project.

Related Publication:

K. Foyevtsova, et al. "Ab initio quantum Monte Carlo calculations of spin superexchange in cuprates: The benchmarking case for Ca_2CuO_3," *Physical Review X* **4**: 031003 (2014). doi: http://dx.doi.org/10.1103/PhysRevX.4.031003.

3.2.4 Modeling the Skin Barrier: Michael Klein, Temple University, INCITE

Objective: To develop computational modeling techniques that improve understanding of the skin's outermost barrier layer—the stratum corneum—at the molecular level and that explain why skin is permeable to some compounds, impermeable to others, and damaged by still others.

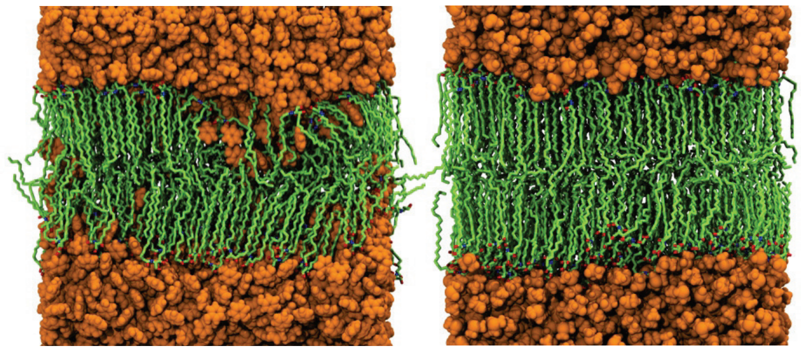

Figure 3.4. Snapshots from simulations of a stratum corneum model (green) with additives (orange) phenol (left) and dimethyl sulfoxide (right) at the same molar concentrations. Notice the high level of disruption in the case of phenol. Image credit: Russell DeVane (Procter & Gamble)

Impact: The project, conducted by researchers at Temple University and Procter & Gamble (P&G), delivered a comprehensive picture in full atomistic detail of the molecular properties of the stratum corneum. This new knowledge gives researchers a foundation for probing what factors impact skin barrier integrity, and it will ultimately lead to a better model for predicting the overall skin permeability of a compound. The multi-billion dollar personal care industry depends on such advancements to design and improve a wide range of skin care and hair care products, as well as goods such as cosmetics, soaps, and deodorants. The data could also enable the pharmaceutical industry to advance the development of drugs that are administered through the skin. Before this project, skin had been examined only with microscopes and relatively simple computational models. Now it has been examined atom by atom.

Accomplishments:

This project achieved a more complete picture of skin-barrier function; the project simulated a one million–atom skin lipid matrix comprising four separate bilayers (Figure 3.4). Each layer was hundreds of square nanometers and made of ceramides (waxy compounds), fatty acids, and cholesterol in water; each shows how skin bio-membrane properties affect the interaction of small molecules, typically used for drugs and personal-care products, with skin. The researchers used 87 million core hours on Titan, achieving several milestones in a matter of weeks that would have taken years with their in-house system

The researchers integrated new levels of detail and complexity into the existing heuristic model created by Jerry Kasting of the University of Cincinnati, adding a more realistic representation of the stratum corneum lipid matrix and using large-scale molecular dynamics simulations to capture how chemicals permeate skin. The enhanced model validated the Kasting model's assumptions regarding molecular weight and octanol-water partition coefficient for a broad range of compounds.

Additionally, modeling of the interaction between skin and large concentrations of selected chemical compounds revealed the key role of ceramides in permeation. Simulations showed how compounds could essentially melt this waxy substance, but chemically instead of thermally.

The work on Titan relied primarily on the GPU-accelerated LAMMPS and GROMACS molecular dynamics codes to track the movements of every molecule in the stratum corneum lipid model to see how it evolves over time.

OLCF Contributions: OLCF computational scientist Mike Brown and other OLCF staff modified LAMMPS to efficiently take advantage of GPUs. Furthermore, Arnold Tharrington, the OLCF project liaison, adjusted LAMMPS to net an additional 10 percent performance boost. The modified code performed 2.5× faster than a comparable CPU-only system (Cray XE6). In cooperation with application developers, OLCF staff worked to compile an optimal version of NAMD. Additionally, staff collaborated to write an automated HPSS archiving tool that facilitated the transfer of data to long-term storage. This team used 71 million Titan core-hours, 87 percent of which required less than 20 percent of Titan, and13 percent of which required between 20 and 60 percent of system resources.

Related publication:

M. Paloncyova, R. H. DeVane, B. P. Murch, K. Berka K., and M. Otyepka. 2014. "Rationalization of reduced penetration of drugs through ceramide gel phase membrane," *Langmuir* **30**(46): 13942–13948 doi: 10.1021/la503289v.

C. M. MacDermaid, R. H. DeVane, M. L. Klein, G. Fiorin. 2014. "Dehydration of multilamellar fatty acid membranes: Towards a computational model of the stratum corneum," *Journal of Chemical Physics* **141**(22): 22D526. doi: 10.1063/1.4902363.

Online Story:

OLCF Staff Writer, "Procter & Gamble and Temple University scientists model skin's makeup" *OLCF News* (November 14, 2014).

3.2.5 Turning Up the Heat from Plasma Exhaust in Tokomak Reactors: C. S. Chang, Princeton Plasma Physics Laboratory, INCITE

Objective: To use very large-scale simulations to achieve a first-principles understanding of the divertor heat-load width physics in present-day tokamak-style fusion energy reactors.

Impact:

One of a few remaining critical ITER physics issues that have not been resolved is the divertor heat-load width. With plasma temperatures of up to 300 million °F, ITER is being designed to exhaust the plasma heat into a divertor chamber with an expected heat load on the divertor surface of up to 20 MW/m^2, assuming that the heat load would spread over a reasonable surface area. However, extrapolating results from experiments in tokamak devices currently available indicates that the exhaust heat could be deposited on the divertor wall within a tightly focused 1 mm wide strip, traveling continuously along the toroidal direction. The empirical prediction is achieved by extrapolating today's observed trends for the heat-load width to be approximately inversely proportional to plasma current ($1/I_P$). This could mean a much higher level of heat load on the narrow strip than the designed material tolerance limit. Since there is little understanding of the underlying physical mechanisms leading to $1/I_P$ behavior of the divertor heat-load width, a simple extrapolation to the burning-fusion device ITER—which is much greater in size and energy—may not be accurate; and confidence in this empirical scaling rule is weak. Since the edge plasma is not in thermal equilibrium, it needs an accurate, kinetic understanding with all the important multiscale physics included: orbit dynamics of charged particles in a curved magnetic field, neutral particle transport and atomic interaction with plasma, nonlinear Coulomb collisions, and blobby micro-turbulence. Such study requires extreme-scale computing and is possible only on Titan at present.

Accomplishments:

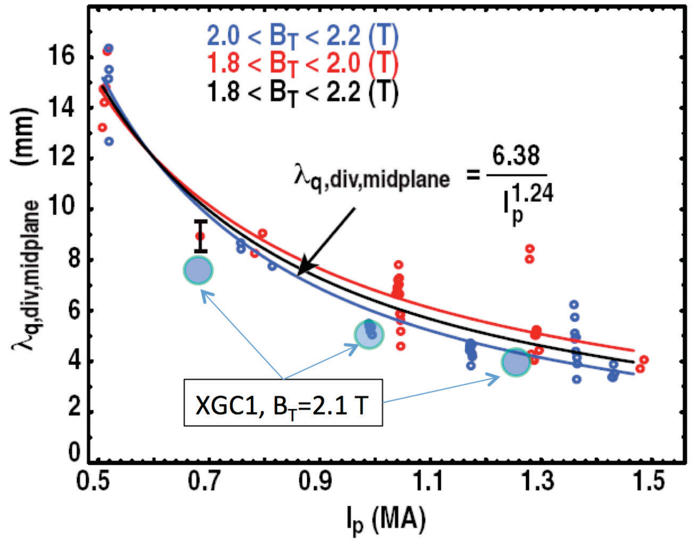

Figure 3.5. Comparison of the XGC1 simulated divertor heat-load widths (three larger light blue circles) with the experimental observation at similar BT values (blue dots) predicting from first principles the inverse dependence on plasma current.

Building upon their success in 2013 in simulating blobby edge turbulence on Titan in a realistic DIII-D geometry, Chang's team challenged the XGC1 code again in 2014 to obtain fundamental understanding of the divertor heat-load width physics that leads to a quantitative explanation of the inverse proportionality between the heat-load width and the plasma current in present-day tokamak devices. All the important multiscale edge physics phenomena have been simulated together, of which the blobby turbulence is an important part. For a more comprehensive validation and understanding, both a conventional aspect ratio tokamak (DIII-D at General Atomics in San Diego) and a tight aspect ratio tokamak (NSTX at Princeton) were targeted (Figure 3.5). The approximately $1/I_P$ scaling of the divertor heat-load width was predicted for each tokamak geometry, and a fundamental physics understanding was obtained. The $1/I_P$ scaling was found to be from the guiding center orbit excursion of the kinetic ions, combined with the pre-sheath effect. The effect of the blobby turbulence on the divertor heat-load spread was found to be a minor factor in current tokamaks. However, in ITER, the blobby turbulence effect is expected to become more important, since the orbital spread size that leads to $1/I_P$ behavior becomes much smaller than the blob size. Thus the naive extrapolation of the present $1/I_P$ scaling to ITER is not expected to hold, a finding that could reduce the predicted exhaust heat intensity and make ITER operation fall within the material design limits. A more extreme-scale, predictive study for full ITER plasma is scheduled on Titan in 2015.

OLCF Contributions:

Chang's ORNL-based collaborators, Ed D'Azevedo, Pat Worley, Scott Klasky, Jong Choi, Nobert Podhosky, and Dave Pugmire (supported by the OLCF, SUPER, SDAV, and the SciDAC project Center for Edge Physics Simulation) optimized the new, fully nonlinear collision operator to a practical level. They engaged increased thread parallelism by employing OpenMP regions within the collision operator, improved the ADIOS I/O speed, and helped with the visualization of output. The latter work was carried out in collaboration with the SciDAC FASTMath Center. As a result, the speed of the nonlinear collision operation increased by approximately a factor of 5. This achievement required runs that would not be possible at any other HPC center: most of the work required 88 percent of the maximal hybrid XK7 capability using GPUs and CPUs simultaneously. The team used 162 million Titan core hours, 91 percent of which was consumed at the capability scale (i.e., greater than 3,750 XK7 nodes), 66 percent at 20–60 percent of the machine size, and 25 percent at 60–100 percent of the machine size. See Section 1.4.1.2 for a more detailed description of the OLCF liaison contributions to this strategic achievement.

Related Publications:

Janghoon Seo, C. S. Chang, S. Ku, J. M. Kwon, W. Choe, Stefan H. Mueller. (2014). "Intrinsic momentum generation by a combined neoclassical and turbulence mechanism in diverted DIII-D plasma edge," *Physics of Plasmas* **21**(9): 092501. doi: 10.1063/1.4894242.

C. S. Chang, J. Boedo, R. Hager, S. Ku, J. Lang, R. Maingi, S. E. Parker, D. Stotler, and S. J. Zweben. 2014. "Gyrokinetic study of edge blobs and divertor heat-load footprint," *Proceedings of the 25th Fusion Energy Conference, IAEA*, October 13–18, 2014, Saint Petersburg, Russian Federation, in press (2015).

3.3 DIRECTOR'S DISCRETIONARY PROGRAM

2014 Operational Assessment Guidance

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

The OLCF allocates time on leadership resources primarily through the INCITE program and through the facility's DD program. The OLCF seeks to maximize 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 make effective use of leadership resources. Further, through the ALCC program, the ASCR office allocates up to 30 percent of the facility's resources.

The goals of the DD program are threefold:

- Prepare for leadership computing competitions (i.e., INCITE and ALCC), e.g., to improve and document application computational readiness.
- Broaden the community of researchers capable of using leadership computing.
- Develop 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 four mission goals of the OLCF:

- To enable high-impact, grand-challenge science and engineering that could not otherwise be performed without the leadership-class computational and data resources
- 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
- To educate and train the next-generation workforce grounded in the application of leadership computing to the most challenging scientific and engineering problems

R&D partnerships are those aligned with DOE and ORNL strategic agendas. They may be entirely new areas with respect to HPC or ones in need of nurturing. Example projects are those associated with the ORNL LDRD 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 Center for Advanced Simulation of Light Water Reactors (CASL), Critical Materials Institute hub led by Ames National Laboratory, ACME program, CNMS, and large experimental facilities such as the SNS at ORNL and the ATLAS and ALICE Experiments at CERN and Brookhaven National Laboratory. Also included in this broad category are projects that came to the OLCF through the Accelerating Competitiveness through Computational ExceLlence (ACCEL) Industrial HPC Partnerships Program, providing opportunities for industrial researchers to access the leadership systems to carry out work that would not otherwise be possible. See section 3.4 for more information about this program. Through ACCEL, the OLCF is achieving the original Congressional intent for the Leadership Computing Program by providing "Leadership Systems, on a competitive, merit-reviewed basis, […] to researchers in

United States industry, institutions of higher education, national laboratories, and other Federal agencies."[9]

The DD program is also accessible to the general HPC community to carry out porting and development exercises for nascent and less-efficient applications. These performance enhancement projects range in scope from immediate INCITE preparation—designed to allow investigators the opportunity to test their codes' computational readiness on INCITE platforms—to somewhat longer-term projects involving improvement in algorithms and implementations. The ORAU–ORNL High-Performance Computing Grant Program awards modest grants of research funding, provided by ORAU, and modest grants of leadership computing time, provided by OLCF, to ORAU-member-university faculty members. The program was established in 2009 to encourage new and expand existing research initiatives among ORAU member institutions using HPC systems. The program is a competitive grant program managed and funded by ORAU and open only to ORAU's member institutions.[10]

The following are examples of DD program outcomes in expanding the leadership computing science community. Of the 145 DD projects operational at OLCF in 2014, 39 were in support of the development of proposals submitted to the 2015 INCITE Call for Proposals or the 2014–-2015 ALCC Call for Proposals. Twelve of the proposals submitted were awarded allocations at OLCF through INCITE or ALCC. See Appendix E for a complete list of the 2014 DD projects.

The OLCF DD program also supports a variety of "data projects" that require data storage and bandwidth capabilities but few compute resources (see Section 4.2). Ongoing data projects include the Earth System Grid Federation, Data Sharing Project for the Center for Exascale Simulation of Combustion in Turbulence codesign project, 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 in DD projects.

The Resource Utilization Council makes the final decision on DD applications, using written reviews from subject matter experts. The actual DD project lifetime is specified upon award: allocations are typically for 1 year or less. The average size of a DD award is roughly 3 million Titan core-hours but can range from tens of thousands to 20 million hours or more.

Since its inception in 2006, the DD program has granted allocations in virtually all areas of science identified by DOE as strategic for the nation (Table 3.2).

Table 3.2. Director's Discretionary program: Domain allocation distribution (percent)

Time period	Biology	Chemistry	Computer Science	Earth Science	Engineering	Fusion	Materials Science	Nuclear Energy	Physics
2008	19	8	28	4	8	15	3	1	14
2009	5	3	19	6	8	6	33	1	19
2010	9	6	10	8	19	6	16	3	23
2011	7	1	10	19	14	0	9	13	26
2012	6	1	21	14	25	5	10	1	18
2013	9	4	15	15	12	8	14	4	19
2014	17	4	22	14	16	2	13	1	11

In 2014, the OLCF DD program participants fully utilized the 10 percent of the resource set aside for these DD program goals, consuming 464,653,171 Titan core hours.

---

[9] Department of Energy High-End Computing Revitalization Act of 2004. Public Law 108–423—NOV. 30, 2004.
[10] http://www.orau.org/university-partnerships/faculty-student-programs/hpc/default.aspx.

3.4 INDUSTRIAL HPC PARTNERSHIPS PROGRAM: ACCEL

ACCEL completed its sixth year in 2014, advancing complex science and engineering solutions and helping DOE meet its goal of increasing the community of researchers able to use next-generation leadership computing resources

Thirty-four industrial projects were under way during 2014. These projects used 262,694,244 million hours, representing approximately 6 percent of the total hours that Titan delivered in 2014.

- In 2014, 51 percent of the industrial project hours were allocated through INCITE, 24 percent via ALCC, and 25 percent through the OLCF DD program.
- Of 34 projects, 18 were new. These firms received awards via INCITE (2 projects), ALCC (2 projects), and DD (14 projects).

New industrial users who received awards (all DD) included Rolls Royce, NVIDIA, Tennessee Valley Authority, Arkema, and oil and gas firm TOTAL.

Ford and GM built on the experience they gained using leadership computing in earlier DD awards to compete successfully for their first ALCC awards (Ford in partnership with ORNL researchers) to continue work intended to help these firms meet looming new fuel efficiency standards. This is another example of OLCF's industrial partnership program, ACCEL, helping firms gain the skills they need to advance in their ability to use leadership computing.

Seattle-based small business Ramgen Power Systems also offers a compelling example of how ACCEL is helping industry reap competitive benefits from using leadership computing. Ramgen used Titan to continue developing innovative, shock wave–based compression systems for carbon capture and sequestration (CCS), work originally begun through an ALCC award on Jaguar to help contribute to DOE goals for reducing CCS costs. In July 2014, the company announced it was being purchased by Dresser-Rand, one of the world's largest suppliers of custom-engineered rotating equipment solutions for the oil, gas, chemical, petrochemical, process, and power generation industries. Ramgen indicated that it was purchased in part because of the advanced modeling and simulation capabilities Ramgen engineers have acquired using OLCF resources, and the firm's resultant ability to apply large-scale computational fluid dynamics simulations to accelerate development of turbo-machinery technology. That new capability helped make Ramgen an attractive acquisition for a world-leading manufacturer.

3.4.1 Katrisk Taps Titan to Model Flooding Risks: Dag Lohmann, Katrisk, Director's Discretionary[11]

California startup KatRisk proved again that big problems are not the purview of big companies and that small companies also can use access to leadership computing. KatRisk used Titan to create an unprecedented product: flood risk maps covering the globe.

No area is immune to the devastation of flooding; the risks extend worldwide and often result in loss of life and livelihood. In the United States, flood maps from the Federal Emergency Management Agency (FEMA) do not cover the whole country. And even these maps do not contain all the information needed to determine a location's risk. For instance, while they tell you whether your property is in danger of flooding, current FEMA maps for the most part do not indicate whether the flooding will be an inch, a foot, or 10 feet.

The situation is even worse in most other countries, where it is often very unclear who is and is not in a flood zone.

To complicate the issue even further, many flooding maps do not cover all the ways an area can become flooded. While they typically do address the dangers posed by overflowing rivers (known as fluvial or riverine flooding), they often do not consider the flash flooding associated with heavy rain

[11] https://www.olcf.ornl.gov/2014/11/11/start-up-firm-taps-titan-to-model-flooding-risks-worldwide/

(known as pluvial flooding). Pluvial flooding may flow into existing rivers, or it may form its own rivers. The dangers associated with it help explain why 20 to 40 percent of US flood claims come from outside FEMA flood zones.

KatRisk's models address these gaps and proved to be a natural fit for GPUs. But the company's internal cluster would not produce timely flood maps for the entire world. Recognizing this, KatRisk applied for and received a DD allocation of 5 million hours on Titan. The KatRisk team was able to harness as many as 1,500 GPUs at a time and reach a performance of 6 petaflop single-precision operations.

With the power of Titan, the team developed for the first time fluvial and pluvial high-resolution, global flood maps. Titan enabled Katrisk to reduce 18 months of computation time to 5 months for its US and Asia maps, accelerating delivery time to the insurance industry by a full year. These maps give insurance firms a better understanding of a region's real flood risks and the ability, therefore, to reflect that risk more accurately in insurance products. This drives prudent decision making among homeowners. And while the decision not to build in a flood zone might be based on the cost of flood insurance, it could, coincidentally, end up saving lives as well. KatRisk plans to complete by end of 2015 flood maps for the rest of the world.

KatRisk also is providing free information on potential losses from floods, as well as hurricanes in the United States and typhoons in Asia, through an online tool called "Katalyser." To use the tool, a property owner can visit http://katmaps.katalyser.com/and enter information such as the latitude and longitude of the property.

3.4.2 Industrial Partnership User Agreement

Ford, GM, and KatRisk all benefited from a unique user agreement ORNL developed with DOE that provides companies more flexibility in the kinds of problems they can run at the OLCF. Standard DOE user agreements require that new information generated from the project (i.e., project output) be considered completely proprietary or completely nonproprietary. Proprietary projects pay a cost recovery fee, and nonproprietary projects are expected to publish their results/new intellectual property.
Under ACCEL, ORNL found that company projects often generate both kinds of intellectual property. This special industrial partnership user agreement, available only at the OLCF, permits an industrial project to have a blend of both proprietary and nonproprietary output, and allows the firm to keep the proprietary portion confidential as long as the firm can commit in advance to publish meaningful science results nonetheless. In addition, this user agreement does not permit technical collaboration (owing to the presence of proprietary output). This agreement is particularly useful for small and medium-size businesses that do not have separate R&D departments accustomed to performing fundamental, nonproprietary scientific research. However, large firms also appreciate the flexibility of the agreement, as it enables them to bring real-world problems to the OLCF instead of sanitized problems or problems from open scientific literature. This agreement has been implemented so successfully on a number of industrial projects that received OLCF DD allocations that DOE is now in the process of broadening its applicability to ALCC projects at the OLCF and at ALCF and NERSC.

4. INNOVATION

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

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

Since the facility's inception in 2004, OLCF staff have provided leadership in the HPC community, spearheading the creation and development of tools and policies necessary for computing and computational science. Recently, this leadership was recognized. In a paper nominated as a best paper finalist at the SC14 supercomputing conference, OLCF staff presented a state-of-practice summary of the facility's decade-long experience in acquiring, deploying, benchmarking, developing, and operating large-scale parallel file systems (PFS). [12] The OLCF's account provides a valuable best-practice guide for the HPC community, highlighting how to make key design choices and avoid pitfalls in choosing and hosting PFS. The paper illustrates the OLCF's experience in evaluating and, when necessary, influencing vendor roadmaps for PFS architectures. Now, looking ahead, the OLCF is blazing new trails in the effective use of the Titan heterogeneous system. It is not possible to highlight all of the innovative work carried out by the OLCF. Instead, this report will focus on several key strategic areas of operations in 2014: GPU activity, data analysis and workflow, and I/O optimization.

4.1 GPU ACTIVITY

The Titan system provides the largest extant heterogeneous architecture for computing and computational science. As shown in Section 2, Business Results, Titan usage is high, delivering on the promise of a system well suited for capability simulations for science. This success is due in part to innovations in tracking and reporting the activity of the nodes and measuring their high degree of reliability, as well as creating novel ways to balance the workload over the entire node.

4.1.1 GPU Statistics and Reporting

In 2013, the OLCF reported the number of GPU-enabled jobs on Titan by using information about GPU-centric software libraries. This initial approach counted a batch job as GPU-enabled if a binary that

[12] Sarp Oral, James Simmons, Jason Hill, Dustin Leverman, Feiyi Wang, Matt Ezell, Ross Miller, Douglas Fuller, Raghul Gunasekaran, Youngjae Kim, Saurabh Gupta, Devesh Tiwari, Sudharshan S. Vazhkudai, James H. Rogers, David Dillow, Galen M. Shipman, Arthur S. Bland, "Best practices and lessons learned from deploying and operating large-scale data-centric parallel file systems," *Proceedings of Supercomputing 2014* (SC14): 27th IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis, New Orleans, November 2014.

executed within the batch job had been linked against a GPU-centric software library. This technique, while enlightening, had two shortcomings:

- Batch jobs that did not link against any GPU-centric libraries but ran code on the GPUs via other mechanisms were not counted as GPU-enabled.
- Batch jobs that linked against GPU-centric libraries but did not actually run code on the GPUs were still counted as GPU-enabled.

In light of these issues, software developers within the OLCF User Support Group were tasked in 2014 with defining and computing improved statistics for GPU activity on Titan.

By March, Cray's Resource Utilization Reporting (RUR) tool was successfully installed on Titan and began to generate GPU statistics logs for all compute jobs (i.e., apruns on Cray machines). RUR can report the exact number of GPU seconds a compute job has accumulated by interfacing with a GPU hardware timer. Each GPU exposes exactly one of these timers. By August, OLCF developers had programmed a new web application for storing and reporting on this new information. By correlating the RUR logs with other logs of interest, the developers were able to generate complete compute job records that were then archived into the new reporting application for future retrieval.

With these new tools in place, the following two statistics can now be computed on Titan for any time span and any subset of compute jobs:

1. Percentage of GPU-enabled time. This is defined as the percentage of consumed time, within a queried timeframe, that employed GPUs. This statistic is comparable to the previous year's calculated GPU-enabled statistic, but staff know that a job used a GPU instead of inferring that it did via its associated GPU libraries.
2. Percentage of GPU activity. This is defined as the percentage of compute job runtime during which a GPU was actually in use (applies only to GPU-enabled jobs).

The GPU activity statistics for different batch jobs can vary significantly based on the computational code employed and its algorithmic requirements. So, while not completely appropriate for drawing generalized system-wide conclusions, these statistics are useful to OLCF staff in identifying projects that could benefit from extra assistance in GPU performance optimization tasks.

Upon initial inspection of the GPU activity for a few selected scientific applications, inconsistencies were noted with respect to reported GPU activity versus expected GPU activity. Further investigation found that (1) the Cray RUR tool returns limited and/or incomplete information for batch jobs that terminate abnormally for instance, because of wallclock limits reached, hardware failure or software error), and (2) because of known limitations of the tool, RUR generally underreports GPU time for batch jobs that do not use an exclusive-process GPU mode. Because of the latter limitation, the OLCF reports GPU activity on exclusive-process compute jobs only; this compute mode comprises more than 93% of all job runtime on Titan for 2014, and approximately 83% of all GPU-enabled runtime.

Work to refine GPU statistics reporting continues. Currently, the OLCF is developing a customized RUR output plugin to streamline the collection of compute job data—which is currently spread across the many different logs—eliminating the need for complex correlation of log files. OLCF developers are also working to provide more GPU-centric reports and to expand the application's data exploration capabilities.

GPU activity statistics for the OLCF in 2014 are provided in Section 2, Business Results.

4.1.2 Understanding, Quantifying, and Analyzing GPU Reliability on Titan

The performance efficiency of GPUs is well understood. However, the reliability characteristics of GPUs in a large-scale computing system have received less attention. Assessing, understanding, and ultimately further optimizing GPU reliability will lead to greater scientific productivity and higher operational efficiency. Therefore, the OLCF has made a long-term investment in analyzing the reliability characteristics of GPUs and how they impact system operations and applications. OLCF staff recently conducted a large-scale field study on GPU error characterization, quantification, and impact. To the best of our knowledge, this is the first detailed study of GPU reliability on a large scale. Results appeared in the *Proceedings of the 21st International Symposium on High-Performance Computer Architecture*.[13]

The study shows that Titan GPUs experience failures at a very low rate, lower, in fact, than vendor estimates. High-standard acceptance tests and rigorous testing of error-prone cards on another cluster during the production phase have enabled the OLCF to identify and eliminate bad cards, increasing the mean time to application interruption significantly (more than 40 hours system-wide for CY 2014).

The rate of occurrence of soft errors that cause GPU applications to terminate (i.e., double bit errors) is fairly low as well (one per week). This has again been possible because of careful and proactive management of GPU cards that exhibit reoccurring bit errors. The OLCF staff performs stress testing of the GPU cards that exhibit double bit errors on a separate cluster before putting them back in the machine. The center also studied the impact of single bit errors on the GPU cards. Single bit errors do not impact an application, as they are automatically corrected by the ECC logic deployed in the GPU cards. It is interesting that almost 98% of the single bit errors have occurred in only 10 GPU cards (0.05% of the whole system). This suggests that a few cards may be significantly more prone to the recurrence of single bit errors. The center also demonstrated that the GPU cards showing high single bit errors can be identified early by checking the memory structure in which the errors are occurring. In addition, 98% of the single bit errors occur on the GPU L2 cache, not the GDDR5 memory. This finding is useful for both GPU architects and the system operations team in identifying such cards early during the production phase.

OLCF staff have also identified the impact of temperature on GPU system failures. A certain type of GPU system failure related to PCIe lane degrades is relatively more frequent in hotter cages. The operations team can manually schedule large-scale GPU jobs in the lower cages (which are relatively cooler in temperature than upper cages in a cabinet) so the large-scale jobs are less likely to experience temperature-related failure.

Staff have also conducted high-energy neutron-beam tests in collaboration with the Federal University of Rio Grande do Sul to study the resilience characteristics of the GPUs deployed on Titan (NVIDIA K20x GPUs). A neutron strike may flip one or more unprotected bits. The probability that a bit will actually be flipped depends on several factors, including the cell design and chip area. If a cell design is highly resilient, there is less likelihood a bit will be flipped. These radiation tests confirmed the finding about the stability of Titan GPUs and demonstrate that the GPU used in Titan is significantly more soft-error resilient than previous generations of GPUs designed by the same vendor.

The OLCF continues to work with vendors to improve current and future generations of GPUs by helping the vendors identify which memory structures in the GPU architecture need more protection in the event of soft errors, and showing how to identify cards more prone to soft errors early in production. OLCF staff have also found multiple inconsistencies related to vendor-provided GPU error logging tools and are working with the vendor to resolve those. Staff are also engaged with vendors to develop more fine-grained and low-overhead methods to collect GPU-related errors. This effort is expected to have significant impact on managing GPUs, their error characteristics, and their impact on production jobs. This study is also helping vendors estimate the impact of GPU errors on future-generation supercomputers (Summit and beyond).

---

[13] Devesh Tiwari, Saurabh Gupta, Jim Rogers, Don Maxwell, Paolo Rech, Sudharshan Vazhkudai, Daniel Oliveira, Dave Londo, Nathan Debardeleben, Philippe Navaux, Luigi Carro, Arthur S. Bland, "Understanding GPU errors on large-scale HPC systems and the implications for system design and operation," in the *Proceedings of the 21st International Symposium on High-Performance Computer Architecture* (HPCA), February 2015.

4.1.3 Functional Partitioning for Efficient End-to-End Computing on Heterogeneous Nodes

Smart resource utilization is critical to both optimal HPC operations and improved end-to-end computing of applications. Titan's heterogeneous node architecture presents a unique opportunity to address both these issues. Typically, the accelerators are the primary compute engines and the CPUs act as masters, orchestrating the offloading of work onto the accelerators and moving the result/output data back to the main memory. An ideal balance is not always achieved, and CPU usage may be proportionally greater than GPU usage or vice versa. Or, because of scaling constraints or resource bottlenecks, scientific applications may not use all of the compute cores present on the node.

The OLCF has developed a functional partitioning (FP) runtime framework that can optimize the usage of resources on a compute node to expedite an application's own end-to-end workflow.[14]

The FP runtime framework is initiated by the main application simulation to co-execute a set of tasks on a pre-specified number of CPU cores. The runtime framework creates a daemon process—the FP-agent—per node. Selected post-processing and analytics operations from an end-to-end workflow, called FP-task(s), are offloaded to use the unused resources (e.g., CPU, GPU, memory). The FP-agent is responsible for several resource management activities, such as setting up a communication infrastructure for the simulation processes and the FP-tasks to interact with each other, launching the FP-tasks, provisioning resources for the tasks, monitoring their progress, identifying slackening in GPU usage, and interference control.

Communication between the simulation and the FP-tasks is accomplished via a shared memory transport atop the Common Communication Interface (CCI). CCI provides a network abstraction allowing better performance, scalability, and low latency in large-scale deployments. This is essential for communicating simulation output to the FP-task and getting the intermediate results back from the FP-task. If the FP-agent detects that GPU resources are partially or fully available, and a GPU version of the same FP-task is also available, the FP-task can be scheduled on the GPU to maintain high GPU utilization. Therefore, the FP-agent is responsible for intelligent scheduling of tasks while minimizing the performance impact on the main scientific application through interference control.

In addition to the ability to launch data analytics as FP-tasks, the framework allows performance of I/O aggregation via an FP-task—thereby streamlining I/O—instead of multiple processes each performing I/O. To this end, staff used the Mercury I/O forwarding library that captures POSIX calls and forwards them to a process, and they developed a CCI plugin for the task.

Future tasks also can be imagined using the FP framework, such as feature extraction for visualization, health checks of a running simulation, and co-execution of reliability tasks with the main application simulation.

An example of the benefit of FP can be shown by recent work on the Community Earth System Model (CESM), one of the widely used applications at the OLCF. In addition to performing prognostic calculations, CESM computes a set of diagnostics used for further analysis at the conclusion of the experiment. The software used for post-processing and data analytics would often need to reinterpolate the model output to a standard structured grid before performing any analytics. Using the FP framework at a large scale with CESM, researchers can run the regridding tasks on output files from the ongoing simulation on underused CPU cores on the same node, in a pipelined fashion. This avoids queue wait time, I/O time to read the dataset, and computation time spent in post-processing the output files. In Figure 4.1, the baseline "CESM" bar shows performance. The next bar shows a typical scenario in which a CESM run is followed by several reinterpolations in series (post-processing time). Since these two tasks are performed on different clusters, a scheduling delay is also observed; it is variable and can be significant, depending on queue wait times. The last bar is the result when reinterpolation is performed in

[14] Karan Sapra, Saurabh Gupta, Ross Miller, Valentine Anantharaj, Scott Atchley, Sudharshan S. Vazhkudai, Devesh Tiwari, Melissa C. Smith, "End-to-End Computing using Functional Partitioning: A Community Earth System Model (CESM) Case Study", Smoky Mountain Conference, Gatlinburg, TN, Sept 2014. (poster)

situ using the FP framework. FP schedules the reinterpolation on the underused CPU cores, and a significant portion of the latency of reinterpolation is hidden successfully.

Thus OLCF developmental efforts and performance evaluation show that FP can enable higher node utilization, which leads to more efficient computing operations and overall higher scientific productivity. The OLCF approach is different from extant techniques on in situ analysis in that it provides a framework for on-the-fly analysis on-node and dynamically uses under-utilized resources therein.

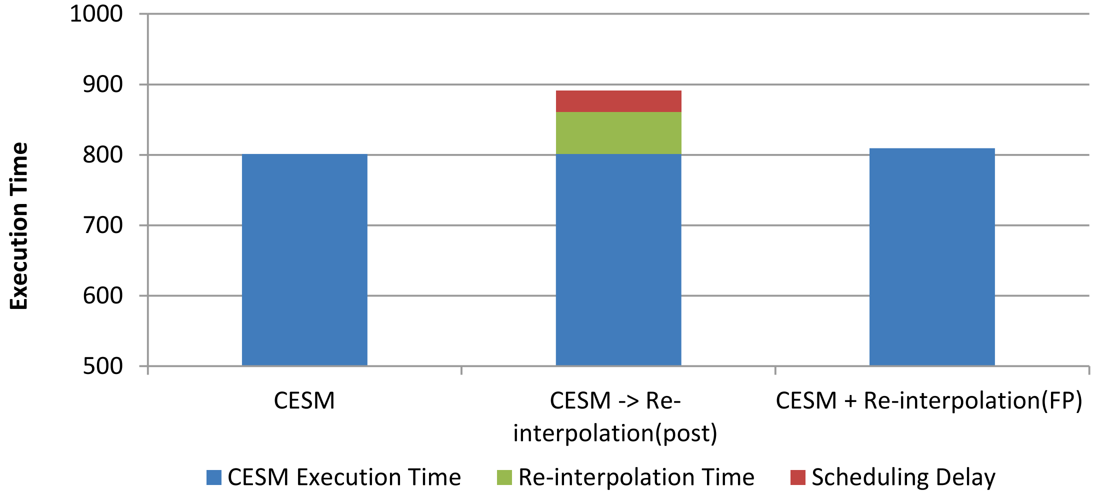

Figure 4.1. Breakdown of execution time for 320-node run of the atmospheric model in CESM for baseline (CESM), baseline followed by reinterpolation (CESM ≥ reinterpolation[post]), and CESM with FP performing in situ reinterpolation (CESM + reinterpolation[FP]).

4.2 DATA ANALYSIS AND WORKFLOW

Increasingly, the OLCF is coupling its unique computational and data resources with experiment and observation data across a broad range of scientific domains. This coupling is driven by a number of factors, including the need for large-scale simulation-based analysis, near real-time analysis requiring massive ensemble runs, and large-scale data storage resources.

Eleven DOE science data pilot projects were featured at the DOE national laboratory booth at the SC14 supercomputing conference. Seven of the 11 leveraged the capabilities of the OLCF. Although the workflow of each data pilot project had unique components, common requirements emerged; many of them could be met only by building upon the scalable computing and data storage capabilities of the OLCF. Six of those pilot projects are summarized in the following sections. The OLCF continues to invest in innovations to promote workflow. One of these activities—PYRAMID—is described in Section 4.2.7.

4.2.1 Near Real-Time Analysis of Experiment

Alexander Hexemer, staff scientist at Lawrence Berkeley National Laboratory (LBNL) and Craig Tull, group leader of the Science Software Systems Group at LBNL, demonstrated the use of Titan to facilitate near real-time analysis of organic photovoltaics (OPV) using x-ray scattering at the Advanced Light Source (ALS). As data were collected at the ALS, data movement to and subsequent analysis on Titan were triggered using more than 8,000 compute nodes running HipGISAXS[15], a massively parallel high performance x-ray-scattering data analysis code. This analysis was used to solve an inverse problem, allowing scientists to understand the OPV material structures from scattering data in the context of

[15] http://portal.nersc.gov/project/als/hipgisaxs/

theoretical models and then drive the next stage of the experiment while the data was being collected. Moving the data to ORNL made sense because only Titan has the computational capability to run HipGISAXS in real time with the data streaming from ALS experiments. This demonstration required co-scheduling of computational resources with the ALS experiment, remote triggering of analysis running on Titan, high-performance data transfer over the Energy Sciences network (ESNet) from the experiment end station, and near–real-time feedback of analysis results through a web portal interface.

4.2.2 Materials Science Data Analysis

Thomas Proffen, division director of Neutron Data Analysis and Visualization in ORNL's Neutron Sciences Directorate, demonstrated the use of Titan to explore and classify features in large volumes of neutron scattering data from the Spallation Neutron Source (SNS). Diffuse scattering contains information about disorder in materials, a critical component in understanding material function. Advances in diffuse scattering instruments have enabled movement from two-dimensional images to 3D volumes of measured scattering, making analysis of this data a significant challenge for scientists. Traditionally, interpretation of diffuse scattering data was manual and was reserved to a small number of experts in the field. A project at SNS demonstrated the use of image processing and machine learning techniques to identify features of interest in $Sr_3(Ru_{1-x}Mn_x)_2O_7$ and then classify these features into material properties such as stacking faults. Much of the analysis required by this project was conducted on Titan at the OLCF. Image processing and feature detection were built using the advanced I/O processing capabilities of the ADIOS system and the visualization capabilities of VisIt (Figure 4.2). Using the swift parallel scripting language, a workflow was constructed that enabled large-scale ensemble analysis with the DISCUS[16] simulation package. Using these capabilities, features identified using advanced image processing techniques were mapped to potential material defects in information from thousands of simulations (Figure 4.3).

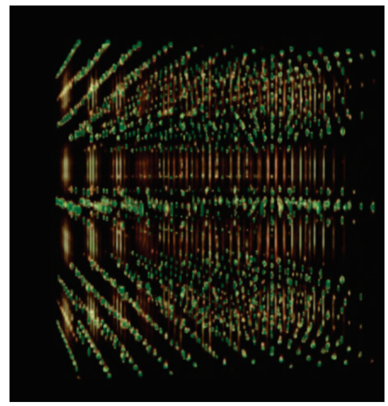

Figure 4.2. Feature identification in neutron scattering data using the ADIOS workflow coupled with the VisIt visualization environment.

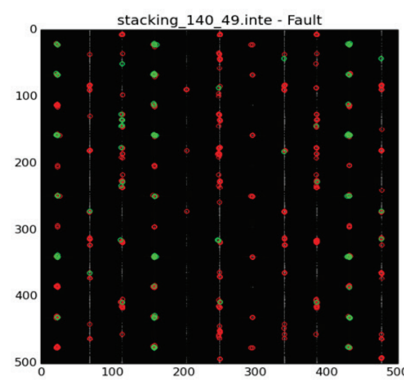

Figure 4.3. Classification of stacking faults.

This demonstration required advanced workflow capabilities coupled with Titan to enable near real-time feedback. Future work will enable remote triggering of analysis and subsequent feedback to the user at the beam line.

4.2.3 Near Real-Time Analysis

Rick Archibald of the Center for Applied Mathematics at ORNL and Sergei Kalinin of the Center for Nanophase Materials Science (CNMS) at ORNL demonstrated near real-time analysis of atomically resolved images using advanced mathematics for feature detection and classification. Using high-throughput scanning transmission electron microscopy, the ferroic oxide $LaCoO_3$ (a material relevant to next-generation battery technology) was imaged with atomic resolution, resulting in more than 1,000 images of the material. A single imaging of $LaCoO_3$ is illustrated in Figure 4.4.

[16] http://discus.sourceforge.net

The current state of the art at the CNMS required many hours and hands-on analysis of these images to capture essential properties of the material such as lattice structure and unit cell changes over time. Using a novel approach that coupled advances in image processing with modeled properties of the material structure, a new workflow was developed. This workflow incorporated scalable algorithms for Titan capable of identifying the essential properties of the material structure (Figure 4.5). Using this new workflow, processing time was reduced to 20 seconds. As new instruments are brought on line at the CNMS that generate thousands of individual images per minute, scalable analysis will be a critical component of experiments using those systems. Future work will extend these methods to additional unit cell configurations, resulting in a library of scalable analytics that can be triggered on demand as part of the experimental workflow at the CNMS.

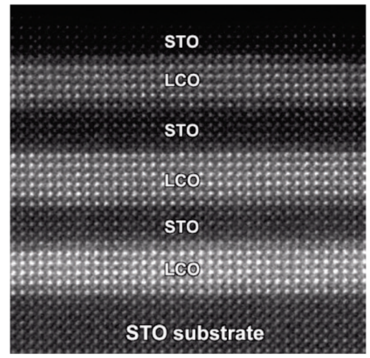

Figure 4.4. $LaCoO_3$ on $SrTiO_3$ (STO) substrate.

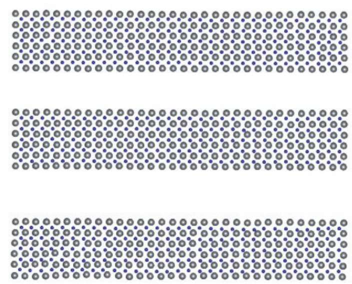

Figure 4.5. Identified lanthanum and cobalt atoms and lattice structure.

4.2.4 Analysis Workloads

Kenneth Read of ORNL's Physics Division demonstrated the use of Titan coupled with the worldwide Large Hadron Collider (LHC) computing grid for analysis that will improve measurements and understanding of the recently discovered Higgs boson and its decay modes and facilitate searches for supersymmetry. This coupling was made possible by recent advances in the ASCR-funded Big PanDA workflow management system (Figure 4.6). This workflow demonstration included remote job invocation on Titan from the LHC computing grid, real-time burst-mode job submission based on Titan compute cycle availability, high-throughput Geant4 and Monte Carlo simulation, and data transfer of results from Titan to tier-2 LHC computing grid sites for availability throughout the grid.

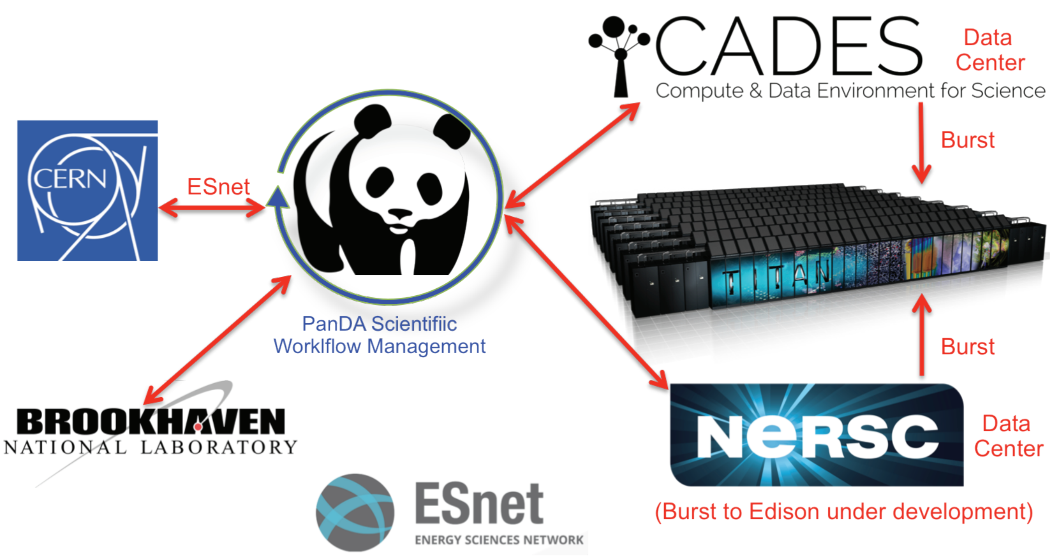

Figure 4.6. Big PanDA workflow, coupling opportunistic compute capacity on Titan with the worldwide LHC computing grid.

4.2.5 Cosmology Simulations

Peter Nugent of the Computational Research Division at LBNL demonstrated the use of Titan for one of the highest-resolution cosmology simulations to date to understand and constrain systematic uncertainties in cosmological observations. These and other simulation-based analysis approaches to understanding cosmological observations are of increasing importance, as the ability to generate massive observation datasets is far outstripping the ability to effectively analyze and understand them. Instruments and observation surveys such as the Dark Energy Survey, the Sloan Digital Sky Survey, and the Large Synoptic Survey telescope are generating catalogs that will range in size from one to several hundred PB. Simulation will play a critical role in these endeavors, from understanding the basic theory of cosmological probes to the production of "mock skies" used to test the observation/analysis chain for data processing workflows that identify, categorize, and catalog features of interest within these digital images (Figure 4.7).

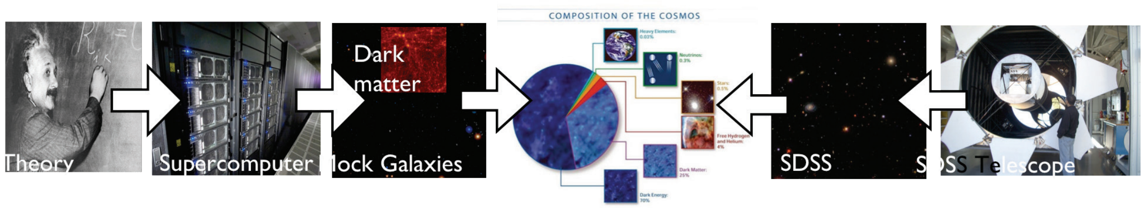

Figure 4.7. Roles of simulation in cosmology.

As part of this demonstration and of an awarded INCITE allocation, the OLCF was used to run a 0.55 trillion-particle cosmology simulation, from which Figure 4.8 illustrates a visualization. This simulation generated over 7 PB of raw data for subsequent analysis. This raw "level I" data contains particle and field information from the simulation campaign. The level I data was further analyzed to generate level II data that captured salient characteristics of interest at a much lower data volume. This data reduction technique enables effective exploratory analysis of the simulation by bringing the dataset to a more reasonably managed size of tens of TB. Once reduced, the dataset was made available for collaborative analysis via a virtualized computing environment based on Docker + OpenStack.

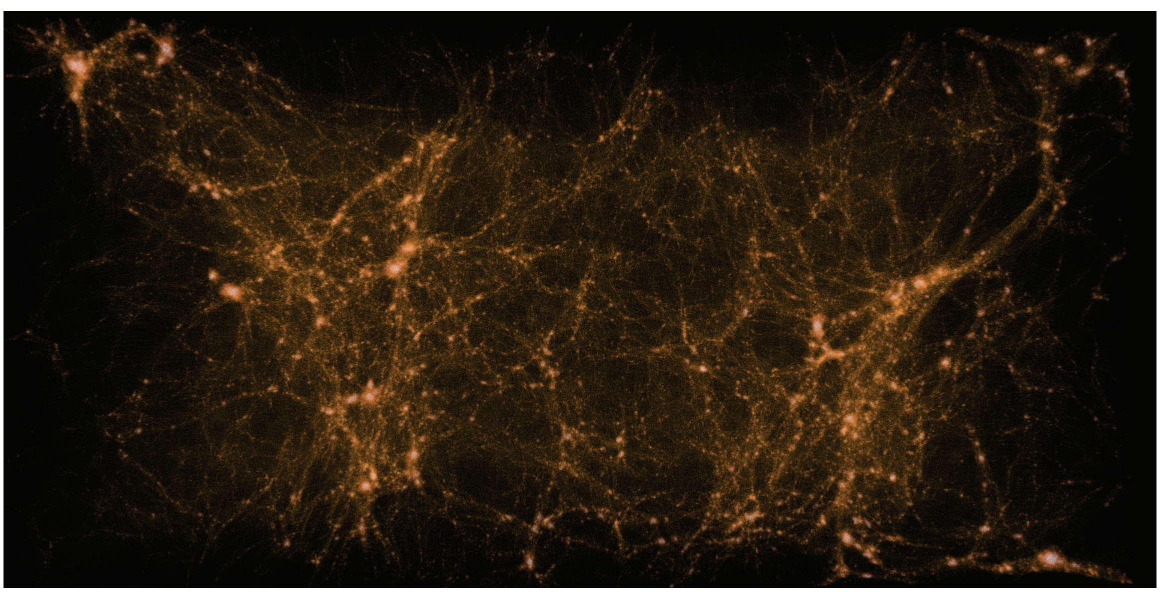

Figure 4.8. Visualization of 1/16384 of the volume of the 0.55 trillion-particle simulation on Titan.

4.2.6 Complex Nanostructures

Simon Billinge of the Condensed Matter Physics and Materials Science Department at Brookhaven National Laboratory demonstrated the use of Titan coupled with scattering data from the National Synchrotron Light Source, SNS, and the Advanced Photon Source to elucidate the fundamental structure of complex nanostructures. Traditional crystallography provides a mechanism to identify material structure based on x-ray, neutron, or electron scattering by solving the inverse problem for a simulated scattering pattern. This approach works well for crystalline structures but quickly becomes intractable for more complex nanostructures in which a one-dimensional scattering pattern is insufficient to constrain the set of possible 3D nanostructures that could give rise to the pattern. For example, a single scattering pattern of a small tetrahedral CdSe quantum dot of 87 atoms used in this demonstration could result from more than 1,042 different nanostructure configurations when uncertainty bounds are unconstrained. Using a novel optimization algorithm known as Liga to constrain the solution space through an optimization of optimizations approach drops the number of candidate solutions to 103. When additional experimental data is added (such as small-angle scattering), the problem is further constrained, providing a tractable set of potential solutions to the inverse problem for further analysis (Figure 4.9).

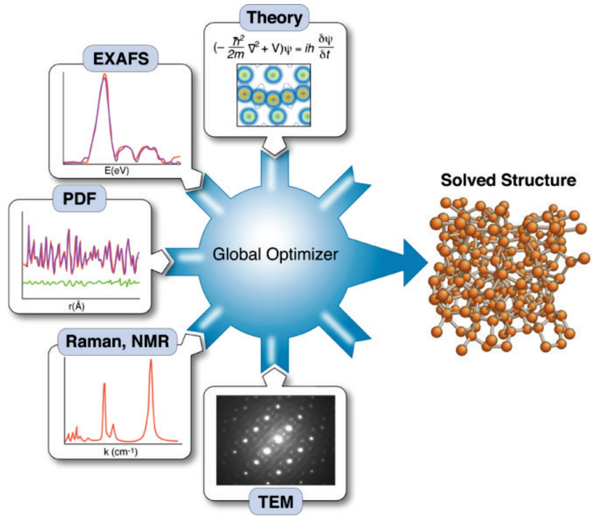

Figure 4.9. Schematic of complex modeling approach.

This type of complex modeling requires significant computational capabilities available only at leadership computing facilities. In this demonstration configuration, 3 weeks of computation time on a modern desktop computer is required to generate and evaluate a single candidate solution. With potentially thousands of solutions, that is intractable. Using Titan, the candidate solution space can be adequately explored in a single day (Figure 4.10).

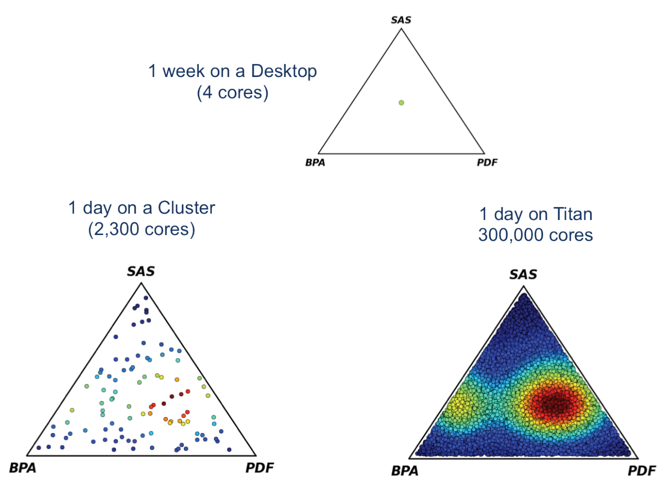

Figure 4.10. Scaling the candidate solution space evaluation.

4.2.7 PYRAMID: Building System Support for Complex Application Workflows

It is natural, for a large supercomputer like Titan, to think primarily of user workflows as comprising primarily single large jobs that are queued and executed through the batch system. But the reality is that most computational science research involves a mix of jobs—numerous smaller jobs that build up to the large individual jobs that we traditionally think of as “supercomputing.” More and more, these simulations are inter-related in more complex ways. A collection of jobs might comprise a “parameter sweep,” in which selected parameters of the simulation are systematically varied to explore the response or quantify the uncertainties in the simulation results, or optimizations to hone in on specific parameter values. Sometimes the workflows are more complex. For example, a simulation might dynamically spawn subtasks to refine regions of the simulation, material property information, or other key data. Or a loosely coupled multiphysics simulation may be able to expose more parallelism if multiple components, and even multiple time steps, are executed concurrently (while satisfying the necessary data dependencies, of course). Regardless of how it arises, the need to launch and manage multiple simulation tasks simultaneously within a single batch job, giving each a portion of the total number of nodes requested, is a common feature of an increasing number of user workflows . Unfortunately, the native task launch mechanism on Titan, Cray/ALPS, was designed more for a traditional supercomputing workflow with small numbers of large simulation tasks per batch job. Attempting to manage workflows with large numbers of smaller concurrent tasks can lead to resource exhaustion on the service nodes, adversely impacting not only the user with the complex workflow, but also other users whose jobs happen to be relying on that service node.

The goal of PYRAMID is to provide a lighter-weight alternative to the native Cray/ALPS task management system, to support complex workflows with more concurrent tasks. It is based on the Scalable runTime Component Infrastructure (STCI) runtime system. STCI has been under development with OLCF support for some time with the primary goal of providing a more resilient runtime to support a variety of tool and communication middleware needs. STCI relies on a hierarchy of agents that manage different aspects of the lifecycle of the computational job. The PYRAMID project has extended STCI to two purposes: (1) support persistence of key portions of the runtime infrastructure across multiple computational tasks and (2) provide two-level resource management, which separates the assignment of job resources to computational tasks and the deployment of the task to the assigned resources. Developers also added the capability to deploy multiple tasks to a node, as well as tools to check the status of individual tasks. Experiments comparing the enhanced STCI infrastructure with the ALPS infrastructure show 14× to 18× better throughput for many short tasks, with constant memory usage on the login node (STCI), as opposed to increasing memory usage (ALPS). The PYRAMID approach reduces the impact on login/service nodes as the number of tasks in the workflow increases, which avoids overloading these shared resources.

4.3 I/O EFFICIENCY

4.3.1 Balanced Placement of I/O

The Spider file system offers a high peak throughput. However, applications do not always achieve such high rates for a variety of reasons, such as contention at the shared storage, diverse access patterns, and delays at multiple levels from the compute node to the storage. A large I/O subsystem such as Spider has been observed to suffer from severe contention, and there exists a significant load imbalance among the different components in the storage system.

To address the I/O load imbalance and contention issues, OLCF staff have developed a topology-aware, balanced I/O placement strategy (libPIO) based on a site-defined, tunable, weighted cost function of selectable resources.[17] The strategy is topology-aware because it assigns weight factors to separate resource components depending upon which storage layer they belong to, as different storage layers have different degrees of impact on the end-to- end performance. It is balanced, as it keeps track of usage of all storage components and balances the load along the end-to-end I/O path. For evaluation purposes, the algorithm was implemented as an easy-to-use, user-space library. The OLCF performed extensive experiments of libPIO on Titan and the Lustre-based Spider II file and storage system.

Results demonstrate, with both synthetic benchmarks and a real-world scientific application (S3D), that the proposed strategy can significantly improve I/O throughput regardless of the layout of the compute node allocation. Synthetic benchmark results indicate that the proposed strategy can improve I/O

---

[17] Feiyi Wang, Sarp Oral, Saurabh Gupta, Devesh Tiwari and Sudharshan Vazhkudai, “Improving large-scale storage system performance via topology-aware and balanced data placement,” *Proceedings of the 20th IEEE Int'l Conference on Parallel and Distributed Systems* (ICPADS), Hsinchu, Taiwan, December 2014.

performance by up to 50% and even more in some cases. For example, libPIO was easily integrated with a large-scale scientific application, S3D, a high-fidelity turbulent reacting flow solver. S3D writes the state of the simulation to the file system, which is later used for analysis. Only 30 lines of code needed to be added/modified in the checkpoint subroutine of the application to integrate the placement library. Primary additions were init(), nid2ost(), and finalize(). An lfs setstripe call was also added to provide the desired placement of output files on the Lustre OSTs. Evaluation results show up to 20% improvement in I/O performance at fairly large-scale runs (3,750 nodes). Although tests were conducted in a production environment, substantial gains in I/O performance were observed. The ease of implementation with minimal code changes suggests that libPIO can be widely adopted by scientific users and middleware I/O libraries.

4.3.2 IOSI: Automatic I/O Signature Extraction from Noisy Server-Side Traces

Competing workloads on a shared storage system like Spider cause I/O resource contention and application performance vagaries. This problem is already evident in Spider, and is likely to become acute at the exascale. More interaction between application I/O requirements and system software tools will help alleviate the I/O bottleneck and move toward I/O-aware job scheduling. However, this requires rich techniques to capture application I/O characteristics, which remain evasive in production systems.

Traditionally, I/O characteristics have been obtained using client-side tracing tools, with drawbacks such as non-trivial instrumentation/development costs, large trace traffic, and inconsistent adoption. The OLCF has developed a novel approach, I/O Signature Identifier (IOSI), to characterize the I/O behavior of data-intensive applications. IOSI extracts signatures from noisy, zero-overhead server-side I/O throughput logs that are already collected on Spider, without interfering with the compiling/execution of applications. IOSI was evaluated using the Spider storage system, the S3D turbulence application, and benchmark-based pseudo-applications. Experiments confirmed that IOSI effectively extracts an application's I/O signature despite significant server-side noise. Compared with client-side tracing tools, IOSI is transparent and interface-agnostic and incurs no overhead. Compared with alternative data alignment techniques (e.g., dynamic time warping), it offers higher signature accuracy and shorter processing times.[18]

The OLCF team was the first to study storage-server-side logs for signature extraction in the HPC domain. The extracted signature can be used in the future to perform I/O-aware scheduling tasks, such as staggering applications so as to not interfere with one another's I/O.

4.3.3 Lazy Checkpointing

Checkpointing is a process in which a long-running application will periodically save information about the state of its progress to permanent storage. The perception among users regarding how frequently they should checkpoint varies widely and subjectively. Excessive checkpointing means less time spent solving the problem at hand, and more time generating I/O on a shared resource. The balance is delicate, as system efficiency is compromised by any unnecessary I/O activity. As computing resources approach exascale—a hundred-fold increase over current supercomputers—increasing or even maintaining efficiency during checkpointing will be paramount.

---

[18] Yang Liu, Raghul Gunasekaran, Xiaosong Ma, Sudharshan S. Vazhkudai, "Automatic identification of applications I/O signatures from noisy server-side traces," *Proceedings of the 12th USENIX Conference on File and Storage Technologies* (FAST '14), Santa Clara, California, February 2014.

OLCF staff have designed a new checkpointing technique that specifically addresses this issue.[19] The checkpointing technique, Lazy Checkpointing, reduces the I/O overhead and thus lowers the barriers for OLCF users to adopt checkpointing for their large-scale runs. Lazy Checkpointing takes advantage of the key trait that failures have a strong temporal locality. The probability of a failure is high soon after a failure has occurred (i.e., more failures occur on the heels of a failure). The temporal locality in failures indicates that a significant fraction of failures is likely to occur within a relatively short time period after a failure occurs, compared with the MTTF of the system. Therefore, as the time interval after the last observed failure increases, the necessity to checkpoint may decrease. Of course, there is a limit to every good thing—Lazy Checkpointing puts upper bounds on how relaxed one can be before disrupting the balance of checkpointing versus an eventual hardware failure.

The OLCF developed detailed analytical models and simulations to study the efficacy of Lazy Checkpointing. Results indicate that Lazy Checkpointing reduces the I/O overhead significantly, by up to 30%, as the need for checkpointing decreases as time elapses between two failures. A prototype implementation of Lazy Checkpointing requires, as input, different system and job-related parameters such as the job size, the checkpoint size, the OLCF file system performance, and the MTTF for the Titan system. The facility is reaching out to users of OLCF applications to help them perform smarter checkpoints while running their applications on Titan.

4.3.4 I/O Harness: New Testing Procedures and Introduction of a Systems Evaluation Team

I/O efficiency can be materially affected in unexpected ways. The OLCF staff members' innovative response to improving I/O efficiency resulted in a new operational procedure and team, which has been shared as best practices.

Titan is unique in its scale and its computing ecosystem, which includes the Atlas center-wide file system that serves the high-speed storage needs of the OLCF for seamless data access between computing and visualization in the scientific workflow. The Atlas file system went through a major hardware and Lustre software upgrade to enable improved performance, user functionality, and operational improvements and was deployed to users in January 2014. Although significant testing was performed using synthetic tests and benchmarks on Atlas before deployment, bugs not seen in earlier testing manifested because of the size of the workload and diversity of applications running on Titan during production. User applications were adversely impacted by severe performance degradation. The OLCF staff rolled back the Lustre software to a known, stable version to restore functionality to Titan.

To mitigate future system degradation due to hardware and/or software new to the Titan ecosystem and to find a stable path to upgrade Atlas as desired, the OLCF created an integrated project team. The purpose was to provide a more robust systems testing approach that more closely mimics user workload and behavior on Titan. The team consisted of members from four groups at the OLCF: Scientific Computing, User Assistance, HPC Operations, and Technology Integration. A new application-centric test suite was created, hardened, and used repeatedly during a 9-month period to identify the issues with the Lustre software. The integrated team worked with the software developers and identified a stable and performant version that was successfully deployed in October 2014.

Since system maintenance and testing competes with providing production hours to OLCF users, testing windows are compressed to 1 day per week; and the OLCF integrated project team often must work late hours to make optimal use of the testing time. The new OLCF systems evaluation team uses this more robust systems testing approach for significant software upgrades to the OS and other programming environment changes to improve user productivity and overall experience on Titan. As seen in Section 1, User Results, overall user satisfaction in 2014 reached its highest level ever. The increase in satisfaction is

---

[19] Devesh Tiwari, Saurabh Gupta, Sudharshan S. Vazhkudai, "Lazy checkpointing: Exploiting temporal locality in failures to mitigate checkpointing overheads on extreme-scale systems," *Proceedings of the 44th Annual IEEE/IFIP Int'l Conference on Dependable Systems and Networks* (DSN 2014), Atlanta, June 2014. The paper was one of the three finalists for the best paper award.

in no small part due to the responsiveness and proactiveness of the OLCF in identifying and resolving unanticipated challenges.

4.4 APPLICATION DEVELOPMENT

As described in Section 1, User Support, the OLCF has an ongoing effort to assist users in porting code to diverse architectures and preparing them for next-generation resources. The OLCF also invests in innovative tools to promote this activity. One such tool, HERCULES, is used to identify patterns in code and systematically transform them in ways that facilitate acceleration, as are MiniApps designed to provide proxies for application testing. Also important to facility operations is the development of innovative solutions for internal applications. In 2014, OLCF staff were recognized by ORNL for their cross-division activities to enhance the OLCF resource tracking system.

4.4.1 HERCULES: Helping Users Understand and Transform their Software

Writing software to run at the largest scales on cutting-edge machines is challenging. Adapting existing software to run well in new environments can be even more so—the code base must be understood sufficiently to identify the code regions requiring attention, and then those areas must be transformed while their context is maintained within the larger program. Adapting a large code that might have been written largely by others, and a hardware platform that is new to them, such as Titan with its GPU accelerators, can be an overwhelming task. Performance analysis tools are often used to identify the "hot spots" in the code that could benefit most from acceleration, but such tools do not help rewrite the code to actually use the accelerators (for example). That is the purpose of HERCULES.

HERCULES is a compiler-based tool for identifying patterns in code and systematically transforming them. It can assist with the porting process in a very direct fashion. It can scan the code for patterns that have already been identified as being amenable to acceleration, thus helping the user identify and understand the code base. Further, it can transform the identified code regions to actually run on the accelerator.

HERCULES is quite general and can be used for a wide range of purposes that involve searching code for patterns and transforming them. Patterns could, for example, identify conditional statements within loops (which might reduce the performance of the loop on accelerator or vector architectures) or highlight cases in which the return codes from calls to well-known library functions (such as MPI calls) are not being checked (a robustness issue). Restructuring code for performance reasons (optimization) or to improve its quality or maintainability (refactoring) are common concerns for software developers. A new capability added to HERCULES in 2014 provides scriptable loop transformations that can help restructure complex loop nests to adapt to different memory hierarchies and accelerator work distribution schedules. Developers often try to use simpler text-based scripting tools to automate widespread transformations or searches. But this approach can be error-prone because such tools do not understand the syntax and semantics of the target code, and purely text-based patterns can easily be fooled by differences in formatting of the code, as opposed to its semantic content. As a compiler-based tool, HERCULES can properly interpret the code, ensuring that searches and transformations can be applied systematically and reliably across large code bases.

Initially developed under ORNL's Laboratory Directed R&D (LDRD) program, and more recently supported by the OLCF, HERCULES continues to grow as new patterns and transformations are added and researchers find new ways to take advantage of its capabilities.

4.4.2 MiniApps: Preparing Applications for Future Architectures

Along with efforts carried out under the aegis of CAAR to prepare applications for future architectures, the OLCF has undertaken to ensure that the operational experience gathered by the center is fed back into the co-design process, including interactions with vendors, computer scientists, and applied mathematicians. Bronson Messer is leading an ORNL LDRD project involving members of the Scientific Computing group that will produce and promulgate a set of reduced, proxy applications (MiniApps) based on large-scale application codes currently supported at OLCF. The first set of OLCF MiniApps includes examples derived from the Denovo, AORSA, and Chimera codes. The MiniApps are designed to encapsulate the details of the most important (i.e., most time-consuming and/or unique) facets of their progenitor applications. One focus of the project is to determine the extent to which the individual MiniApps reflect the "true" performance of the modeled applications—determined by making direct comparisons to the original application performance characteristics, a task Scientific Computing is uniquely qualified to undertake. Because of their relative compactness compared with a production application, the MiniApps will be used by the OLCF in several ways:

- Evaluation of programming models. Multiple versions of each MiniApp are under development using several different specific programming models (e.g., OpenACC, CUDA, OpenMP) to make quantitative statements regarding the performance gains or losses inherent in the choice of a programming model.
- Evaluation of performance portability between architectures. Because performance portability between next-generation architectures is becoming increasingly important, the MiniApps will be an avenue for evaluating the software development effort necessary to achieve portability.

The project expects to release the produced MiniApps as open source software in 2015, allowing for future significant interaction between external research communities and vendors without incurring the significant learning costs a production application often requires. In the past year, the team has produced an initial set of MiniApps, including versions written with OpenMP, OpenACC, and CUDA. It has performed a series of performance tests on these versions, comparing their performance with that of the original application codes on the Cray XK7, Cray XC30, and an Intel Phi-based cluster at NICS. In 2015, they will produce a full set of implementations—using all of the applicable programming models—for each MiniApp.

4.4.3 CI/D System for Automated Testing and Improved Deployment

ORNL recently recognized NCCS staff for their achievement in promoting the development, testing, and deployment of internal mission-critical applications. As the complexity of production codebases has grown, so have the inherent challenges around how to write, review, test, and deploy code quickly, efficiently, and securely. These challenges dramatically affect the OLCF's Resource Allocation Tracking System (RATS), an internal web application that stores much of the NCCS operational data.

In 2014, members of the User Assistance and HPC Operations groups formed a small DevOps Working Group. On May 20, 2014, they successfully tested and deployed a new version of the RATS application into production using the new procedures. The outcomes have been dramatic, and RATS is now able to perform the following tasks.

- Integrate application code development, version control, and automated testing environments. As a result of improved version control, application releases can now be rolled back easily if major problems are encountered, preventing service downtimes and unwanted emergency prioritization of critical application fixes.
- Queue up code review for important application code changes. Peer code review greatly increases the maintainability and security of the NCCS mission-critical applications.
- Validate application correctness before deployment. The CI/D system launches a suite of regression tests against the application's code base automatically every time new code is added to the repository. Automated regression testing greatly increases the reliability of the NCCS mission-critical applications.

- Decrease application deployment time. Before the new DevOps procedures were in place, total deployment time for a new RATS version, including the troubleshooting of related deployment issues, could take up to 250 FTE hours over the span of a week. With the new procedures in place, major changes to RATS are typically deployed in 1 FTE hour.

The integration of RATS into the CI/D system has made the application more flexible, secure, and efficient. The OLCF spends significantly less time per feature request throughout the development and deployment lifecycle, and development mistakes have been minimized. Issues that are important to facility management, such as reporting flexibility, are now much simpler and safer to implement. Perhaps as important, better communication exists within division groups as a result of the cross-cutting nature of these projects, and the OLCF as a whole is able to better prioritize internal development tasks.

Risk Management

HIGH PERFORMANCE COMPUTING FACILITY
2014 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2015

5. RISK MANAGEMENT

CHARGE QUESTION 5: Is the facility effectively managing risk?

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

The major risks currently being tracked are listed and described. Any mitigation planned for or implemented are included in the descriptions. As of this writing, the OLCF has two "high" operational risks:

1. The facility will continue having difficulty finding sufficient staff. To address the risk, staff recruiting and retention efforts have received increased emphasis.
2. OLCF's proportion of the shared costs to run and maintain the Building 5600 Central Energy Plant (CEP) will increase significantly if other participating users do not replace decommissioned computing systems. Some CEP equipment may have to be taken out of service to keep overall CEP costs down.

5.1 RISK MANAGEMENT SUMMARY

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

Operations and project meetings are held weekly and risk, which is continually being assessed and monitored, is usually discussed at the meetings. At least monthly, specific risk meetings are held, attended by the federal project director, facility management, OLCF group leaders, and others as required. 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 are stated in the "Risk Notes" narrative section of the register when appropriate. Early and late risk impact dates are recorded as well. Risk owners are expected to be proactive in tracking any trigger conditions and the impact horizons of the risks for which they are responsible, and to bring appropriate attention to management of those risks, whatever the risk-rating level.

The OLCF reports current high- and medium-level risks to the DOE program office as part of its monthly operations report. At the time of this writing, 25 active entries are in the OLCF operations risk register. They fall into two general categories: risks for the entire facility and risks particular to some aspect of it. Across-the-board risks are concerned with such things as safety, funding/expenses, and staffing. More focused risks are concerned with reliability, availability, and use of the system or its components (e.g., the computing platforms, power and cooling, storage, networks, software, and user interaction). In addition to operational risks, at the time of this report, there are 43 tracked risks for the OLCF-4 project.

Costs for handling risks are integrated within 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 CY 2014

5.2.1 ID# 1006—Inability to Acquire Sufficient Staff

Risk Owner	Arthur S. Bland, OLCF Project Director		
Probability	High		
Impact	***Cost:*** Low	***Schedule:*** Low	***Scope/Tech:*** Low
Rating	HIGH		
Status	Accepting the risk		

The OLCF has difficulty in acquiring adequate qualified staff because of a highly competitive job market. The risk is that desired work outcomes will not be achieved; some important tasks may be postponed or eliminated, and/or more current staff will become dissatisfied from overwork or missed opportunities to work on preferred assignments. The effect could be missed performance metrics, user dissatisfaction, or increased staff dissatisfaction.

Trigger: Open positions >10% of available positions

Although the cost, schedule, and technical impact ratings are all Low, the risk is rated High because of "other" impacts, such as those to the OLCF's or ORNL's reputation as a preferred place to work.

The OLCF has increased its emphasis on both recruitment of new staff and retention of existing staff. Should management become aware that work outcomes might be impaired, temporary help may be obtained from other ORNL resources, or contracts may be sought with external sources.

5.2.2 ID# 361—Scientists Decline to Port to Heterogeneous Architecture

Risk Owner	Jack C. Wells, NCCS Director of Science		
Probability	Medium		
Impact	***Cost:*** Low	***Schedule:*** Medium	***Scope/Tech:*** Low
Rating	MEDIUM		
Status	Mitigating the risk		

Common to all programming models is the need to structure and/or restructure codes to express increased hierarchical parallelism on today's hybrid multicore architectures. This is necessary on all high-performance architectures to achieve good performance. Beyond this restructuring, a user needs to use relatively new programming models to "offload" the computation to the GPU in GPU-accelerated hybrid architectures. The risk is that some users will decline to port or will delay porting of their applications to this new architecture because of the difficulty or cost. As a result, the OLCF would expect to see a decrease in the number and/or quality of proposals submitted to allocation programs such as INCITE and ALCC.

Trigger: A decrease in the number and/or quality of proposals submitted to allocation programs such as INCITE and ALCC

The original risk evaluation rated this risk as High. Mitigation with outreach, training, and the availability of libraries and development tools has ameliorated some initial user resistance. The marked improvement of compiler directive technology from Cray, CAPS, and PGI (including the OpenACC standardization) is overcoming some technical barriers for computational scientists to port and achieve acceptable performance running on hybrid, accelerated architectures. Additionally, the Software Tools team is leveraging LDRD and other investments to develop tools to assist users in porting their codes. Of the 30 INCITE projects awarded time at the OLCF for 2014, 16 had a computational readiness score greater than or equal to 4 out of 5. Likewise, the 30 INCITE projects awarded time at the OLCF for 2015, 15 had a CR score greater than or equal to 4 out of 5. To merit a CR score of 3 or greater, the applicant must demonstrate significant efforts in porting to the heterogeneous architecture. Many applications teams appear to be porting their codes.

5.2.3 ID# 407—Loss of Key Personnel

Risk Owner	Arthur S. Bland, OLCF Project Director		
Probability	Medium		
Impact	***Cost:*** Low	***Schedule:*** Low	***Scope/Tech:*** Low
Rating	LOW		
Status	Mitigating the risk		

Much of the effort within the OLCF is provided by highly trained and highly experienced staff. The loss of critical skill sets or knowledge in certain technical and managerial areas may hinder ongoing progress.

Trigger: Intelligence on potential or actual loss of key personnel

The center will increase recruiting and hiring in some key areas to maintain a qualified workforce and provide the planned 78 FTE level of effort. It will also continually work on succession plans for key staff through which the backup person is at least partially trained in the skills needed to complete the work to minimize the impact of a loss. Careful planning, along with its documentation in as much detail as possible, helps to mitigate the impact of a loss.

5.2.4 ID# 1063—Programming Environment Tools May Be Insufficient

Risk Owner	David E. Bernholdt, Group Leader, Computer Science Research		
Probability	Medium		
Impact	***Cost:*** Medium	***Schedule:*** Low	***Scope/Tech:*** Low
Rating	MEDIUM		
Status	Mitigating the risk		

This risk was formerly ID 906. It was recategorized. Titan relies on GPU accelerators for the bulk of its computational capability. The programming environment may not provide users with tools with which they are familiar, comfortable, and experienced and may not offer the levels of performance expected on the new system. If the programming environment is not productive for the users, they may withdraw from using the OLCF in favor of other centers.

Trigger: Concerns reported by user-application liaisons

The OLCF created a Software Tools Group within the NCCS to own the problem. The center surveyed users on their requirements in this area and on the adequacy of the tools available or planned. It found that for most of the primary tools from the OLCF-2 environment, there are plans to extend useful functionality for the next-generation system. Where it found gaps, the center initiated contracts with vendors to accelerate their development and to add key functionality needed. The center has developed portable programming models (through vendor partners) such as the directive-based OpenACC standard and the OpenMP directives for accelerators. Today, the OLCF is a member of the OpenMP and OpenACC standards committees to push for needed improvements and eventual consolidation of those programming standards. It is also offering training to its users in how to use the programming models as well as the programming tools. The center has also contracted with Allinea and Dresden to have on-site user support to assist users with the tools.

5.2.5 ID# 917—Robust Support Will Not Be Available to Ensure Portability of Restructured Applications

Risk Owner	Tjerk Straatsma, Group Leader Scientific Computing		
Probability	Medium		
Impact	***Cost:*** Medium	***Schedule:*** Low	***Scope/Tech:*** Medium
Rating	MEDIUM		
Status	Mitigating the risk		

The programming model that the center proposes requires a restructuring to utilize the standard distributed memory technologies in use today (e.g., MPI, Global Arrays) and then a thread-based model (e.g., OpenMP or Pthreads) on the node that captures larger-granularity work than is typically done in current applications. In the case of OpenMP, the compiler can facilitate and optimize this thread level of concurrency. This restructuring is agnostic to the particular multicore architecture and is required to expose more concurrency in the algorithmic space. OLCF experience to date shows that the center almost always enhances the performance with this kind of restructuring. The use of directives-based methods will allow the lowest level of concurrency to concomitantly be exposed (e.g., vector- or streaming-level programming). This means that the bottom level of concurrency can be directly generated by a compiler. The center expects that this kind of restructuring will work effectively with portable performance on relevant near-term architectures (e.g., IBM BG/Q, Cray Hybrid, and general GPU-based commodity cluster installations). However, restructured applications will be able to make use of several programming models—CUDA, OpenCL, OpenACC, or even PTX and other library-based approaches (e.g., OLCF's Geryon library)—to expose the lowest (vector-like) level of concurrency.

The risk is that robust versions of OpenACC will not be available for other contemporary platforms. Also, there is a risk that OpenCL could be lacking on OLCF-3's platform and OpenCL would remain lacking on the Titan platform.

The effect would be that applications run on Titan could be developmental "dead ends" because of poor performance, lack of a full set of features, or other problems. Users would have to work around these issues or change programming models.

Trigger: Intelligence on deficiencies in support applications

Multiple instantiations of compiler infrastructure tools will be adopted to maximize the exposure of multiple levels of concurrency in the applications. This will be complemented by publishing the case studies and experience gained from working with the six CAAR applications, coupled with the appropriate training of the OLCF user community. Work with vendors continues to improve compiler technology and other tools. The OLCF, ALCF, and NERSC are working together on code portability. The three centers have developed a plan for collaboration and workshops and training for users.

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

5.3.1 ID# 124—Storage System Reliability and Performance Problems

Risk Owner	Sudharshan S. Vazhkudai, Group Leader, Technology Integration		
Probability	Low		
Impact	***Cost:*** Low	***Schedule:*** Low	***Scope/Tech:*** Medium
Rating	LOW		
Status	Accepting the risk		

Hardware or software bugs could cause the storage system to exhibit reduced reliability and performance.

Trigger: Acceptance testing identifies issues in the file system

Significant hardware issues occurred in FY 2013. OLCF worked with the vendor to address the redesign and replacement of the motherboards on the object storage servers in FY13Q4, at no cost to the OLCF. Risk remains active because it continues to be a concern. It has been a year since the storage system upgrade and several months since the controller upgrade. The storage system has been stable. The potential for additional hardware or software problems remain, therefore this risk item remains active.

5.3.2 ID# 359—Lustre Updates May Cause File System to Become Unstable at Larger Scales

Risk Owner	Sudharshan S. Vazhkudai, Group Leader, Technology Integration		
Probability	Low		
Impact	***Cost:*** Low	***Schedule:*** Low	***Scope/Tech:*** Low
Rating	LOW		
Status	Mitigating the risk		

The acquisition of the dedicated Lustre efforts by larger companies is relevant to the issue of long-term support for Lustre, which is critical to the success of the OLCF. The lack of long-term support could have seriously hindered the ability to field a production-quality file system for the OLCF machines.

Trigger: Intelligence of loss of support or another corporate acquisition threat or action

The OLCF has been conducting large-scale test shots of Lustre 2.5 throughout FY 2014 to test it at scale and perform bug fixing. The center has pushed several of its patches to Intel for adoption into the main source tree. The OLCF has deployed Lustre 2.5 on the Atlas file system and has been testing it for several months. The Atlas file systems have gone into production on several clusters and on Titan in October of 2014. As these systems begin to use Atlas, the center will be able to identify any instability at a larger scale. OLCF staff have seen a few issues with Lustre at scale, including in January of 2014 and October of 2014. They are continuing to work with the Lustre development community and are monitoring the performance and stability of Lustre.

5.4 RISKS RETIRED DURING THE CURRENT YEAR

5.4.1 ID# 1101—Funding Inadequate to Cover Projected Spend Plan (FY 2014)

Risk Owner	Arthur S. Bland, OLCF Project Director		
Probability	Medium		
Impact	***Cost:*** High	***Schedule:*** Medium	***Scope/Tech:*** High
Rating	HIGH		
Status	**RETIRED**: Risk did not occur		

Annual budgets are set with guidance from the ASCR office, but actual allocated funds are unknown until Congress passes funding bills. Continuing resolutions have been common, and often several months elapse before actual funding is resolved.

Trigger: Intelligence on congressional or DOE funding capabilities and priorities

The risk did not occur. In January 2014, Congress passed and the president signed an appropriations bill meeting the budget requirements of the OLCF.

5.5 MAJOR NEW OR RECHARACTERIZED RISKS SINCE LAST REVIEW

5.5.1 ID# 721—Lustre Metadata Performance Continues to Impact Applications

Risk Owner	Sudharshan S. Vazhkudai, Group Leader, Technology Integration		
Probability	Medium		
Impact	***Cost:*** Low	***Schedule:*** Low	***Scope/Tech:*** Low
Rating	LOW		
Status	Mitigating the risk		

Metadata performance is critical to a wide variety of leadership applications. Its performance depends on many factors, all of which need to be optimized. Lustre performance has been stymied by not being able to scale beyond a single server and by limited performance on the server. There is a risk that single metadata server performance will not be adequate and may adversely affect both applications and interactive users. This risk has already occurred and will continue to affect performance.

Trigger: Direct observations reported by users or staff

The OLCF is working with other major Lustre stakeholders through OpenSFS to develop features to improve single metadata server performance and follow-on support of multiple metadata servers for the Lustre file system. The center has deployed Lustre 2.4, which has the Distributed Namespace (DNE) feature meant to alleviate the metadata bottleneck. The center will hold off on turning on the DNE feature until later in 2015 because of schedule delays caused by performance issues in Lustre. After DNE is turned on, the center will be able to determine if it alleviates the metadata performance bottleneck. Additional measures such as multiple file systems have been deployed that reduce the load on the metadata server.

The probability of the original appreciation of the risk was changed from Medium to High and the scope impact was changed from Medium to Low. These changes did not affect the overall original risk rating of Medium. The overall residual risk rating of Low did not change because it is believed that the mitigation efforts will be effective. The risk continues to be monitored.

5.5.2 ID# 997 Problems with System Reliability, Diagnosis and Recovery in a Large Hybrid System May Arise

Risk Owner	James Rogers, NCCS Director of Operations		
Probability	Low		
Impact	***Cost:*** Low	***Schedule:*** Low	***Scope/Tech:*** Low
Rating	LOW		
Status	Mitigating the risk		

NVIDIA will supply the diagnostics for Glacier. There is concern that, because of the complexities involved, these diagnostics may not be developed sufficiently to identify failures of large systems and cannot be integrated into Cray reporting mechanisms.

Trigger: Intelligence on actual or likely problems

A risk event of this type occurred in July 2012. Based on the mitigation activities implemented through 2014, the risk was recategorized from High to Low. The response included failure analysis, significant mechanical rework of the SXM, aggressive electrical and mechanical stress testing of the reworked GPU/SXM design, individual repair of every single GPU/SXM, and multi-phase testing of repaired components before they were returned to service. These activities generated new protocols/diagnostics for identifying GPU/SXM FRU failures.

NVIDIA device driver revisions provide high-fidelity per-GPU statistics; Cray provides a mechanism for aggregating these statistics to understand GPU computational and memory utilization.

Additionally, OLCF and its subcontractors have developed and implemented mechanisms that monitor/report both correctable and uncorrectable errors on CPUs, CPU main memory (DIMMS), GPUs, and GPU memory. These events are correlated against hardware maintenance activities by Cray, and they can be used to identify failing components (correctable errors above acceptable error thresholds). Statistical analysis reveals that, since the repair in 2012, fewer than 10 physical GPU/SXMs have accumulated more than 90% of the known errors system wide. These results demonstrate the significant stability of the system through CY 2014.

The ability to identify and remove failing nodes from service and successfully route the system contributes to high system availability, well above the established metrics for success. The risk continues to be monitored.

5.5.3 ID #1142 OLCF Cost Increases from Fewer Computer Room Customers

Risk Owner	James Rogers, NCCS Director of Operations		
Probability	High		
Impact	***Cost:*** High	***Schedule:*** Low	***Scope/Tech:*** Low
Rating	HIGH		
Status	Mitigating the risk		

The OLCF Building 5600 CEP has been built and maintained to support the total capacity of the 5600 data center. Recently, a large computer in the data center was decommissioned; therefore, operational costs are divided among fewer customers. Fewer customers will likely result in higher cost to the remaining data center customers, especially the OLCF, which is the largest computer center customer.

Trigger: Loss of major data center customers

ORNL will look to reduce maintenance costs by reducing the equipment that needs to be maintained. It is also possible that this cost could be shared with ORNL Facilities and Operations, as it might be willing to contribute to the maintenance costs if doing so is in the best interest of the laboratory.

5.6 MAJOR RISKS FOR NEXT YEAR

The major risks for next year will be similar to the major risks tracked this year. However, as Titan has now completed acceptance and users are effectively using the system, many of the risks of such a new architecture are lower than the level at which the center had rated them last year.

6. SITE OFFICE SAFETY METRICS

CHARGE QUESTION 6: Has the facility incorporated site office safety recommendations appropriately?
OLCF RESPONSE: Yes.

6.1 SUMMARY

ORNL is committed to operating under the DOE safety regulations specified in 10 C.F R. 851,[20] "Worker Safety and Health Program" as well as applicable DOE Orders and Standards. These safety requirements are incorporated into the ORNL contract, as required compliance documents. To implement these safety requirements in a consistent manner across ORNL, UT-Battelle (UT-B) 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 within ORNL operations and R&D to implement integrated safety management (ISM) functions and principles. The use of ORNL's ISM process culminates with the development and implementation of research safety summaries (RSSs), which are reviewed and approved by the ORNL Safety Services Division, line managers, and the research staff.

An RSS provides the means by which ORNL management and staff can plan and conduct research in a safe manner. It is used to control work, train participants, and provide information regarding operations and emergency services if ever needed. Under a work control review system, work plans are also written before maintenance work is allowed to proceed, to ensure that the 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 on RSSs, work plans, and subcontracts, as well as inspections of job sites throughout each year. Lessons learned, safety snapshots, and management assessments are conducted and recorded into the Assessment and Commitment Tracking System (ACTS). ACTS provides feedback for the completion of the ORNL ISM process. The DOE ORNL Site Office (OSO) participates in the field implementation and documentation of all of the operational safety reviews, and partners with the ORNL Offices of Institutional Planning and Integrated Performance Management and Safety Service Division on some independent safety management system assessments.

The culture of safety at ORNL is reflected in the above processes, which seek to reduce and prevent injuries to our personnel and their potential exposure to hazards associated with the operation of the facility. The OLCF works closely with the OSO and Regina Chung, the Federal Project Director, who

---

[20] 10 C.F.R. 851 outlines the requirements for a worker safety/health program to ensure that DOE contractors and their workers operate a safe workplace. Additionally, 10 C.F.R. 851 establishes procedures for investigating whether a violation of a requirement of this part has occurred, for determining the nature and extent of any such violation, and for imposing an appropriate remedy.

solicited the following feedback from the OSO staff in the Operations and Oversight Division regarding OLCF's safety culture.

- DOE had established a target Total Recordable Cases (TRC) rate and Days Away, Restricted, or Transferred (DART) rate for UT-B's operation and management of ORNL in the FY 2014 Performance Evaluation and Measurement Plan (PEMP). A review of the monthly Safety Charts and the TRC and DART summary documents submitted to OSO indicated that overall FY 2014 TRC and DART rates have increased from FY 2013. However, operations of OLCF in the National Center for Computational Sciences remained safe, efficient, and effective. It has resulted in zero TRCs and zero DARTs in FY 2014.
- A notable outcome to "Develop and execute an implementation plan that is responsive to the recommendations from the UT-B Electrical Safety Task Force" was established by the DOE/OSO in the 2014 PEMP. This outcome was achieved by UT-B as reported in the FY 2014 Performance Evaluation Report. Significant progress was made during FY 2014 in completing the actions, including revision to the Electrical Safety SBMS subject area and revising Work/Project Planning and Control Subject Area implementation tools.
- A review of the OLCF-4 project Health and Safety Plan and the Hazard Analysis Report in preparation for the DOE Office of Science/Office of Project Assessment's CD-2/3b Review was conducted in August 2014. Comments from the DOE Site Office Subject Matter Experts were fully integrated. The OSO continued to work closely with UT-B OLCF operations staff to ensure that any operational issues were addressed in a timely manner.

Cyber Security

HIGH PERFORMANCE COMPUTING FACILITY
2014 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2015

7. CYBER SECURITY

CHARGE QUESTION 7: Has the site been certified to operate (cyber security)?

OLCF RESPONSE: Yes, the most recent OLCF authority to operate (ATO) was granted on June 16, 2014. The current ATO expires on June 16, 2015.

7.1 SUMMARY

All information technology (IT) systems operating for the federal government must have certification and accreditation (C&A) to operate. This involves the development of a policy, the approval of the policy, and an assessment of how well the organization is managing those IT resources—an assessment to determine that the policy is being put into practice.

The OLCF has the authority to operate for 1 year under the ORNL C&A package approved by DOE on June 16, 2014. The ORNL C&A package uses *Recommended Security Controls for Federal Information Systems and Organizations* [National Institute of Standards and Technology Special Publication 800-53, revision 3 (2009)] as a guideline for security controls. The OLCF is accredited at the moderate level o controls for protecting the integrity of the user and system information (FIPS Publication 199), which authorizes the facility to process sensitive, proprietary, and export-controlled data.

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

Summary of the Proposed Metric Values

HIGH PERFORMANCE COMPUTING FACILITY
2014 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2015

8. ACTUAL AND PROPOSED METRIC VALUES

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

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

8.1 SUMMARY

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

Table 8.1. OLCF metrics and actual data for 2014 and proposed metrics and targets for 2015 and 2016

2014 metric and target	2014 actual	2015 metric	2015 target	2016 target	Reporting period
Are the processes for supporting the customers, resolving problems, and outreach effective?					
Customer Metric 1: Customer Satisfaction					
Overall score on the OLCF user survey. Target: Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample	The OLCF exceeded the metric target: 4.6/5.0.	Overall score on the OLCF user survey	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample	Annual
Improvement on results that scored below satisfactory in the previous period. Target: Results will show improvement in at least ½ of questions that scored below satisfactory (3.5) in the previous period.	The OLCF exceeded the metric target: No question scored below satisfactory (3.5/5.0) on the 2014 survey.	Improvement on results that scored below satisfactory in the previous period.	Results will show improvement in at least one-half of the questions that scored below satisfactory (3.5) in the previous period.	Results will show improvement in at least one-half of the questions that scored below satisfactory (3.5) in the previous period.	Annual
Customer Metric 2: Problem Resolution					
OLCF survey results related to problem resolution. Target: Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample	The OLCF exceeded the metric target: 4.6/5.0.	OLCF survey results related to problem resolution	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample	Annual
OLCF user problem resolution time period. Target: 80% of OLCF user problems will be addressed within three business days, by either resolving the problem or informing the user how the problem will be resolved.	The OLCF exceeded the metric target: 90%.	OLCF user problem resolution time period.	Eighty percent of OLCF user problems will be addressed within three business days, by either resolving the problem or informing the user how the problem will be resolved.	Eighty percent of OLCF user problems will be addressed within three business days, by either resolving the problem or informing the user how the problem will be resolved.	Monthly

Table 8.1. (continued)

2014 metric and target	2014 actual	2015 metric	2015 target	2016 target	Reporting period
Customer Metric 3: User Support					
Average of user support ratings. Target: Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample	The OLCF exceeded the metric target: 4.6/5.0	Average of user support ratings	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample	Annual
Is the facility maximizing the use of its HPC systems and other resources consistent with its mission?					
Business Metric 1: System Availability (for a period of one year following a major system upgrade, the targeted scheduled availability is 85% and overall availability is 80%)[a]					
Scheduled Availability. Target: Titan: 90%; HPSS: 95%; External File Systems: 90/95%.	The OLCF exceeded the metric target. Titan: 99.59%; HPSS: 99.83%; /atlas1: 99.73%; /atlas2: 99.50%	Scheduled availability.	Titan: 90%; HPSS 95%; External File Systems 95%	Titan: 90%; HPSS 95%; External File Systems 95%	Monthly
Overall Availability. Target: Titan: 80%/85%/90%; HPSS 90%; External File Systems 85/90%.	The OLCF exceeded the metric target: Titan: 95.8%; HPSS: 98.56%; /atlas1: 99.15%; /atlas2: 98.59%.	Overall availability.	Titan: 90%; HPSS 90%; External File Systems: existing, 90%	Titan: 90%; HPSS 90%; External File Systems: existing, 90%	Monthly
Business Metric 2: Capability Usage					
OLCF will report on capability usage. Target: In the first year of production, at least 30% of the consumed node-hours will be from jobs requesting 20% or more of the available compute nodes	The capability usage was 62.58%. The OLCF exceeded the metric target	OLCF will report on capability usage	In subsequent years, at least 35% of the consumed node-hours will be from jobs requesting 20% or more of the available compute nodes	In subsequent years, at least 35% of the consumed node-hours will be from jobs requesting 20% or more of the available compute nodes	Monthly
N/A[b]	N/A[b]	OLCF will report GPU usage (reference only, no target)	N/A[b]	N/A[b]	Monthly

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

APPENDIX A. RESPONSES TO RECOMMENDATIONS FROM THE 2013 OPERATIONAL ASSESSMENT REVIEW

In February 2014 the OLCF presented the 2013 operational activities of the center to the DOE sponsor. Recommendations provided by reviewers, ORNL actions, and DOE ASCR comments and actions are given in the tables below.

1. **Are the processes for supporting the customers, resolving problems, and outreach effective?**
   Assessment: Yes & Yes

Recommendation	ORNL Action/Comments	HQ Action/Comments
None		

2. **Is the OLCF maximizing the use of its resources consistent with its mission? (Financial data will be covered under this question only for onsite reviews)**
   Assessment: Yes & Yes

Recommendation	ORNL Action/Comments	HQ Action/Comments
OLCF should track and report the use of the GPUs in 2014. With the changes to the Cray software stack in CLE 4.2 UP02, and its installation on Titan in Q1CY14, along with modifications to the job-accounting database, such GPU accounting data will be able to be measured.	We concur. Additional functionality within the resource allocation tracking system is being developed internally that will provide mechanisms for generating reports such as the ones suggested. The target for fully completing these programming/system upgrades is end of CY14. A subset of this functionality is available for the CY14 OA reporting period, and will continue to be further incorporated in to regular operational assessment and reporting activities.	

3. **Is the OLCF enabling scientific achievements consistent with the Department of Energy strategic goals? Specifically applicable to Goal 2: "Maintain a vibrant U.S. effort in science and engineering as a cornerstone of our economic prosperity with clear leadership in strategic areas." Goal 2 includes the targeted outcome: "Continue to develop and deploy high–performance computing hardware and software systems through exascale platforms." Sites may also include contributions to other goals and other targeted outcomes.**
   Assessment: Yes & Yes

Recommendation	ORNL Action/Comments	HQ Action/Comments
None		

4. **Have innovations been implemented that have improved OLCF operations? This includes innovations adopted from, recommended to, or adopted by other Facilities.**
   Assessment: Yes & Yes

Recommendation	ORNL Action/Comments	HQ Action/Comments
None		

5. **Is OLCF effectively managing risk?**
Assessment: Yes & Yes

Recommendation	ORNL Action/Comments	HQ Action/Comments
None		

6. **Has the OLCF incorporated site office safety recommendations appropriately?**
Assessment: Yes & Yes

Recommendation	ORNL Action/Comments	HQ Action/Comments
None		

7. **Does the OLCF have a valid cyber security plan and authority to operate?**
Assessment: Yes & Yes

Recommendation	ORNL Action/Comments	HQ Action/Comments
None		

8. **Are the performance metrics used for the review year and proposed future years sufficient and reasonable for assessing OLCF's Operational performance?**
Assessment: Yes & Yes

Recommendation	ORNL Action/Comments	HQ Action/Comments
The overall availability target for resources past their first year should be 90% to be consistent with other facilities.	The OA target for the OLCF system is set with the sponsor, and may be different from another site based on unique factors to that system. Beginning with the 2014 operational assessment period, the OLCF and their sponsor have agreed to an OA metric that sets a first year target for computing resources of 80%. This increases to 85% in year two and to 90% in year three through the end of life for that particular compute resource. This is reflected in the business result tables as 80%/85%/90% for compute resources. HPSS OA remains 90%. External file systems OA remains 85% in year 1 and 90% in year two through the end of life for that resource 85%/90%.	
The scheduled availability target for resources past their first year should be 95% to be consistent with the other facilities.	The SA target for the OLCF system is set with the sponsor at 90% for HPC resources and external file systems. HPSS SA is 95%.	

What is your overall assessment of the Facility Operational performance?

Recommendation	ORNL Action/Comments	HQ Action/Comments
The OLCF did an exceptional job this year – deploying Titan while still managing to deliver value to its users. They did a great job addressing issues with the machine. OLCF continues to exhibit effective management strategies for providing scientific capabilities to a varied and demanding user community. This report provides a sharp and detailed description of its planning process, risk management strategies, and overall operations. We expect that OLCF will continue to have direct and significant impact on innovation and driving discoveries.	Thank you.	
Excellent. In addition to successfully rolling out a new HPC system, the OLCF facility continues to innovate and provide an effective platform for science done on the cutting edge of HPC. The OLCF has done a great job in operating the facility in 2013. In spite of major issues with the main computational resource (Titan), the facility managed to reach the availability and utilization metrics. The center has also done an excellent job in getting users prepared to use the novel architecture of the system at scale and provides very good response to user issues and problems. The facility has also done a good job in outreach and industrial partnerships. The innovative work done at the center over the 2013 period was impressive and will not only assist facility staff with managing the systems effectively, but also improve the user experience overall.	Thank you.	

9. What is your overall assessment of the Facility Operational performance?

Recommendation	ORNL Action/Comments	HQ Action/Comments
OLCF continues to execute in an exemplary fashion. Despite significant challenges in deploying Titan, OLCF has been successful in delivering substantial computing resources to their user community, enabling science accomplishments, advancing the frontiers of application performance for GPU computing, and providing other innovations that will be broadly useful within the HPC community.	Thank you.	
The OLCF is well regarded for its effectiveness in managing its user facility. OLCF enables scientific capabilities and advancements that do not exist in the private sector or in academia, and should be commended for the excellent support it provides to its user community. The report provides a thorough description of its planning, risk management strategies, and overall operations. We expect that OLCF will continue to have direct and significant impact on innovation and driving discoveries.	Thank you.	

APPENDIX B. TRAINING, WORKSHOPS, AND SEMINARS

Event Type	Description	Date	Participants
User Con Call	OLCF User Conference Call	January 7, 2014	22
Seminar Series	Gleipnir: A tool for memory access pattern profiling, overview, and experience	January 21, 2014	15
Workshop/training	Compiler Optimization Webinar	January 23, 2014	14
Workshop/training	Austin Peay Mathematics Programming (Non Honors)	January 23, 2014	30
Workshop/training	Austin Peay Mathematics Programming (Honors)	January 23, 2014	9
Seminar Series	Advancing Models for Multiphase Flow and Transport in Porous Medium Systems	January 27, 2014	18
Workshop/training	Getting Started at OLCF	January 28, 2014	26
User Con Call	OLCF User Conference Call	February 4, 2014	19
Workshop/training	ADIOS Staff Training	February 11, 2014	16
Seminar Series	Data Science Seminar CSE Division	February 14, 2014	-
Seminar Series	Brown Bag Series on Data Visualization and Analysis	February 17, 2014	-
Workshop/training	ADIOS Staff Training	February 21, 2014	7
User Con Call	OLCF User Conference Call	March 4, 2014	14
Seminar Series	STEM School Chattanooga	March 25, 2014	24
Seminar Series	Tennessee Society of Professional Engineers - Knoxville Chapter	March 28, 2014	20
User Con Call	OLCF User Conference Call	April 1, 2014	16
Workshop/training	Intro to HPC	April 8, 2014	16
Workshop/training	INCITE Proposal Writing Webinar	April 22, 2014	39
Workshop/training	Compiler Directives Weekly Lunch Webinars: Week 1	April 28, 2014	20
Workshop/training	Compiler Directives Weekly Lunch Webinars: Week 2	May 5, 2014	18
User Con Call	OLCF User Conference Call	May 6, 2014	25
Workshop/training	Compiler Directives Weekly Lunch Webinars: PGI	May 9, 2014	26
Workshop/training	Intel Compiler Training	May 21, 2014	30
Workshop/training	Paraview Workshop	May 22, 2014	36
User Con Call	OLCF User Conference Call	June 3, 2014	25
Workshop/training	Crash Course in Supercomputers - Summer 2014	June 20, 2014	120
Workshop/training	HPC Fundamental Classes - Week 1	June 23, 2014	48
Workshop/training	HPC Fundamental Classes - Week 2	June 30, 2014	42
Workshop/training	HPC Fundamental Classes: Week 3	July 7, 2014	38
Workshop/training	HPC Fundamentals Class - Week 4	July 21, 2014	24
Workshop/training	Getting Started at OLCF	July 21, 2014	17
Workshop/training	HPC Fundamental Classes -Week 5	August 4, 2014	14
User Con Call	OLCF User Conference Cal	August 5, 2014	15
Seminar Series	The ADF (Amsterdam Density Functional) Modeling Suite 2014	August 14, 2014	-
Seminar Series	GPU-accelerated Quantum ESPRESSO: past, present and future	August 29, 2014	-

Event Type	Description	Date	Participants
User Con Call	OLCF User Conference Cal	September 9, 2014	31
User Con Call	OLCF User Conference Call	October 15, 2014	30
User Con Call	OLCF User Conference Call	December 3, 2014	85
Workshop/training	CAAR: Call for Proposals Webinar	December 10, 2014	31
Workshop/training	TAU Training at ORNL	April 15-16, 2014	33
Workshop/training	Allinea DDT Training	April 28-29, 2014	20
Workshop/training	2014 OLCF Users Meeting	July 22-24, 2014	111
Workshop/training	CASL Student Workshop	June 12-13, 2014	44
Workshop/training	Joint Facilities User Forum on Data-Intensive Computing	June 16-18, 2014	127
Workshop/training	OLCF Hackathon: OpenACC	October 27-31, 2014	45

APPENDIX C. OUTREACH PRODUCTS

Date	Type of Product	Title
1/2/2014	Rebranding	Created appropriate OLCF templates to comply with branding effort
2/5/2014	Publication	ACCEL Brochure
2/14/2014	Highlight	Bleeding Edge of Fusion
2/14/2014	Highlight	OLCF groups work together to evaluate faster data transfer methods
2/14/2014	Highlight	Collaboration Results in Better Parallel File System Tools
2/14/2014	PPT Slide	OLCF groups work together to evaluate faster data transfer methods
2/14/2014	PPT Slide	Collaboration Results in Better Parallel File System Tools
2/14/2014	Quad Chart	Bleeding Edge of Fusion
2/17/2014	Poster	OLCF Poster #1 - DAC
2/17/2014	Poster	OLCF Poster #2 - DAC
2/17/2014	Poster	OLCF Poster #3 - DAC
2/17/2014	Poster	OLCF Poster #4 - DAC
3/18/2014	Highlight	Titan Project Explores the Smallest Building Blocks of Matter
3/18/2014	Highlight	Titan Delivered in 2013
3/18/2014	Highlight	Seeing is Believing: New OLCF visualization lab showing early promise
3/18/2014	Highlight	OLCF Industry User Named Person to Watch in High-Performance Computing for 2014
3/18/2014	PPT Slide	Titan Delivered in 2013
3/18/2014	PPT Slide	Seeing is Believing: New OLCF visualization lab showing early promise
3/18/2014	PPT Slide	OLCF Industry User Named Person to Watch in High-Performance Computing for 2014
3/18/2014	Quad Chart	Titan Project Explores the Smallest Building Blocks of Matter
3/20/2014	Web Article	Remembering OLCF Group Leader
4/11/2014	Highlight	Simulation Solves Mystery of How Liquid-Crystal Thin Films Disintegrate
4/11/2014	Highlight	ORNL Researchers Collaborate to Study Application I/O Behavior
4/11/2014	Highlight	OLCF Meetings with Indian Officials Produce Opportunities for Collaboration
4/11/2014	PPT Slide	ORNL Researchers Collaborate to Study Application I/O Behavior
4/11/2014	PPT Slide	OLCF Meetings with Indian Officials Produce Opportunities for Collaboration
4/11/2014	Quad Chart	Simulation Solves Mystery of How Liquid-Crystal Thin Films Disintegrate
5/6/2014	Highlight	'Engine of Explosion' Discovered at OLCF now Observed in Nearby Supernova Remnant
5/6/2014	Highlight	World's Most Powerful Accelerator Comes to Titan with a High-Tech Scheduler
5/6/2014	Highlight	Everybody's Talking About Titan
5/6/2014	PPT Slide	World's Most Powerful Accelerator Comes to Titan with a High-Tech Scheduler
5/6/2014	PPT Slide	Everybody's Talking About Titan
5/6/2014	Quad Chart	'Engine of Explosion' Discovered at OLCF now Observed in Nearby Supernova Remnant
6/2/2014	Highlight	Going Nuclear: Titan enables next-gen models for existing, tomorrow's fission reactors
6/2/2014	Highlight	Titan's Tiny Counterpart Engages, Educates
6/2/2014	Highlight	Titan Stability Helps Improve User Experience
6/2/2014	Highlight	Say Hello to Bio: OLCF attends Bio-IT in effort to recruit underrepresented field

Date	Type of Product	Title
6/2/2014	PPT Slide	Going Nuclear: Titan enables next-gen models for existing, tomorrow's fission reactors
6/2/2014	PPT Slide	Titan's Tiny Counterpart Engages, Educates
6/2/2014	PPT Slide	Titan Stability Helps Improve User Experience
6/2/2014	PPT Slide	Say Hello to Bio: OLCF attends Bio-IT in effort to recruit underrepresented field
6/4/2014	Poster	Update OLCF Overview Poster for CSGF Fellows Meeting
6/4/2014	Publication	OLCF Titan Fact Sheet for NUFO
6/4/2014	Video	OLCF Video for NUFO
6/13/2014	Highlight	Ramgen Takes Turbomachine Designs for a Supersonic Spin on Titan
6/13/2014	Quad Chart	Ramgen Takes Turbomachine Designs for a Supersonic Spin on Titan
6/30/2014	Highlight	OLCF Displays How Science Works for America at DC Event
6/30/2014	Highlight	Wells meets AAS: OLCF Director of Science details impact, promise of HPC on radio astronomy
6/30/2014	PPT Slide	OLCF Displays How Science Works for America at DC Event
6/30/2014	PPT Slide	Wells meets AAS: OLCF Director of Science details impact, promise of HPC on radio astronomy
7/2/2014	Web Article	Westinghouse–CASL team wins major computing award for reactor core simulations on Titan
7/23/2014	Highlight	Fine-Tuning the Accelerated Future of Climate Modeling
7/23/2014	Highlight	OLCF Explores Future of Fossil Fuels
7/23/2014	Highlight	OLCF Staff Members Take Leading Roles at Joint Facilities User Forum
7/23/2014	Highlight	UT students experience supercomputing on Titan
7/23/2014	PPT Slide	OLCF Explores Future of Fossil Fuels
7/23/2014	PPT Slide	OLCF Staff Members Take Leading Roles at Joint Facilities User Forum
7/23/2014	PPT Slide	UT students experience supercomputing on Titan
7/23/2014	Quad Chart	Fine-Tuning the Accelerated Future of Climate Modeling
7/29/2014	Poster	SciDAC Poster #1
7/29/2014	Poster	SciDAC Poster #2
8/20/2014	Highlight	Reconnecting the Dots: Researchers use Titan to unravel the many mysteries of high-density plasmas
8/20/2014	Highlight	Titan proves to be more energy-efficient than its predecessor
8/20/2014	Highlight	OLCF Connects with CSGF Fellows at Annual Program Review
8/20/2014	Highlight	Changes to This Year's User Meeting Brought Major Success
8/20/2014	PPT Slide	Titan proves to be more energy-efficient than its predecessor
8/20/2014	PPT Slide	OLCF Connects with CSGF Fellows at Annual Program Review
8/20/2014	PPT Slide	Changes to This Year's User Meeting Brought Major Success
8/20/2014	Quad Chart	Reconnecting the Dots: Researchers use Titan to unravel the many mysteries of high-density plasmas
8/25/2014	Highlight	Oak Ridge Supercomputer Turns the Tide for Consumer Products Research
8/25/2014	Quad Chart	Oak Ridge Supercomputer Turns the Tide for Consumer Products Research
8/30/2014	PPT Slide	National Lab Day Placards
9/8/2014	Publication	Titan Fact Sheet for National Lab Day
9/15/2014	Highlight	OLCF Researcher to Work with Clean Combustion Center at Saudi University
9/15/2014	Highlight	OLCF User Group and Executive Board Established

Date	Type of Product	Title
9/15/2014	Highlight	HPC Fundamentals Course Offers ORNL Staff an Introduction to Supercomputing
9/15/2014	PPT Slide	OLCF Researcher to Work with Clean Combustion Center at Saudi University
9/15/2014	PPT Slide	OLCF User Group and Executive Board Established
9/15/2014	PPT Slide	HPC Fundamentals Course Offers ORNL Staff an Introduction to Supercomputing
10/14/2014	Highlight	Unravelling Enzymes
10/14/2014	Highlight	Tiny Titan Suits Up for Science on the Hill
10/14/2014	Highlight	Supernova Summer School Along the Road to Exascale
10/14/2014	Highlight	Interns Gain Skills, Experience from Summer at ORNL
10/14/2014	PPT Slide	Tiny Titan Suits Up for Science on the Hill
10/14/2014	PPT Slide	Supernova Summer School Along the Road to Exascale
10/14/2014	PPT Slide	Interns Gain Skills, Experience from Summer at ORNL
10/14/2014	Quad Chart	Unravelling Enzymes
10/30/2014	Web Article	Iron-Based Superconductor
11/11/2014	Highlight	Start-up Firm Taps Titan to Model Flooding Risks Worldwide
11/11/2014	Highlight	The Complexities of Combustion
11/11/2014	Highlight	Investigating the Earth's Inner Workings
11/11/2014	Highlight	New Tool to Help Verify Integrity of Archival Data Holdings
11/11/2014	Highlight	Navigating in a Virtual World
11/11/2014	Highlight	OLCF's SciComp Team provides valuable assist for Princeton researcher
11/11/2014	PPT Slide	New Tool to Help Verify Integrity of Archival Data Holdings
11/11/2014	PPT Slide	Navigating in a Virtual World
11/11/2014	PPT Slide	OLCF's SciComp Team provides valuable assist for Princeton researcher
11/11/2014	Quad Chart	Start-up Firm Taps Titan to Model Flooding Risks Worldwide
11/11/2014	Quad Chart	The Complexities of Combustion
11/11/2014	Quad Chart	Investigating the Earth's Inner Workings
11/14/2014	Fact Sheets	INCITE Fact Sheets
11/14/2014	Highlight	Procter & Gamble and Temple University scientists model skin's makeup
11/14/2014	Highlight	Spiraling Back in Time
11/14/2014	News Release	Oak Ridge to Acquire Next Generation Supercomputer
11/14/2014	Poster	SC14 Poster #1
11/14/2014	Poster	SC14 Poster #2
11/14/2014	Poster	SC14 Poster #3
11/14/2014	Poster	SC14 Poster #4
11/14/2014	Quad Chart	Procter & Gamble and Temple University scientists model skin's makeup
11/14/2014	Quad Chart	Spiraling Back in Time
11/14/2014	Video	OLCF Video for Summit Release
11/14/2014	Video	OLCF Video for Summit Release (for NVIDIA version)
11/14/2014	Web/Soc. Media	OLCF SC Blog/Social Media Presence
11/14/2014	Website	Summit: Scale new heights. Discover new solutions.
11/14/2014	Website	SC14 Website
11/18/2014	Web Article	Two OLCF Partners Win Major HPC Award for Innovation Excellence
11/18/2014	Web/Soc. Media	OLCF has Big Presence in the Big Easy at SC14

Date	Type of Product	Title
11/25/2014	Poster	INCITE 2015 Projects Poster
12/8/2014	Highlight	Researchers Call on Titan to Help Solve the Shaky Future of Rare Earth Magnets
12/8/2014	Highlight	OLCF Staff Streamlines Application Development and Deployment
12/8/2014	Highlight	OLCF Hosts First OpenACC Hackathon
12/8/2014	Highlight	OLCF has Big Presence in the Big Easy at SC14
12/8/2014	Highlight	Staff Awarded for Ground-Breaking Research, Outreach, and Cost Savings
12/8/2014	PPT Slide	OLCF Staff Streamlines Application Development and Deployment
12/8/2014	PPT Slide	OLCF Hosts First OpenACC Hackathon
12/8/2014	PPT Slide	OLCF has Big Presence in the Big Easy at SC14
12/8/2014	PPT Slide	Staff Awarded for Ground-Breaking Research, Outreach, and Cost Savings
12/8/2014	Quad Chart	Researchers Call on Titan to Help Solve the Shaky Future of Rare Earth Magnets
12/15/2014	Publication	Oak Ridge Leadership Computing Facility Annual Report 2013-14
12/16/2014	Poster	ALCC Projects Poster 2013-14

APPENDIX D. BUSINESS RESULTS FORMULAS

Scheduled Availability

2014 Operational Assessment Guidance

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

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

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

Overall Availability

2014 Operational Assessment Guidance

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

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

Mean Time to Interrupt

2014 Operational Assessment Guidance

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

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

where

time in period is start time–end time,

start time = end of last outage prior to reporting period,

end time = start of first outage after reporting period (if available) or start of the last outage in the reporting period.

Mean Time to Failure

2014 Operational Assessment Guidance

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

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

where

time in period is start time–end time,

start time = end of last outage prior to reporting period,

end time = start of first outage after reporting period (if available) or start of the last outage in the reporting period.

System Utilization

2014 Operational Assessment Guidance

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

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

APPENDIX E. DIRECTOR'S DISCRETIONARY LIST

PI	Institution	Allocation	Utilization	Project Name
Pratul Agarwal	ORNL	8,000,000	17,983,241	Characterizing the Conformational Sub-states for Developing Hyper-catalytic Enzymes
Kiran Alapaty	Environmental Protection Agency	1,800,000	0	MPAS-US EPA
Edoardo Apra	PNNL	1,000,000	148,708	NWChem Porting to Titan
Moetasim Ashfaq	ORNL	13,600,000	4,087,766	A hierarchical regional modeling framework for decadal-scale hydro-climatic predictions and impact assessments
Moetasim Ashfaq	ORNL	3,000,000	3,089,905	Towards the Development of an Integrated Energy-Water Risk Assessment Tool for Probable Maximum Precipitation and Flood
Dominique Aubert	University Strasbourg	2,000,000	518,006	BEMMA : Benchmarking Emma
Nitin Bangera	MIND Research Network	50,000	0	GPU Accelerated Forward Solutions for the EEG and MEG
Wolfgang Bangerth	Texas A and M University	2,000,000	1,156	deal.II/Aspect
Charlotte Barbier	ORNL	1,000,000	115,198	Large Scale Hydraulic Fracture Simulations
Rodney Bartlett	ORNL	5,000,000	5,157,593	Coupled-cluster Studies of the Active Site of the Cytochrome P450 Enzyme
Jerome Baudry	UT-Knoxville	17,000,000	15,365,869	Massive ensemble docking for drug toxicity prediction
Jerome Baudry	UT-Knoxville	10,000,000	23,086	Drugging the " Undruggable"
David Beratan	Duke University	2,600,000	781,522	Photoinduced Electron Transfer Between Semiconducting Nanoparticles
Amitava Bhattacharjee	University of Iowa	8,000,000	20,099,104	Magnetic Reconnection and Laboratory Astrophysics with Laser-Produced Plasmas
Simon Billinge	BNL	1,000,000	0	Nanostructure Complex Modeling
William Bird	TVA	1,000,000	91,178	TVA CASL Test Stand - Evaluation of Lower Plenum Flow Anomaly using VERA/Hydra-TH
Michael Brim	ORNL	1,500,000	12,934	Lustre-Vision
Aurel Bulgac	University of Washington	15,000,000	188,053	Fission Fragment Yields within Time-Dependent Density Functional Theory
Michael Bussmann	Helmholtz-Zentrum Dresden-Rossendorf	750,000	944,736	Laser-Wakefield Simulations Using PICONGPU

PI	Institution	Allocation	Utilization	Project Name
John Bussoletti	Boeing	3,000,000	0	Exploration Of New Computing Technology For Modeling High Lift Systems Of Commercial Aircraft
Henri Calandra	Total	5,000,000	0	Advance computing for Geoscience applications
Luigi Capone	Texas A and M University	1,000,000	0	HYDRA GPUs architecture migration task I
Richard Casey	Colorado State University	3,750,000	8,207,320	Large-Scale Metagenomic and Bioinformatic Data Analysis of Semiconductor-based Next Generation DNA Sequencing
Barbara Chapman	University of Houston	219,449	0	A similarity-based analysis tool for pattern derivation and large scale program restructuring
Xiaolin Cheng	ORNL	7,500,000	13,598,732	Computational Study of Cellulose Synthase via Enhanced Sampling in High Performance Computing
Scott Christley	University of Chicago	2,000,000	12	HPC Lung Model of Pulmonary Acinus
Michael Clark	NVidia	2,000,000	22,463	Petascale Cross Correlation
Todd Coleman	University of California San Diego	500,000	0	Efficient Bayesian Inference Methods with Distributed Convex Optimization
Jason Cope	DataDirect Networks	1,000,000	1,876	Assessing the Scalability of DataDirect Networks 'Iron Monkey' Burst Buffer on Titan
Jacques Corbeil	KatRiskm LLC	750,000	31,743	Next Generation De Novo Assembler
Andrew Corrigan	Department of Defense	700,000	450,399	Benchmarking the Jet Engine Noise Reduction (JENRE) code on Titan
Bruce D’Amora	IBM	500,000	8,533	CORAL Benchmarking
Erik Deumens	University of Florida	5,000,000	178,798	Predicting and improving the performance in the super instruction architecture
Russell DeVane	University of Wisconsin	6,000,000	11,409,030	Disruption of lamellar lipid systems induced by small molecule permeants
Nikolay Dokholyan	University of North Carolina Chapel Hill	500,000	159	Characterization of structure and dynamics of clinically relevant membrane proteins by means of molecular dynamics simulations
Dipangkar Dutta	Mississippi State University	1,470,000	277	A New Search for the Neutron Electric Dipole Moment
Hazim El-Mounayri	Indiana University	700,000	250,420	AFM-based nanomachining of 3D structures
Eirik Endeve	ORNL	5,000,000	0	Fast Algorithms for Multiphasic Modeling of Non- Equilibrium Transport
Katherine Evans	ORNL	11,000,000	6,477,780	Transition and Preparation for DOE Accelerated Climate Model for Energy
Mark Fahey	ORNL	1,000	0	Automatic Library Tracking Database

PI	Institution	Allocation	Utilization	Project Name
Fernanda Foertter	ORNL	5,000,000	0	Developing Scalable Heterogeneous Computing Training Code Examples
Benoit Forget	Massachusetts Institute of Technology	6,000,000	3,229,888	OpenMC/TH
Sampath Gajawada	ORNL	57,600	0	TGS Titan Benchmark Runs
Manjunath Gorentla Venkata	ORNL	5,100,000	9,700,410	Achieving Near Real-time Imputation in Thousands of Samples using HPC Systems
Manjunath Gorentla Venkata	ORNL	2,000,000	4,671,946	Enhancing Application Performance and Resiliency of Extreme-Scale Systems
David Green	ORNL	2,000,000	103,847	AORSA-VORPAL Coupling Development for Radio-Frequency Heating of Fusion Plasmas
James Hack	ORNL	23,520,000	1,344	Ultra High Resolution Global Climate Simulation to Explore and Quantify Predictive Skill for Climate Means, Variability and Extremes
Dana Hammond	NASA-LARC	500,000	71	Scaling of FUN3D for 2014 INCITE Proposal
David Hart	National Center for Atmospheric Research	200,000	0	NCAR benchmarking for climate and weather models
Cory Hauck	ORNL	2,000,000	1,708,441	Moment Methods for Linear Kinetic Equations
Remco Havenith	Groningen University	7,000,000	1,636,105	Computational Modeling of Organic Photovoltaics
Feng He	University of Wisconsin	300,000	0	Implications of the Early Anthropogenic Hypothesis
Katrin Heitmann	ANL	2,000,000	18,205	Dark Universe
Oscar Hernandez	ORNL	200,000	1,348,751	SCALPERF
Judith Hill	ORNL	40,000,000	18,247,051	Computational Partnerships
Judith Hill	ORNL	100,000	16,700	Computational Science Graduate Fellowship Program
Forrest Hoffman	ORNL	6,000,000	4,635,853	DOE SciDAC-3 ACES4BGC Partnership Project Gen2ESM Foundry
Travis Humble	ORNL	3,000,000	0	Quantum Computing Simulations
Travis Humble	ORNL	2,000,000	1,053,003	Jade Adiabatic Device Emulator (JADE)
Antti-Pekka Hynninen	NREL	500,000	7,339	High performance GPU molecular dynamics engine in CHARMM
Daniel Jacobson	Stellenbosch University	1,000,000	2,081	Scaling up of Parallelized Ortholog Detection Algorithms for Comparative Genomics of Bacterial Genomes
Jacek Jakowski	UT-Knoxville	100,000	40,387	Electronic structure calculation methods on accelerators

PI	Institution	Allocation	Utilization	Project Name
Balaji Jayaraman	Pennsylvania State University	4,000,000	3,396,331	Towards High-fidelity Petascale Computations of Atmospheric Turbulence-driven Wind Turbine Aerodynamics and Wakes
Thomas Jenkins	Tech-X	6,000,000	208,846	Extended magnetohydrodynamic simulations of toroidal fusion plasmas
Terry Jones	ORNL	3,000,000	0	HPC Colony II
Terry Jones	ORNL	3,000,000	1,916,790	Colony
Terry Jones	ORNL	3,000,000	0	Hobbes: Operating System and Runtime Research for Extreme Scale
Alexander Kamil	Columbia University	2,000,000	3	Identify Organisms from a Stream of DNA Sequences
George Karniadakis	Massachusetts Institute of Technology	1,000,000	523,050	Development of GPU-Accelerated Mesoscale Simulations
David Keffer	UT-Knoxville	1,000,000	173,903	Computational study of novel lignin-derived carbon composite Li-ion anodes
Jeongnim Kim	ORNL	3,000,000	72,782	QMC Glue
Gerhard Klimeck	Purdue University	5,000,000	2	Nanoelectronics Modeling at the Petascale on Heterogeneous Systems
Alexei Kritsuk	University of California	3,000,000	2,104,260	High-resolution Simulations of Compressible MHD turbulence on GPU
Predrag Krstic	ORNL	1,470,000	23	Science of the Plasma-Material Interface at Extreme Conditions
Predrag Krstic	ORNL	1,000,000	351,191	Gas-Liquid-Solid Interfaces for Energy Applications
Tang-Wei Kuo	General Motors	9,000,000	9,905,077	Multi-hole injector optimization for spark-ignited direct-injection gasoline engines
Jakub Kurzak	UT-Knoxville	200,000	623	Bench Testing Environment for Automated Software Tuning
William Laidig	Procter & Gamble	3,000,000	16,394	Can Supercomputing Help Mechanismic Understanding of a Novel Catalyst for a Strategic Raw Material?
Gregory Laskowski	GE Global Research	3,000,000	32,356	Adjoint-based techniques for LES
James Lewis	West Virginia University	1,000,000	889	High-Throughput Design of Delafossite Oxide Materials for Photovoltaics
Xiaoyi Li	United Technologies Research Center	1,000,000	27,764	High fidelity direct numerical simulation of sprays for realistic injection applications in industry
Xiaoye Li	LBNL	20,000,000	10,641,884	Enabling Next-Generation Light Source Data Analysis through Massive Parallelism
Lucas Lindsay	Department of Defense	300,000	91,340	First principles thermal transport and thermoelectric properties of materials

PI	Institution	Allocation	Utilization	Project Name
Ping Liu	BNL	800,000	6	CO2 Activation through Heterogeneous Catalysis
Philip Locascio	University of Oxford	2,750,000	37,914,261	Molecular action of the membrane bound dimer Osmolarity sensor of E.coli, EnvZ
Dag Lohmann	KatRiskm LLC	2,000,000	1,825,374	The Cost of Global Climate Extremes
Dag Lohmann	KatRiskm LLC	5,000,000	3,532,343	Worldwide Flood Maps
Carlos Lopez	Vanderbilt University	2,500,000	396,830	Reducing unidentifiability in cell-signaling network models through multidimensional analysis and molecular simulation
Salil Mahajan	ORNL	2,000,000	1,945,729	Impact of Aerosols and Air-sea Interactions on CESM Biases in the Western Pacific Warm Pool Region
Guillermo Maldonaldo	UT-Knoxville	1,200,000	0	Fuel/Core Heavy Metal Design Optimization to Improve Performance/Safety of LSCR
Tiziana Matteo	Carnegie Mellon University	2,940,000	5,858,562	Petascale Cosmology with P-Gadget
Jens Meiler	Vanderbilt University	10,000,000	11,068,814	Investigation of the Mechanism of Cellulose Synthesis in Plants Using Neutron Scattering and High Performance Computing
Dominic Meiser	Tech-X	3,000,000	1,530,473	TxHPCG
Bronson Messer	ORNL	1,000,000	0	CORAL Benchmarking
Bronson Messer	ORNL	6,000,000	8,011,101	Explosive Nucleosynthesis and Deflagration to Detonation in Type Ia Supernovae
Vittorio Michelassi	GE Global Research	10,000,000	8,921,468	HIPSTAR-G-01
Misun Min	ANL	500,000	450,537	Nek-HOM (Codes for High Order Methods)
Brian Mitchell	GE Global Research	1,500,000	10,770	TACOMA GPU Port
Benson Muite	King Abudallah University of Science and Technology	1,000,000	1,132,829	Numerical investigations of semilinear partial differential equations
Chris Mundy	PNNL	11,760,000	29,525	Control of Complex Transformations with Advanced Molecular Simulation
Marco Nardelli	University of North Texas	700,000	23,840,584	Ab initio infrastructure for high-throughput computational materials
Joshua New	ORNL	175,000	761,252	Autotune E+ Buildings
Joshua New	ORNL	2,000,000	0	Big Data Mining for Building Analytics
Jeffrey Nichols	ORNL	13,000,000	173,046	Holmes - Cryptographic Key Discovery
Frank Noe	Zuse Institute Berlin	6,000,000	78	Adaptive molecular simulation of the immunological synapse
Cevdet Noyan	Columbia University	5,000,000	732,576	Analysis of Powder X-ray Diffraction (PXRD) Profiles of Polycrystalline Nanopowders in Kinematic Regime

PI	Institution	Allocation	Utilization	Project Name
Mark Oxley	Vanderbilt University	5,000,000	360,567	Simulation of atomic-resolution electron energy loss spectra on the meso-scale
Chongle Pan	ORNL	4,500,000	40,139	Large-scale metagenomics analysis for biosurveillance and environmental microbiology
Sergey Panitkin	BNL	10,500,000	2,023,170	Next Generation Workload Management System
Sreekanth Pannala	ORNL	1,000,000	5	Computational Infrastructure for parallel simulations of Cycle-to-Cycle variations of in-cylinder combustion
Sreekanth Pannala	ORNL	1,470,000	625,449	Using Solid Particles as Heat Transfer Fluid in CSP Plants
Robert Patton	ORNL	2,000,000	1,173	Scalable Deep Learning Systems for Exascale Data Analysis
B. Pettitt	University of Houston	100,000	3,037	Solubility of Peptides and Proteins
Michael Pindzola	Auburn University	2,000,000	2,780,741	Atomic and Molecular Collisions for Astrophysical and Laboratory Physics
Roy Primus	GE Global Research	4,500,000	415,654	Application of High Performance Computing for Simulating Cycle to Cycle Variation in Dual Fuel Combustion Engines
David Pugmire	ORNL	2,000,000	440,781	SDAV
Yevgeniy Puzyrev	Vanderbilt University	1,000,000	332,961	Flexural phonons and mechanical properties of two-dimensional materials
Vahid Ranjbar	Brookhaven National Lab	1,000,000	0	Spin Tracking for RHIC
Kenneth Read	ORNL	300,000	32,762	Probing Fluctuating Initial Conditions of Heavy-Ion Collisions
Bhanu Rekepalli	UT-Knoxville	1,000,000	0	Developing highly scalable parallel Bioinformatics applications on Titan
Duane Rosenberg	Oak Ridge National Laboratory	2,000,000	0	Small scale statistics and intermittency in rotating strongly stratified turbulence: Verification and connection to Bolgiano-Obokhov phenomenology
Bhagawan Sahu	Global Foundries US Inc.	5,880,000	14,486	Density Functional Studies of Si/SiGe interface structures
Vaidyanathan Sankaran	United Technologies Research Center	2,000,000	395,948	Towards Combustor Simulation Using Large Eddy Simulation and Graphical Processing Units
Ramanan Sankaran	ORNL	6,000,000	0	Porting of the RAPTOR code to GPU-accelerated nodes and scalability studies for Large Eddy Simulation (LES) of high-pressure liquid hydrocarbon fuel injection processes
Andreas Schaefer	Friedrich-Alexander-Universitaet Erlangen-Nuernberg (FAU)	1,470,000	37,800	LibGeoDecomp

PI	Institution	Allocation	Utilization	Project Name
Gustavo Seabra	Universide de Federal de Pernambuco	561,467	213	Elucidation of the Molecular Mechanism of Enzymatic Reactions by Molecular Dynamics and Hybrid Quantum Mechanical and Molecular Mechanics Simulations
Bradley Settlemyer	ORNL	2,000,000	1,220,888	Towards a Resilient and Scalable Infrastructure for Big Data
Steven Shannon	North Carolina State University	1,470,000	4,185	Particle-In-Cell Simulation of Radio Frequency Field Structure Near Plasma Facing Antenna Components
Yuji Shinano	Zuse Institute Berlin	0	0	ParaSCIP
Yuji Shinano	Zuse Institute Berlin	5,000,000	1,295,254	ParaSCIP
Galen Shipman	ORNL	200,000	4,377	Data Intensive Science Incubators
Galen Shipman	ORNL	0	5	ESG
Galen Shipman	ORNL	8,000,000	14,871,401	Accelerating Materials Modeling with Leadership Computing
Adam Simpson	ORNL	100,000	1,336	Supercomputing in the Classroom
Srdjan Simunovic	ORNL	750,000	0	Validating Predictive Modeling of Carbon Fiber Composites In Automotive Crash Applications
Mohammed Sourouri	ORNL	405,000	26,052	User-Friendly programming of GPU Clusters
Sriraj Srinivasan	Arkema	2,000,000	155,783	High Performance Catalysts in Fluorination of Climate Impacting Fluorogases
Ravichandra Srinivasan	Ramgen Power Systems	2,000,000	964,692	Visualization of Tip Injection Phenomena in the Near Stall Regime of a Transonic Fan Stage
Ravichandra Srinivasan	Ramgen Power Systems	32,000,000	21,570,062	Compressible Flow Turbomachinery Optimization: Numerical Tools Advancement
Ashok Srinivasan	Florida State University	439,459	0	Accelerating Quantum Monte Carlo on Massively Parallel Computing Platforms
Gabriel Staffelbach	CERFACS	4,000,000	712,267	Tolentucoco: Towards LES in Turbines using code compiling
Jirina Stone	University of Oxford	8,000,000	0	Phase transitions in high density matter in neutron stars and supernovae
Bobby Sumpter	ORNL	8,000,000	2,889,238	Understanding and Manipulating Surface Mediated Interactions
Alexander Szalay	Johns Hopkins University	40,000	0	Demonstrations of Data-Scope at 100 Gbps Across a National Data-Intensive Computational Science Testbed at SC13
Jim Tallman	GE Global Research	2,940,000	2,180,271	Tacoma Scalability for INCITE-sized problems
William Tang	Princeton University	8,000,000	884,068	Extreme Scale PIC Research on Advanced Architectures

PI	Institution	Allocation	Utilization	Project Name
Rong Tian	Chinese Academy of Sciences Institute of Computing Technology	1,000,000	177,108	Scalable and fault tolerant meshfree/particle simulation
Georgia Tourassi	ORNL	1,000,000	179,838	High Performance Infodemiology Architecture for Knowledge Extraction from Unstructured Health Text
David Trebotich	LANL	4,000,000	10,879,711	Chombo-Crunch
John Turner	ORNL	6,000,000	102,296	Computational Engineering and Energy Sciences (CEES) Group Projects
Gerald Tuskan	ORNL	500,000	0	Poplar genome wide association study - BioEnergy Science Center
Brad VanDerWege	Ford Motor Company	1,470,000	788,222	Cycle-to-Cycle Combustion Variation Modeling
Joost VandeVondele	ETH Zurich	6,000,000	4,101,873	5th rung functionals with CP2K
Nicola Varini	Curtin University	2,400,000	2,603,725	EXX-PETA
Greg Voth	University of Utah	1,500,000	62,942	Energy Storage and Conversion Materials
Yi Wang	ORNL	1,470,000	863,144	CFD Modeling of Industrial Scale Fire Growth and Suppression
Cai-Zhuang Wang	Ameslab	2,000,000	1,994,073	Structure prediction for energy-related complex materials
Joshua Webb	Caterpillar Inc	10,000,000	0	GPU Enhancement of Weld Distortion Prediction
Tao Wei	Lamar University	1,000,000	21,547	Anti-biofouling Material Design
Theresa Windus	Iowa State University	5,000,000	6,421,846	Critical Materials Institute: Separations Science
Xiaobiao Xu	Florida State University	3,000,000	0	Simulating the circulation of North Atlantic Ocean at 1/50° resolution
Pui-kuen Yeung	Georgia Institute of Technology	7,500,000	1,535,934	Scale-Similarity and Turbulence Mixing: Schmidt number effects and new algorithmic developments
Martin Zacharias	Entrepreneurial University	3,000,000	4,235,571	Molecular dynamics of amyloid formation at atomic resolution
Martin Zacharias	Technical University of Munich	9,000,000	0	Molecular Dynamics of amyloid formation at atomic resolution: Test of simulation and boundary conditions
Cyril Zeller	NVIDIA	13,500,000	13,028,918	CoDesign
Gecheng Zha	University of Miami	600,000	1,750,172	Large Eddy Simulation of Base Drag Reduction Using a novel Passive Jet Flow control method
Xiaoguang Zhang	ORNL	1,700,000	1,091,748	A Comprehensive Theoretical/Numerical Tool for Electron Transport in Mesoscale-Heterostructures
Leonid Zhigilei	University of Virginia	4,000,000	58,377,347	Atomistic simulations of laser interactions with metals
Jack Ziegler	NREL	3,000,000	108	Determining FCC Catalyst Residence Time Distributions in Circulating Biofuel Reactors

PI	Institution	Allocation	Utilization	Project Name
Simon Zwart	Leiden University	4,000,000	3,174,533	The Fine Structure of the Milky Way Galaxy
Scott Klasky	ORNL	5,000,000	5,710,741	Adios Titan Access
Misun Min	ANL	500,322	0	Nek-HOM (Codes for High Order Methods)
Stephen Poole	ORNL	4,000,000	0	Durmstrang