

US Department of Energy, Office of Science High Performance Computing Facility Operational Assessment 2015 Oak Ridge Leadership Computing Facility



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

Ashley D. Barker
David E. Bernholdt
Arthur S. Bland
Jeff D. Gary
James J. Hack
Stephen T. McNally
James H. Rogers
Brian Smith
T. P. Straatsma
Sreenivas R. Sukumar
Kevin G. Thach
Suzy Tichenor
Sudharshan S. Vazhkudai
Jack C. Wells

March 2016

DOCUMENT AVAILABILITY

Reports produced after January 1, 1996, are generally available free via US Department of Energy (DOE) SciTech Connect.

Website <http://www.osti.gov/scitech/>

Reports produced before January 1, 1996, may be purchased by members of the public from the following source:

National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone 703-605-6000 (1-800-553-6847)
TDD 703-487-4639
Fax 703-605-6900
E-mail info@ntis.gov
Website <http://www.ntis.gov/help/ordermethods.aspx>

Reports are available to DOE employees, DOE contractors, Energy Technology Data Exchange representatives, and International Nuclear Information System representatives from the following source:

Office of Scientific and Technical Information
PO Box 62
Oak Ridge, TN 37831
Telephone 865-576-8401
Fax 865-576-5728
E-mail reports@osti.gov
Website <http://www.osti.gov/contact.html>

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

Oak Ridge Leadership Computing Facility

**HIGH PERFORMANCE COMPUTING FACILITY OPERATIONAL
ASSESSMENT 2015
OAK RIDGE LEADERSHIP COMPUTING FACILITY**

Ashley D. Barker
David E. Bernholdt
Arthur S. Bland
Jeff D. Gary
James J. Hack
Stephen T. McNally
James H. Rogers
Brian Smith
T. P. Straatsma
Sreenivas R. Sukumar
Kevin G. Thach
Suzy Tichenor
Sudharshan S. Vazhkudai
Jack C. Wells

Date Published: March 2016

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

CONTENTS

	Page
LIST OF FIGURES	vi
LIST OF TABLES	vii
ABBREVIATED TERMS	viii
EXECUTIVE SUMMARY	xi
1. USER RESULTS	1
1.1 USER RESULTS SUMMARY	1
1.2 USER SUPPORT METRICS	2
1.2.1 Overall Satisfaction Rating for the Facility	2
1.2.2 Average Rating across All User Support Questions	3
1.2.3 Improvement on Past Year Unsatisfactory Ratings	3
1.2.4 Assessing the Effectiveness of the OLCF User Survey	4
1.3 PROBLEM RESOLUTION METRICS	5
1.3.1 Problem Resolution Metric Summary	5
1.4 USER SUPPORT AND OUTREACH	6
1.4.1 User Support	7
1.4.2 User Assistance and Outreach Analysts	9
1.4.3 Scientific Liaisons	9
1.4.4 Data Liaisons	11
1.4.5 Visualization Liaisons	13
1.4.6 OLCF User Group and Executive Board	15
1.4.7 Training, Education, and Workshops	15
1.4.8 Training and Outreach Activities for the Future Workforce	16
1.4.9 Outreach	18
1.5 LOOKING FORWARD	19
1.5.1 Application Portability	19
1.5.2 Application Readiness and Early Science	19
1.5.3 CSEEN Postdoctoral Program	20
2. BUSINESS RESULTS	22
2.1 BUSINESS RESULTS SUMMARY	22
2.2 CRAY XK7 (TITAN) RESOURCE SUMMARY	22
2.3 CRAY XC30 (EOS) RESOURCE SUMMARY	22
2.4 LUSTRE FILE SYSTEMS (SPIDER II) RESOURCE SUMMARY	22
2.5 DATA ANALYSIS AND VISUALIZATION CLUSTER (RHEA) RESOURCE SUMMARY	23
2.6 HIGH PERFORMANCE STORAGE SYSTEM RESOURCE SUMMARY	23
2.7 VISUALIZATION RESOURCE SUMMARY	23
2.8 OLCF COMPUTATIONAL AND DATA RESOURCE SUMMARY	23
2.8.1 OLCF HPC Resource Production Schedule	23
2.8.2 Business Results Snapshot	24
2.9 RESOURCE AVAILABILITY	26
2.9.1 Scheduled Availability	26
2.9.2 Overall Availability	27
2.9.3 Mean Time to Interrupt	27
2.9.4 Mean Time to Failure	28
2.10 RESOURCE UTILIZATION	28
2.10.1 Resource Utilization Snapshot	28

2.10.2	Total System Utilization	29
2.11	CAPABILITY UTILIZATION.....	30
2.12	GPU USAGE	31
2.12.1	Measuring GPU-Enablement and GPU-Activity	32
2.12.2	Case Study: Transitioning from Discretionary Program to INCITE	34
2.13	PARALLEL DU TOOL SUITE	34
2.14	DDNTOOL	35
2.15	BALANCED PLACEMENT I/O.....	35
2.16	PARALLEL DATA TRANSFER.....	36
3.	STRATEGIC RESULTS	37
3.1	SCIENCE OUTPUT	37
3.1.1	OLCF Publications Report.....	38
3.2	SCIENTIFIC ACCOMPLISHMENTS.....	38
3.2.1	Flowing Toward Red Blood Cell Breakthroughs: George Karniadakis, Brown University, INCITE.....	39
3.2.2	Researchers Mine Information from Next-Generation Subsurface Flow Simulations: James McClure, Virginia Polytechnic Institute and State University, INCITE.....	40
3.2.3	Fusion Researchers Use Titan to Burst Helium Bubbles: Brian Wirth, Oak Ridge National Laboratory, ASCR Leadership Computing Challenge	41
3.2.4	Titan Takes on the Big One: Thomas Jordan, Southern California Earthquake Center, INCITE.....	42
3.2.5	Researchers Build Bacteria’s Photosynthetic Engine: Klaus Schulten, University of Illinois at Urbana-Champaign, INCITE	43
3.3	DIRECTOR’S DISCRETIONARY PROGRAM	44
3.3.1	Mechanical Properties of Nanoscopic Lipid Domains, Director’s Discretionary.....	45
3.4	INDUSTRIAL HPC PARTNERSHIPS PROGRAM: ACCEL	46
3.4.1	FM Global: Fighting Fires with Modeling and Simulation	46
3.4.2	GE Successfully Develops New Predictive Simulation Methods for Gas Turbines	47
4.	INNOVATION	49
4.1	DATA ANALYSIS AND WORKFLOWS	49
4.1.1	OLCF—State of the Workflow Update	49
4.1.2	OLCF and CADES collaborate in Health Sciences	50
4.1.3	BigPanDA Workflow Management on Titan	51
4.2	I/O EFFICIENCY	52
4.2.1	Parallel File System Profiling with fprof.....	52
4.2.2	Checkpoint Advisory Tool.....	52
4.2.3	HPSS Operational Investments and Improvements	53
4.2.4	I/O Test Harness	54
4.3	APPLICATION DEVELOPMENT	54
4.3.1	GUIDO: Grand Unified Information Directory for OLCF Operations.....	54
4.3.2	GLEIPNIR: Understanding Application Memory Usage Patterns and Scalability.....	55
4.4	RESOURCE MANAGEMENT	56
4.4.1	Dual-Ended Scheduling	56
5.	RISK MANAGEMENT	58
5.1	RISK MANAGEMENT SUMMARY	58
5.2	MAJOR RISKS TRACKED IN CY 2015	59
5.2.1	ID# 361—Scientists Decline to Port to Heterogeneous Architecture	59
5.2.2	ID# 407—Loss of Key Personnel	59
5.2.3	ID# 721—Lustre Metadata Performance Continues to Impact Applications	60

5.2.4	ID# 1063—Programming Environment Tools May Be Insufficient	60
5.2.5	ID# 917—Robust Support Will Not Be Available to Ensure Portability of Restructured Applications.....	61
5.3	RISKS THAT OCCURRED DURING THE CURRENT YEAR AND THE EFFECTIVENESS OF THEIR MITIGATION	61
5.3.1	ID #408 – System Outages from External Causes	61
5.3.2	ID #1142 OLCF Cost Increases from Fewer Computer Room Customers.....	62
5.4	RISKS RETIRED DURING THE CURRENT YEAR.....	62
5.4.1	ID# 124—Storage System Reliability and Performance Problems	62
5.4.2	ID# 1093—Impact of Load Swings on Power Grid	63
5.4.3	ID# 1143—ORNL-Furnished equipment	63
5.5	NEW OR RECHARACTERIZED RISKS SINCE LAST REVIEW	64
5.5.1	ID# 997—Problems May Arise with System Reliability, Diagnosis, and Recovery in a Large Hybrid System	64
5.5.2	ID# 1006—Inability to Acquire Sufficient Staff	64
5.5.3	New Risk: ID #1145 —Changes from External Project Managers Cause Development Impacts to HPSS.....	65
5.5.4	New Risk: ID #1146 —Suboptimal Resource Utilization in HPSS Causes Lack of Storage and Greater Expenses	65
5.5.5	New Risk: ID #1147 —Data Integrity Issues in Archival Data in HPSS	66
5.6	MAJOR RISKS FOR NEXT YEAR	66
6.	SITE OFFICE SAFETY METRICS.....	67
6.1	SUMMARY.....	67
7.	CYBER SECURITY	69
7.1	SUMMARY	69
7.2	SAFEGUARDS AND SECURITY SURVEY	69
8.	ACTUAL AND PROPOSED METRIC VALUES	70
8.1	SUMMARY	70
APPENDIX A. RESPONSES TO RECOMMENDATIONS FROM THE 2014 OPERATIONAL ASSESSMENT REVIEW		A-1
APPENDIX B. TRAINING, WORKSHOPS, AND SEMINARS		B-1
APPENDIX C. OUTREACH PRODUCTS.....		C-1
APPENDIX D. BUSINESS RESULTS FORMULAS		D-1
APPENDIX E. DIRECTOR’S DISCRETIONARY PROJECTS ENABLED (AT ANY POINT) IN CY 2015.....		E-1

LIST OF FIGURES

Figure	Page
Figure 1.1. Categorization of help desk tickets.....	6
Figure 1.2. The software (subroutine) of CLM.....	7
Figure 1.3. Proton density after laser impact on a spherical solid density target: irradiated by an ultra-short, high intensity laser (not in picture), the intense electro-magnetic field rips electrons apart from their ions and creates a plasma.	11
Figure 1.4. Graphical representation of ACME ESM workflow.	13
Figure 1.5. OLCF visualization on the cover of the May 2015 <i>Physics of Fluids</i>	14
Figure 1.6. Participants from the BigNeuron Hackathon, including scientific, data, and visualization liaison.	14
Figure 1.7. Robert French (left) and Adam Simpson (right) take visitors through the steps of operating Tiny Titan at an exhibit at the American Museum of Science and Energy.	18
Figure 1.8. Robert French (center) discusses Tiny Titan with Science Bowl participants with Carolyn Lauzon (right) and OLCF’s Adam Simpson (left).	18
Figure 2.1. 2015 XK7 resource utilization—Titan core-hours by program.	29
Figure 2.2. Capability usage by job size bins and project type.	31
Figure 2.3. GPU-enabled percentage of compute time for the DD, ALCC, and INCITE user programs.	33
Figure 2.4. GPU-activity for the DD, ALCC, and INCITE user programs.	33
Figure 2.5. Percentage of GPU-enabled compute time by month for TUR105.	34
Figure 2.6. Splunk-provided sample of performance data obtained from DDNTTool.	35
Figure 3.1. Red blood cells (red) and circulating tumor cells (green) traveling through a microfluidic cell sorting device as simulated by uDeviceX.	39
Figure 3.2. Simulations on Titan enable detailed tracking of two fluid phase systems during flow through porous media.	40
Figure 3.3. Bubble pressure as a function of the number of helium atoms added to a growing, sub-surface gas bubble in tungsten, revealing significant pressure drops as the bubble expands through loop punching before eventually bursting and releasing helium.	41
Figure 3.4. The CyberShake seismic hazard map shows the magnitude, or level of shaking, for the Los Angeles region, defined by the amount of change of a surface or structure in a 2 second period, with a 2% probability of increasing within the next 50 years.	42
Figure 3.5. Researchers used experimental data to create a chromatophore model that contained about 16,000 lipids and 101 proteins, including the five major types of proteins that contribute to the clockwork of processes resulting in the conversion of light energy to ATP.	43
Figure 3.6. The dynamics of contrast-enhanced unilamellar lipid vesicles reveal the local mechanical properties of nanoscopic domains.	46
Figure 4.1. File size distribution of the Atlas1 partition of the Spider II file system.	52

LIST OF TABLES

Table	Page
Table 1.1. 2015 user result metrics summary	2
Table 1.2. OLCF user support summary: Metric targets and calendar year results	2
Table 1.3. Satisfaction rates by program type for key indicators.....	3
Table 1.4. User survey participation	4
Table 1.5. Statistical analysis of key results	5
Table 1.6. Problem resolution metric summary	6
Table 1.7. Applications in the Center for Accelerated Application Readiness	20
Table 2.1. OLCF production computer systems, 2015	24
Table 2.2. OLCF HPC system production dates, 2008–present	24
Table 2.3. OLCF business results summary for HPC systems	25
Table 2.4. OLCF business results summary for Eos.....	25
Table 2.5. OLCF business results summary for HPSS	25
Table 2.6. OLCF business results summary for the external Lustre file system—Spider II.....	26
Table 2.7. OLCF business results summary: Scheduled availability	27
Table 2.8. OLCF business results summary: Overall availability.....	27
Table 2.9. OLCF business results summary: Mean time to interrupt	28
Table 2.10. OLCF business results summary: Mean time to failure.....	28
Table 2.11. The 2015 allocated program performance on the OLCF resources	30
Table 2.12. OLCF capability usage on the Cray XK7 system	30
Table 2.13. 2015 GPU-enabled and CPU-only usage by program	33
Table 3.1. Summary of unique OLCF publications over the period 2012–2015	38
Table 8.1. OLCF metrics and actual data for 2015 and proposed metrics and targets for 2016 and 2017	71

ABBREVIATED TERMS

3-D	three-dimensional
ACCEL	Accelerating Competitiveness through Computational Excellence
ACME	Accelerated Climate Model for Energy
ACTS	Assessment and Commitment Tracking System
ADIOS	Adaptable IO System
ALCC	ASCR Leadership Computing Challenge
ALCF	Argonne Leadership Computing Facility
ALICE	A Large Ion Collider Experiment
ALPS	Application Level Placement Scheduler
ALTD	Automatic Library Tracking Database
ANL	Argonne National Laboratory
ASCR	Advanced Scientific Computing Research
ATO	authority to operate
ATP	adenosine triphosphate
BNL	Brookhaven National Laboratory
BPIO	balanced placement in/out
BRAIN	Brain Research through Advancing Innovative Neurotechnologies
C&A	certification and accreditation
CAAR	Center for Accelerated Application Readiness
CADES	Compute and Data Environment for Science
CASL	Center for Advanced Simulation of Light Water Reactors
CESM	Community Earth System Model
CLM	Community Land Model
CNMS	Center for Nanophase Materials Science
CPU	central processing unit
CR	computational readiness
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
DIMM	dual in-line memory module
DNE	Distributed Namespace
DNS	direct numerical simulation
DOE	Department of Energy
DOE-OSO	DOE ORNL Site Office
DOI	digital object identifier (used to provide a lasting link to an online document)
ESGF	Earth System Grid Federation
EVEREST	Exploratory Visualization Environment for Research in Science and Technology
FY	fiscal year
GB	gigabyte
GPU	graphics processing unit
GUI	graphical user interface
HACC	Hardware/Hybrid Accelerated Cosmology Code
HEP	High Energy Physics

HiPSTAR	High-Performance Solver for Turbulence and Aeroacoustics Research
HOMME	High-Order Methods Modeling Environment
HPC	high-performance computing
HPSS	High Performance Storage System
HZDR	Helmholtz Zentrum Dresden Rossendorf (Helmholtz Association of German Research Centre)
INCITE	Innovative and Novel Computational Impact on Theory and Experiment
ISM	integrated safety management
IT	information technology
JIF	journal impact factor
LANL	Los Alamos National Laboratory
LBNL	Lawrence Berkeley National Laboratory
LCF	Leadership Computing Facility
LDRD	Laboratory Directed R&D
LHC	Large Hadron Collider
LLNL	Lawrence Livermore National Laboratory
LNET	Lustre networking
LSMS	Locally Self-Consistent Multiple Scattering
MB	megabyte
MPI	Message Passing Interface
MTTF	mean time to failure
MTTI	mean time to interrupt
MUI	multiscale universal interface
NCCS	National Center for Computational Sciences
NCSA	National Center for Supercomputing Applications
NERSC	National Energy Research Scientific Computing Center
NICS	National Institute of Computational Sciences
NLM	National Library of Medicine
OA	overall availability
OAR	Operational Assessment Report
OLAP	Online Analytical Processing
OLCF	Oak Ridge Leadership Computing Facility
ORISE	Oak Ridge Institute for Science and Education
ORNL	Oak Ridge National Laboratory
OSO	ORNL Site Office
OST	object storage target
OUG	OLCF Users Group
ParRep	Parallel Replica
PB	petabyte
PDACS	Portal for Data and Analysis Services for Cosmological Simulations
PETSc	Portable, Extensible Toolkit for Scientific Computing
PIC	particle-in-cell
PM	preventative maintenance
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
SA	scheduled availability
SBMS	Standards-Based Management System
SC	sickle cell anemia
SCEC	Southern California Earthquake Center
SciComp	Scientific Computing Group
SciDAC	Scientific Discovery through Advanced Computing
SM	streaming multiprocessor
SMP	symmetric multiprocessing
SNS	Spallation Neutron Source
TB	terabyte
TRC	Total Recordable Cases
TVA	Tennessee Valley Authority
UAO	User Assistance and Outreach
UIUC	University of Illinois at Urbana-Champaign
UT-B	UT-Battelle
UTK	University of Tennessee–Knoxville
WL-LSMS	Wang-Landau Locally Self-Consistent Multiple Scattering
WMS	workload management system

Executive Summary

HIGH PERFORMANCE COMPUTING FACILITY 2015 OPERATIONAL ASSESSMENT OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2016

EXECUTIVE SUMMARY

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

Calendar year (CY) 2015 was filled with outstanding operational results and accomplishments: a very high rating from users on overall satisfaction that ties the highest-ever mark set in CY 2014; the greatest number of core-hours delivered to research projects; the largest percentage of capability usage since the OLCF began tracking the metric in 2009; and success in delivering on the allocation of 60, 30, and 10% of core hours offered for the INCITE (Innovative and Novel Computational Impact on Theory and Experiment), ALCC (Advanced Scientific Computing Research Leadership Computing Challenge), and Director's Discretionary programs, respectively. These accomplishments, coupled with the extremely high utilization rate, represent the fulfillment of the promise of Titan: maximum use by maximum-size simulations.

The impact of all of these successes and more is reflected in the accomplishments of OLCF users, with publications this year in notable journals *Nature*, *Nature Materials*, *Nature Chemistry*, *Nature Physics*, *Nature Climate Change*, *ACS Nano*, *Journal of the American Chemical Society*, and *Physical Review Letters*, as well as many others. The achievements included in the *2015 OLCF Operational Assessment Report* reflect first-ever or largest simulations in their communities; for example Titan enabled engineers in Los Angeles and the surrounding region to design and begin building improved critical infrastructure by enabling the highest-resolution Cybershake map for Southern California to date.

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 workflows was featured at the Department of Energy (DOE) booth at SC15, for the second year in a row, highlighting work with researchers from the National Library of Medicine coupled with unique computational and data resources serving experimental and observational data across facilities.

Effective operations of the OLCF play a key role in the scientific missions and accomplishments of its users. Building on the exemplary year of 2014, as shown by the 2014 Operational Assessment Report (OAR) review committee response in Appendix A, this 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 CY 2015, which, unless otherwise specified, denotes January 1, 2015, through December 31, 2015.

COMMUNICATIONS WITH KEY STAKEHOLDERS

Communication with the Program Office

The OLCF regularly communicates with the Advanced Scientific Computing Research Program Office through a series of 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 OAR, 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 sponsors, potential users, and the public.

Communication with the User Community

The OLCF's communications with users take a wide variety of forms and are tailored to the objectives of relating science results to the larger community and helping users to more efficiently and effectively use OLCF systems. The OLCF offers many training and educational opportunities throughout the year for both current facility users and the next generation of HPC users (see Section 1.4.5).

The impact of OLCF communications is assessed as part of an annual user survey. The mean rating for users' overall satisfaction with OLCF communications was 4.4 in 2015. Ninety-four percent of respondents (282) rated their overall satisfaction with communications from the OLCF as "satisfied" or "very satisfied. Ninety-six 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
- automated notifications of system outages
- OLCF website
- conference calls
- OLCF User Council and Executive Board
- one-on-one interactions through liaisons and analysts
- social networking vehicles

SUMMARY OF 2015 METRICS

In consultation with the DOE program sponsor and as proposed in the 2014 OAR, a series of metrics and targets were identified to assess the operational performance of the OLCF in CY 2015. The 2015 metrics, target values, and actual results as of December 31, 2015, are noted throughout this report and are summarized in Section 8. The OLCF exceeded all of the metric targets.

RESPONSES TO RECOMMENDATIONS FROM THE 2014 OPERATIONAL ASSESSMENT REVIEW

The OLCF did not receive any operational recommendations resulting from the 2014 Operational Assessment Review. The OLCF continued to operate at an exemplary level in 2015 as staff worked to further refine GPU usage data collection and reporting techniques as recommended in the 2013 Operational Assessment Review.

OPERATIONAL REALIGNMENTS TO BETTER SERVE OLCF STAKEHOLDERS

Recognizing a need within the user community, the OLCF responded with several strategic hires in CY 2015. A new group was formed to focus on and assist users with data analysis, visualization, and workflow support. The new group, Advanced Data and Workflows, consists of data and visualization scientists who have been instrumental in OLCF's rich history of user collaboration and operational excellence. The Advanced Data and Workflows group is led by Sreenivas Rangan Sukumar, who comes to the OLCF from ORNL's Computational Science and Engineering Division, where he conducted transformational work in data science.

Kathlyn Boudwin left her role as the OLCF Deputy Project Director to serve as the Associate Projects Director for the Computer and Computational Sciences Directorate at ORNL. The OLCF hired Justin Whitt as Deputy Project Director. He came to OLCF from the National Institute for Computational Sciences (NICS), within which he held multiple leadership roles. He brings years of management experience from academia, private industry, and HPC to the OLCF.

The National Center for Computational Sciences (NCCS), the ORNL division that hosts the OLCF, recognized a need for focused oversight of both present and future systems. To that end, Jim Rogers moved from the NCCS Director of Operations position to the position of Director of Computing and Facilities. In this role, he focuses on future systems and the facilities that will power, cool, and host those systems. The NCCS hired Stephen McNally as Operations Manager to focus on the day-to-day operation of the existing systems and facilities. As Operations Manager, he ensures that the center operates efficiently and remains responsive to changing user needs. He also comes from NICS, where he served as the HPC Operations Group Leader, managing and directing technical staff in an HPC environment.

User Results

HIGH PERFORMANCE COMPUTING FACILITY 2015 OPERATIONAL ASSESSMENT OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2016

1. USER RESULTS

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

OLCF RESPONSE: Yes. In 2015, the Oak Ridge Leadership Computing Facility (OLCF) supported 1,176 users on over 316 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 in 2015 was 4.6. Overall ratings for the OLCF were positive; 95% of users reported being “satisfied” or “very satisfied.”

The center measures its performance using a series of quantifiable metrics. The metric targets are structured to ensure that users are provided prompt and effective support and that the 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 2015 with 92% of tickets being resolved within 3 business days. The OLCF continues to enhance its technical 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 scientific liaisons; multiple channels for communication with users, including the OLCF User Council; and training programs, user workshops, and tools to reach and train both current facility users and the next generation of computer and computational scientists. The success of these activities and identification of areas for development are tracked through the annual OLCF user survey.

In an effort to promote continual improvement at the OLCF, users are sent surveys soliciting their feedback regarding support services and their experience as users of the facility. The 2015 survey was launched on October 6, 2015, and remained open for participation through December 2, 2015. The survey was sent to 873 users 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 who logged into an OLCF system between January 1, 2015, and October 1, 2015. OLCF staff members were excluded from participation. A total of 308 users completed the survey, for an overall response rate of 35%.

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 were 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 respondents' responses to other items that were directly related 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 2015 survey can be found on the OLCF website.¹

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

Table 1.1. 2015 user result metrics summary

Metric description	2014 target	2014 actual	2015 target	2015 actual
Overall OLCF Satisfaction score on the user survey	3.5/5.0	4.6/5.0	3.5/5.0	4.6/5.0
Show improvement on results that scored below satisfactory in the previous period	Results will show improvement in at least 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.	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 2015 survey.
OLCF survey results related to problem resolution.	3.5/5.0	4.6/5.0	3.5/5.0	4.6/5.0
Percentage of user problems addressed within 3 business days	80%	90%	80%	92%
Average of all user support services ratings	3.5/5.0	4.6/5.0	3.5/5.0	4.6/5.0

1.2 USER SUPPORT METRICS

The OLCF exceeded all of the User Support metrics for 2015. 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 2014		CY 2015	
	Target	Actual	Target	Actual
Overall OLCF Satisfaction rating	3.5/5.0	4.6/5.0	3.5/5.0	4.6/5.0
Average of all user support services ratings	3.5/5.0	4.6/5.0	3.5/5.0	4.6/5.0

1.2.1 Overall Satisfaction Rating for the Facility

Users were asked to rate their satisfaction on a 5-point scale, where a score of 5 indicates a rating of “very satisfied” and a score of 1 indicates a rating of “very dissatisfied.” The metrics agreed upon by the DOE OLCF program manager define 3.5/5.0 to be satisfactory.

¹ <http://www.olcf.ornl.gov/media-center/center-reports/2015-outreach-survey/>

Overall ratings for the OLCF were positive with 95% of users responding that they are satisfied or very satisfied with the OLCF overall.

Key indicators from the survey, including overall satisfaction, are shown in Table 1.3. They are summarized and broken out by program. The data show 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	4.7/5.0	4.6/5.0	4.6/5.0
Overall satisfaction with support services	4.4/5.0	4.5/5.0	4.4/5.0	4.4/5.0
Overall satisfaction with compute resources	4.6/5.0	4.6/5.0	4.5/5.0	4.6/5.0
Overall satisfaction with data resources	4.4/5.0	4.5/5.0	4.4/5.0	4.5/5.0

1.2.2 Average Rating across All User Support Questions

The calculated mean of all answers to all user support services questions on the 2015 survey was 4.6/5.0, indicating that the OLCF exceeded the 2015 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 the survey respondents. The following comments are samples from the survey:

“OLCF give[s] excellent support to industry users. My experience has always been Very Satisfied”

“The OLCF support pages have been improved in terms of finding helpful resources. Thank you!”

“The support staff and responsiveness of help@olcf.ornl.gov are excellent. They resolve any issues I have very quickly.”

“The overall user experience is the best I have ever had.”

1.2.3 Improvement on Past Year Unsatisfactory Ratings

Each year the OLCF works to show improvement on no less than half of any questions that scored below satisfactory (3.5/5.0) in the previous year’s survey. All questions scored above 3.5 on both the 2014 and 2015 surveys. However, based on feedback received from the 2014 User Survey in conjunction with other feedback channels, the OLCF took the following actions in 2015 to enhance the user experience at the OLCF.

- Increased the memory on each of Rhea’s 512 compute nodes from 64 to 128 GB
- Added nine GPU nodes for Rhea
- Simplified the Rhea queue policy, making the scheduling environment more intuitive and easier to use
- Increased home directory quotas from 10 to 50 GB
- Increased Globus transfer window to 72 hours
- Added new system status indicators to the OLCF website for the OLCF file systems

- Added 18 PB of disk cache to the High Performance Storage System (HPSS) to decrease the time required for storing, accessing, or retrieving data
- Upgraded from 10-gigabit to 40-gigabit Ethernet connectivity for all the HPSS movers to help speed up the migration of data from the disk cache to the long-term tape storage
- Released a new tool called GLprof, built upon Gleipnir (Section 4.4.2), to help users more easily manage cache behavior
- Installed Darshan for profiling the I/O behavior of applications on Titan, Eos, and Rhea

1.2.4 Assessing the Effectiveness of the OLCF User Survey

Before sending the user 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. Jack Wells (OLCF Director of Science) sent the initial survey invitation on October 6, 2015, and subsequent follow-up reminders were sent by Jim Hack (National Center for Computational Sciences [NCCS] Director), Buddy Bland (OLCF Project Director), ORISE, the OLCF User Council, and individual members of the OLCF. The survey was advertised on the OLCF website and 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 of one to two years. The number of users in the more than two years category increased in 2015. (Table 1.4).

Table 1.4. User survey participation

	2014 survey	2015 survey
Total number of respondents (Total percentage responding to survey)	312 (32%)	308 (35%)
New users (OLCF user <1 year)	26%	27%
OLCF user 1–2 years	27%	22%
OLCF user >2 years	48%	51%

1.2.4.1 Statistical Analysis of the Results

A statistical analysis of key survey areas is shown in Table 1.5. The results reflect overall satisfaction with the facility, services, and computational resources.

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 rating was above two deviations from the mean.

Table 1.5. Statistical analysis of key results

Survey topic	Number of survey respondents	Number of survey responses to this question	Mean	Variance	SD
Overall satisfaction with the OLCF	308	304	4.6	0.55	0.74
Overall satisfaction with Titan	308	257	4.5	0.30	0.55
Overall satisfaction with data resources	308	237	4.4	0.71	0.84
Overall satisfaction with support services	308	298	4.4	0.74	0.86
Overall satisfaction with communications	308	282	4.4	0.42	0.65
Overall satisfaction with training	308	247	4.3	0.48	0.69
Overall satisfaction with the website	308	272	4.4	0.40	0.63

1.3 PROBLEM RESOLUTION METRICS

The operational assessment review 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. 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; then provide a workaround to the user, if possible, and track the issue to resolution with the product vendor, where resolution could be a bug fix or workaround acknowledgement.

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 2,395 tickets in response to user queries for CY 2015. On average, the OLCF fielded roughly 200 tickets per month, which matches the actual monthly ticket distribution for the year. The center exceeded the problem-resolution metric and responded to 92% of the queries within 3 business days (Table 1.6).

Table 1.6. Problem resolution metric summary

Survey Area	CY 2014		CY 2015	
	Target	Actual	Target	Actual
Percentage of problems addressed in 3 business days	80%	90%	80%	92%
Average of problem resolution ratings	3.5/5.0	4.6/5.0	3.5/5.0	4.6/5.0

Tickets are categorized by the most common types. The top three reported categories in 2015 were account issues/queries, running jobs, and job compilations (Figure 1.1).




Figure 1.1. Categorization of help desk tickets.

1.4 USER SUPPORT AND OUTREACH

The requested OAR 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, provided in Appendices B and C, respectively

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

- A user support staff made up of account management liaisons, User Assistance and Outreach (UAO) analysts, Scientific Computing Group (SciComp) liaisons, data liaisons, and visualization liaisons.
- Multiple vehicles to communicate with users, sponsors, and vendors

- Developing and delivering training to current and future users
- The strong outreach component needed to interface with the next generation of HPC users, the external media, and the public

1.4.1 User Support

The OLCF recognizes that users of HPC facilities have a 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; liaisons in respective scientific, data, and visualization areas; and computer scientists who assist on issues surrounding the programming environments and tools. The following list details some of the high-level support activities of CY 2015 and the specific OLCF staff resources available to assist users.

- **Applying and Improving the OpenACC Standard.** OLCF staff often partner closely with users who are attempting to work at the cutting edge of available tools. As Professor Michael Zingale and his group at Stony Brook University recently set out to port their supernova modeling code MAESTRO to GPUs using OpenACC, they encountered some challenges. ORNL’s representative to the OpenACC standards committee, Oscar Hernandez, quickly stepped in to assist. He worked with Zingale’s team and the compiler vendor to isolate several bugs and identify enhancements to the implementation of version 2.0 of the OpenACC standard, which was then new. He also identified several minor deficiencies in the standard and worked with the standards committee to revise the language specification itself to allow a simpler approach to managing data between the host and the GPU device.

As a result, Zingale’s group has an initial version of the MAESTRO code that shows a 400% speedup in the portions of the code ported to the GPUs and offers the potential to increase the physical fidelity of the model. Changes to the OpenACC standard have appeared as part of version 2.5, which was released in November 2015.

- **On-site Vendor Staff Assist Users and Enhance Tools.** The OLCF partners with vendor staff to improve time to solution and collaboration between users and vendors through matrixed support staff from the Computer Science Research Group. The OLCF discovered this highly effective model and has continued it by hosting Alinea staff member Nick Forrington and Technical University of Dresden researcher Frank Winkler. Forrington and Winkler worked to enable their tools to perform optimally with researcher applications. Figure 1.2 illustrates one result of this partnership; Forrington worked to visually represent the software structure of the Community Land Model (CLM), showing frequently accessed areas of the application.




Figure 1.2. The software (subroutine) of CLM. Image credit: Wang, D. et al. *Environmental Modelling and Software* 55 (2014).

- **Wraprun.** OLCF User Assistants developed a tool named Wraprun to assist Titan users who wish to run ensemble jobs—ones consisting of multiple independent jobs that are run simultaneously—but have difficulty using the default job launcher, Cray’s Aprun. By bundling

the jobs, Wraprun allows users to launch more jobs in less time, increasing efficiency and productivity.

- **Spider II Performance Status Indicator.** The facility has historically provided system status indicators for the compute systems and archive resources on the OLCF website. In the 2014 survey, users noted they would like the same status indicators for the file systems. Through a collaborative effort among multiple groups, the OLCF developed a tool to provide users with a “red/yellow/green” status for the Spider II file system.
- **Improving Globus GridFTP Transfers.** To simplify Globus GridFTP transfers, the OLCF moved away from the use of third-party certificates and deployed a MyProxy x509 PKI server and Certificate Authority that uses OAuth delegation to generate and distribute temporary proxy grid certificates for Globus GridFTP transfers. Users need only approve the delegation using their valid OLCF credentials to generate a 72 hour proxy certificate. This greatly reduces the data transfer complications and GridFTP lead-time as users no longer need to obtain, manage, protect, distribute, and activate their own third-party personal x509 certificates.
- **Improving Communication.** The OLCF uses a set of system-specific email lists to provide automated notifications of system state changes and other notable system events. Because the email frequency is higher than for standard email notifications, membership on the lists is optional. Historically, users had to opt-in to receive the emails; but an evaluation of that process determined that an opt-out method would be more beneficial to the user community. In 2015, the OLCF developed a process to automate populating each system’s high-volume mail list with users who are actively using the system. An active user is defined as one who has connected to a system or used the system’s batch queue within a recent period of time. Maintaining an email list of active users avoids the need for users to sign up and enables the OLCF to email only users actively using a resource without sending unnecessary emails to users not currently using it.
- **Resource and Allocation Tracking System.** The UAO Development team continued making additions and improvements to the OLCF’s in-house customer relationship management software to increase security, track additional information, improve the user interface, provide additional data validation, and increase the functionality and efficiency of the Resource and Allocation Tracking System (RATS) GUI. The team developed and deployed a new API layer with 25 endpoints, which provides a more secure and accessible way for disparate systems and processes to communicate with RATS. Previously, RATS data was accessible only through direct database access, which was found to be inefficient. The team redesigned the allocation GUI and backend to be more intuitive and built mechanisms for identifying and correcting any deficiencies in the allocation dataset.
- **Operational Optimizations in Projects and Accounts Projects.** In 2015, the OLCF analyzed each step in the process of enabling a new project and determined that the Export Control review step was an area for improvement to reduce processing time. The OLCF partnered with the ORNL Export Control office to identify additional resources that are now being used to conduct OLCF reviews in a more prioritized manner. The change occurred in October 2015, and the OLCF has experienced improved turnaround time on Export Control reviews since then. The

2016 INCITE Export Control reviews were all completed before the end of November; that left adequate time to complete all remaining approval steps in time for all INCITE projects to start on January 1, 2016 with at least one user active for each project. The OLCF also contributed to the development of a new ORNL central user agreement database managed by the ORNL Partnerships Directorate. The new database makes it easier to track user agreements that are in progress or to identify user agreements that have already been executed and those set to expire in the near future. Before the creation of the new database, much of the information needed to complete the user agreements was found in spreadsheets that were difficult to manage.

1.4.2 User Assistance and Outreach Analysts

UAO analysts provide the interface for users, address user queries, act as user advocates, cover front-line ticket triage, resolution, and escalation, provide user communications, develop and deliver training and documentation, and install third-party applications for use on the computational and data resources.

1.4.3 Scientific Liaisons

In addition to UAO staff, the OLCF provides certain projects with scientific (1.4.3), data (1.4.4), and visualization (1.4.5) liaisons. These liaisons are a unique OLCF response to HPC problems faced by users. The OLCF actively partners with project principal investigators to assist in improving operational scaling and efficiencies. Several examples of scientific liaison collaboration with and support of users are provided in Sections 1.4.3.1, 1.4.3.2, and 1.4.3.3.

1.4.3.1 Multiscale Simulations of Human Pathologies

The INCITE project “Multiscale Simulations of Human Pathologies,” led by George Karniadakis of Brown University, was one of five finalists selected in 2015 for the prestigious Gordon Bell Prize. This award is given each year to recognize outstanding research efforts in HPC. Karniadakis’s team was recognized for its computational work to simulate human blood flow to develop improved drug delivery methods and better predictors to fight tumor formation and sickle cell anemia.

The OLCF played a significant supportive role in this simulation effort, as liaison Wayne Joubert assisted the team in reaching its project goals. The assistance included guiding the team’s efforts to undertake a major computational campaign requiring the entirety of the Titan resource over a long period of time, helping the team make good use of Titan’s large node count, installing software required by the project, and performing simulation runs needed to demonstrate the scalability of the algorithms and code. Petros Koumoutsakos, one of the project leaders, affirmed that the support provided by Wayne and the OLCF has been critical to this project. Now in its third year, the Karniadakis team will continue to rely on the OLCF for assistance in reaching its science goals.^{2,3}

1.4.3.2 PDACS—Portal for Data and Analysis Services for Cosmological Simulations

The Portal for Data and Analysis Services for Cosmological Simulations (PDACS) has been developed by the Hardware/Hybrid Accelerated Cosmology Code (HACC) team at Argonne National Laboratory over several years. PDACS is an open, web-based platform that allows the download, transfer, and manipulation of simulation data, as well as computational analysis of the data using available resources. The framework enables the wrapping of analysis tools written in a number of languages and makes them available to the user within its workflow system using a convenient and powerful dataflow programming paradigm. The system allows for provenance tracking, implementing a transparent sharing method as well as an important resource for checking the reproducibility of results generated by the workflows. Users are able to submit their own tools to the system, and a test suite is being developed to provide acceptance testing of such tools.

Under INCITE, the HACC team recently completed the “Q Continuum” simulation on Titan, a large-scale structure simulation covering a volume of $1,300^3$ megaparsecs and evolving more than half a trillion particles. The team partnered with OLCF and researchers at Los Alamos National Laboratory (LANL) to

² <http://www.hpcwire.com/2015/10/14/cray-xk7-titan-used-to-simulate-complicated-blood-flow/>

³ <https://www.olcf.ornl.gov/2015/10/13/flowing-toward-red-blood-cell-breakthroughs/>

develop a new in situ/co-scheduling workflow for HACC analysis tasks at the OLCF.⁴ The resulting analysis of these simulation data will enable interesting new science results, including the provision of sophisticated sky maps to the major dark energy surveys. Toward this end, Level II data (i.e., partially analyzed data) from the Q Continuum simulation have been moved to CADES (Compute and Data Environment for Science) and a virtual machine has been deployed on CADES to allow a wide community of researchers to use PDACS, including observational astronomers.

This new workflow was an important first step in making the full science potential of the Q Continuum simulation available. The HACC team previously worked closely with LANL researchers under the Scientific Data Management, Analysis and Visualization SciDAC (Scientific Discovery through Advanced Computing) program to develop a halo finder that can take advantage of Titan's GPUs. These tools are also useful in the PDACS workflow. Working closely with colleagues in CADES, Bronson Messer, the SciComp liaison enabled submission to CADES' new XK7 resource directly from the HACC data virtual machine, fully enabling the PDACS workflow with the maximum possible performance. In the coming year, the team expects to release simulation data and catalogs via PDACS backed by this new capability.

1.4.3.3 Performance Optimization and Capability Improvements of LSMS

Markus Eisenbach used his intimate knowledge of the OLCF to serve as a highly effective principal investigator on the INCITE project "Scalable First Principles Calculations of Materials at Finite Temperatures."

The project actively develops and uses the Locally Self-Consistent Multiple Scattering (LSMS) code for calculating the electronic structures of materials. The INCITE scientific liaison provided support for benchmarking and debugging the LSMS code and scaling it to the full Titan system. The code had already been ported to efficiently utilize the GPUs on Titan and was further optimized for more efficient use of time and energy to solution. The success of this work is evident in both the weak and strong scaling efficiency achieved by LSMS on Titan: the team has demonstrated that the GPU accelerated code shows near-perfect weak scaling from calculations of 16 atoms on 4 nodes of Titan requiring 67.3 s per iteration, to 65,536 atoms on 16,384 Titan nodes requiring 69.9 s per iteration. This opens the opportunity for fully first principles ground-state calculations of metal and alloy systems of up to 500,000 atoms, which enables the study of realistic experimentally accessible nano-systems.

The main work has concentrated on improving the combination of first principles electronic structure calculations with LSMS and the sampling of finite temperature calculations using Wang-Landau Monte-Carlo methods (WL-LSMS). In the past year, the capabilities of the code have been extended to enable the sampling of alloy ordering transitions in addition to the previously available sampling of magnetic order. The OLCF liaison provided important support in debugging this new capability that for the first time enables the direct first principles calculation of phase transitions in alloys without the need for mean field theories or model fitting. The team benchmarked the WL-LSMS code and demonstrated the strong scaling of the code for a 1,024-atom system. The GPUs enabled the WL-LSMS code on Titan to generate 0.0208 samples/s on 257 nodes and 0.9128 samples/node on 12,289 nodes for a parallel efficiency of 92%. In the past, the team has also demonstrated a 7.3× improvement in energy-to-solution of a GPU-enabled version of the code on Titan compared with CPU-only WL-LSMS code. Finally, their intimate knowledge of the code enabled the team to provide important input for the discussion of future programming models for accelerated and threaded architectures, and WL-LSMS has been used to test profiling and debugging tools at scale.

⁴ Christopher Sewell et al. "Large-scale compute-intensive analysis via a combined in-situ and co-scheduling workflow approach," pp. 50:1–50:11 in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC15, New York, USA (2015). ISBN 978-1-4503-3723-6.

1.4.4 Data Liaisons

1.4.4.1 ADIOS-enabled FireFOAM

The OLCF supported FM Global's research and simulations on Titan and Eos. In particular, the FM Global team benefitted from a close collaboration with Norbert Podhorszki, a computational scientist within the OLCF who specializes in end-to-end I/O solutions. He helped the FM Global team with scaling issues related to data movement. By employing ADIOS, the team was able to experience a 12× I/O performance increase compared with the original configuration. In total, the 576 core Eos simulation went from requiring 2 hours and 5 minutes to taking just under 40 minutes.

1.4.4.2 GPU-Driven Discovery of a New Method of Particle Acceleration

Michael Bussmann's goal is to understand and control a new method of particle acceleration and replace particle accelerators with high-powered lasers for radiation therapy, to make the therapy available for treating cancer patients (Figure 1.3). Bussmann, of the Helmholtz Zentrum Dresden Rossendorf (HZDR), and his students developed the PiconGPU code to study accelerator science and are now using this code to aid in this important research for cancer therapy.

To observe atomic-scale interactions happening many times per second, Bussmann's team is using Titan's unique heterogeneous computing capability to enable heavy utilization of the GPUs. This allows them to speed up their research by over an order of magnitude compared with other teams using CPU-only codes. As was mentioned in the 2014 OLCF OAR, their code was a finalist for the 2013 Gordon Bell Prize; and during their simulations, they noticed that the time spent on I/O was a major cause of problems. The team needs to output data during the entire simulation with as many time steps as possible to study physical effects. Consequently, on Titan, hundreds of terabytes of data are generated every hour. ADIOS is used to write out the data on the entire span of a partition of the Spider II file system, achieving a writing speed of over 180 GB/s. The team not only sped up their I/O by 15×, compared with their previous solution of parallel HDF5, but also achieved the fastest observed data transfer on Titan for that application.




Figure 1.3. Proton density after laser impact on a spherical solid density target: irradiated by an ultra-short, high intensity laser (not in picture), the intense electro-magnetic field rips electrons apart from their ions and creates a plasma. By varying the target geometry and laser properties, scientists could find optimal regimes to accelerate high quality, directed ion beams that are currently studied in accompanying experiments. *Image credit: Alex Huebl, HZDR; David Pugmire, ORNL.*

1.4.4.3 Total Embraces ADIOS

Total is a French multinational integrated oil and gas company and one of the six major oil companies in the world. Its business covers oil, gas, and crude oil exploration and production to power generation, transportation, refining, and product trading. Total has one of the fastest computers in the world, delivering over 4.4 petaflops. Because Total needs even more computing power, it has been using Titan to run some of its more complex problems. In their efforts to scale their codes, Total researchers found that the I/O for their seismic imaging efforts needed to be improved, and they investigated the use of ADIOS. In their extensive study, which will be presented at the Houston Oil and Gas meeting in March 2016, they found that ADIOS provides the best I/O framework for their applications.

Currently, these applications take terabytes of data as inputs and output petabytes of result data, along with intermediate data; and they require good performance for structured grid data. For the evaluation, Total used one of the applications (RTM) from its SGI system with OLCF's Titan and compared the performance of the MPI-IO, HDF5, ADIOS and SIONLib I/O implementations. The ADIOS team provided guidance for using ADIOS optimally for Total's use cases. The team's investigation found that ADIOS was easy to use and worked extremely well on the two machines on which they tested it. They will begin a more intensive collaboration with the ORNL team to further integrate their codes with ADIOS.

1.4.4.4 Earth System Grid Federation Node at ORNL in Support of ACME

The Accelerated Climate Model for Energy (ACME) is the flagship DOE earth system model for simulating the coupled climate system at high resolution. The model is designed and is being optimized to exploit DOE leadership-class computer systems. The modeling experiments, facilitated via the INCITE project led by Mark Taylor of Sandia National Laboratories, are designed to address the scientific goals of understanding the interactions of the climate system with the hydrological and biogeochemical cycles and the cryosphere-ocean system. Currently, the ACME project includes more than 100 scientists and technical staff who are involved in model development, performance optimization, model validation, exploratory computational campaign, data management and workflow orchestration. OLCF staff have been integral members of the ACME project since its inception. The ACME model development process and the exploratory computational campaign challenges involve the proper configuration and execution of the model among facilities such as the National Energy Research Scientific Computing Center (NERSC), the Argonne Leadership Computing Facility (ALCF), and OLCF. The ACME workflow helps manage data using the Earth System Grid Federation (ESGF) node⁵ co-developed and deployed at the ORNL CADES facility and managed by John Harney of the Advanced Data and Workflows group, with additional system support from the HPC Operations group. In addition, the Rhea analysis cluster at OLCF is the primary system used by ACME to integrate and synthesize the model results from the various model development and tuning exercises of the computational campaign across geographically disparate computational facilities.

The ACME scientists faced logistical challenges in sharing simulation output and value-added results from a multitude of modeling experiments. The ESGF is a multi-agency international collaborative effort to develop, deploy, and operate a set of geographically distributed "nodes" that employ common federation protocols for publication, search, discovery and downloading of climate data, including simulation output. The data and associated metadata are catalogued, stored, and managed by every node independently of other nodes in the federation; but all users are able to query and access data seamlessly as if the data are in a single global archive.

As shown in Figure 1.4, each ESGF node provides a set of services that facilitate user management and data access according to the distribution policy. To date, the ACME INCITE project has published over 100 TB via the node at ORNL. The data is accessible on the ORNL ESGF data portal, which is also a discovery path to observational data archived by other repositories such as the Carbon Dioxide Information Analysis Center and the Atmospheric Radiation Measurement User Facility. The ACME scientists have been able to rapidly and seamlessly publish simulation output and experimental results using the ORNL ESGF node, resulting in enhanced communication and increased productivity for a collaborative next-generation model development process.

⁵ <https://esg.ccs.ornl.gov>

1.4.5 Visualization Liaisons

1.4.5.1 Enabling Accelerated Graphics on Titan Creates World's Fastest Visualization Cluster

The team of Jamison Daniel, Mike Matheson, and Benjamin Hernandez have deployed ParaView (a scientific visualization suite) on the Titan, Rhea, Focus, and EVEREST clusters. Titan uses general-purpose GPUs consisting of Tesla K20X accelerators, which make up roughly 90% of the system's floating point computational capability. These accelerated graphics cards were originally inaccessible to applications such as ParaView, as the rendering could only be done in a statically linked, cross-compiled, non-threaded version of MESA that was executed on the CPU.

Users have been requesting the ability to use Titan for GPU-accelerated graphics and remote rendering to enable the full graphics capacity of the system for in situ

visualization and analysis. Before this capability was available, software rendering and X-forwarding was the only option, even though 90% of Titan's computational power comes from GPUs.

With this new capability, the OLCF enabled GPU-accelerated in situ workflows through the latest versions of ParaView/Catalyst, and users can run their personalized visualization tools remotely using VirtualGL.

As a result, ParaView visualization and Catalyst in situ visualization jobs run significantly faster than before, and the approach reduces or eliminates data transfers from Titan to the visualization and analysis systems. One of the users taking advantage of these features is Peter Vincent of Imperial College, who is studying the impact of air flow profiles on the acoustic properties of jet engines.

The visualization team also worked with Kitware to develop and deploy a complete modernization of the OpenGL interface in ParaView. The new version of ParaView 5.0 demonstrates performance increases by a factor of more than 1,000 for polygonal models (from 0.3 million triangles per second per processor [single threaded MESA] to 450 million triangles per second per GPU [OpenGL 2]).^{6,7}




Figure 1.4. Graphical representation of ACME ESM workflow.

⁶ <http://blogs.nvidia.com/blog/2015/11/17/titan-largest-gpu-visualization/>

⁷ <http://www.kitware.com/blog/home/post/1008>

1.4.5.2 Largest-ever Vortex Visualization

A scientific visualization for an INCITE project led by Susan Kurien of Los Alamos National Laboratory and supported by OLCF's Mike Matheson and Duane Rosenberg was featured in the American Institute of Physics journal *Physics of Fluids* (Figure 1.5). This project on Bolgiano-Obukov scaling in rotating-stratified turbulence required a full-resolution blowup (at linear grid resolution of 4,096 points) of a highlighted region of a 2-D slice of z-vorticity. With off-the-shelf tools failing, OLCF data liaisons marshaled a workflow using Titan and EVEREST resources (Focus and Lens) to help the users visualize their simulation results. "The figure shows the complex vortex network at the smallest scale of the problem—the largest of its type ever done. We are a featured article in the journal now," the user said.

1.4.5.3 BigNeuron Hackathon

The Brain Research through Advancing Innovative Neurotechnologies (BRAIN) Initiative is part of a new presidential focus on revolutionizing the understanding of the human brain. By accelerating the development and application of innovative technologies, researchers will be able to produce a revolutionary new dynamic picture of the brain that, for the first time, will show how individual cells and complex neural circuits interact in both time and space ([International Neuroinformatics Coordinating Facility](#)).

Attendees representing 13 organizations in government and the private sector worked to scale up neuron reconstruction algorithms at the OLCF to consolidate the knowledge base surrounding



Figure 1.6. Participants from the BigNeuron Hackathon, including scientific, data, and visualization liaisons.

after visual inspection and validation of the data, converge on a standard. Experts from the OLCF's new Advanced Data and Workflows group, along with CADES staff, helped the group leverage OLCF computing resources to render 3-D reconstructions using the parallel-processing infrastructure.




Figure 1.5. OLCF visualization on the cover of the May 2015 *Physics of Fluids*.

computational modeling of neurons.

Researchers at the Allen Institute for Brain Science and ORNL organized the first BigNeuron Data Analytics and Visualization Hackathon (partially funded by the International Neuroinformatics Coordinating Facility) that took place November 16–20 at ORNL (Figure 1.6).⁸ Hanchuan Peng, associate investigator at the Allen Institute and the PI on the OLCF DD project, used the hackathon as an opportunity to consolidate the computational neuroscience knowledge base and develop best practices for all researchers in the field. A major challenge for the group was to reconstruct 3-D neurons from hundreds of terabytes of data using popular methods and,

⁸ <https://www.olcf.ornl.gov/2016/01/05/bigneuron-hackathon-branches-out-at-olcf>

1.4.6 OLCF User Group and Executive Board

The OLCF User Group (OUG), consisting of all users, meets once a month by conference call and webinar to discuss OLCF news, resources, policies, and timely HPC tutorials and techniques. The OUG Executive Board, consisting of 11 members, meets shortly before or after the monthly call to provide the OLCF with in-depth feedback and guidance on topics such as training, facility resources, and policies. The board has also organized working groups to supply user feedback on specific topics, such as the annual user meeting. In 2015, the OUG Executive Board requested that OLCF find opportunities for users to share “user driven content” about lessons learned using OLCF resources. The OLCF worked with the board to identify users to share their content during two of the monthly OUG webinars. These talks were well received, and the OLCF plans to provide more of these opportunities in the coming year.

The tips, training, and user talks from the OUG conference calls/webinars are recorded and the resulting video is linked on the OLCF training events page. The OUG Executive Board requested that the OLCF make the videos searchable by topic and more easily accessed. In response, the OLCF launched a Training channel on Vimeo (<https://vimeo.com/olcf>). The channel is visible to the public and organizes all training videos into one list accessed from the channel’s homepage. Tags allow the videos to be searchable by topic.

More information about the OLCF User Group, including a list of the Executive Board members, can be found at <https://www.olcf.ornl.gov/about-olcf/oug/>.

1.4.7 Training, Education, and Workshops

Workshops, user conference calls, training events, and seminars are integral components of both user assistance and outreach. Training can obviate difficulties in performing science on such large-scale systems, and training events can serve to engage both the public and the user community. In addition to training users to use the resources available at the OLCF, the training program focused on an additional area: software development best practices.

See Appendix B for a summary of training events. A few of the notable 2015 events are highlighted below.

1.4.7.1 Center for Accelerated Readiness (CAAR) Workshops

The OLCF hosted two identical kick-off workshops for members of the new Center for Accelerated Application Readiness (CAAR) partnership projects during the spring of 2015. The workshops brought together the members of the CAAR partnership teams, including core developers of the new CAAR applications, staff from the OLCF, members of the IBM/NVIDIA Center of Excellence at ORNL, and staff from NERSC and ALCF. The kick-off workshops were designed to help the project teams begin preparations to ensure their scientific applications are ready to make effective use of Summit, OLCF’s next-generation supercomputer, when it enters production. A total of 79 people participated in the two events.

1.4.7.2 2015 Hackathons

The OLCF partnered with other institutions in 2015 to offer four GPU Hackathons, in which programmers from around the world gathered to gain GPU application portability expertise straight from the experts. Hackathons were held at the [National Center for Supercomputing Applications](#) (NCSA) at the [University of Illinois at Urbana-Champaign](#) and the [Swiss National Supercomputing Centre](#) (CSCS) in Switzerland, and two were hosted by the OLCF and held locally. The 2015 events totaled more than 173 attendees and included a diversified range of institutions and science disciplines. Much of the continued success of these events results from the dedication of the vendor partners and mentors. Following the formula for success, each of the nine attending scientific programming teams at the OLCF Hackathon was

matched with mentors from vendors IBM, Cray, The Portland Group, and NVIDIA. Also participating were mentors from the OLCF, CSCS, and Cornell University. The lessons learned from this year's Hackathons will aid programmers for years to come as they inevitably must adapt to newer, more diversified architectures. A few weeks after the Hackathon, Fernanda Foertter of the UAO group shared the event results with an even greater audience at the SC15 supercomputing conference in Austin, Texas.

1.4.7.3 2015 OLCF User Meeting

New and long-time Titan users attended the annual OLCF User Meeting in June to learn about, share, and discuss the most recent science OLCF users are conducting on Titan. Ninety people participated in the 3 day event held June 23–25. The meeting coincided with the election of new OUG Executive Board members. First-time board members Hai Ah Nam of LANL, Thomas Maier of ORNL, and Mark Taylor of Sandia National Laboratories were elected to 3 year terms, and Stephane Ethier of Princeton Plasma Physics Laboratory was elected to a second term. The first 2 days of the meeting focused on science produced on Titan and the computational challenges users overcame in the process. On the first day, LANL researcher Luis Sandoval gave the keynote talk, “Helium Bubble Growth in Tungsten under Realistic Rates.” Presentations on the second day included an overview by NVIDIA's Jeff Larkin of GPU computing and the OpenACC programming standard. The third day's activities looked ahead to next-generation supercomputing applications and the OLCF's next petascale machine, Summit.

1.4.7.4 International Workshop on the Lustre Ecosystem

OLCF staff were significant contributors to the First International Workshop on the Lustre Ecosystem held March 3–4 in Annapolis, Maryland. Engagement in this workshop was geared toward sharing operational expertise gained by the OLCF over the past decade of Lustre deployments on Jaguar and Titan. Experts presented in the tutorial sessions about topics in which the OLCF has gained significant operational expertise, such as Lustre networking (LNET) configuration, monitoring and performance analysis, and limitations of standard utilities as parallel file systems grow in size and performance.

On the second day, workshop participants heard presentations from seven different organizations with content focused on four major areas. The first session focused on data-intensive workloads and their impacts on the Lustre file system. The second session focused on managing Lustre file systems in day-to-day operation. The final sessions focused on monitoring tools for Lustre and recent hardware advances that specifically benefit Lustre. This collaborative workshop is the first in an annual series in which Lustre experts from academia, government agencies, and industry gather to share operational experiences and best practices.

1.4.8 Training and Outreach Activities for the Future Workforce

The OLCF maintains a broad program of collaborations, internships, and fellowships for young researchers. Forty-six student interns and postdoctoral researchers were supported from January 1, 2015, through December 31, 2015. More than 25 students completed summer internships at the OLCF in 2015, including Douglas Aaser, Jeremy Anantharaj, Aaron Barlow, Andrew Bowers, Shawn Cox, Marcela Crosariol, Mindy Earnest, Jonathan Freed, Benjamin Klein, William McElmurray, Samuel Migirditch, Broxton Miles, Christopher Muzyn, Sarah Neuwirth, Bin Nie, Miki Nolin, Ellias Palcu, Benjamin Smith, Kevin Song, Kun Tang, Christina Thiessen, Thomas-Allan Tison, Joel Venzke, Lipeng Wan, and Michael Wilder.

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

- Marcela Crosariol joined the UAO Group, coming from Brazil to provide OLCF users with technical support. After helping her mentor, Fernanda Foertter, develop user tutorials for Titan—the OLCF's flagship supercomputer—Crosariol began to develop tutorials for Summit. The tutorials will help users accelerate their programs to make them run faster and more efficiently on

the new system. After completing her degree in system analysis and design from Sao Jose Dos Campos Technological College during the current school year, Crosariol plans to seek another internship at ORNL or apply to a master's program in computer science. Ultimately, she plans to pursue a career combining parallel programming and user outreach.

- Kevin Song interned for the Computer Science Research Group in the OLCF's Computer Science and Mathematics Division under the supervision of Manjunath Gorentla Venkata. Song's primary focus was developing methods of ensuring reliability on supercomputers as they approach exascale. Song graduated from the University of California– Merced in May and is attending graduate school in computer science at the University of Texas at Austin.
- Jeremy Anantharaj, who interned for the Technology Integration Group under the mentorship of Ross Miller, spent his summer analyzing data. He completed a program that collects and analyzes performance data for the entire Atlas file system. He also developed a program that collects and analyzes I/O performance data on a per-job basis. The script collects job data from the file system and, unlike the computer itself, saves the data in a log file format. It enables users to analyze old records and compare them to detect whether there are trends in the speed at which the bandwidth is changing.
- Valentine Anantharaj mentored three undergraduate interns during the summer. Douglas Aaser, a freshman from the University of Tennessee–Knoxville, developed microapps characterizing the I/O patterns of the ACME model for climate simulation. Mindy Earnest, a junior from Clemson University, developed the scripting necessary to configure and run the Community Earth System Model (CESM) to quantify and characterize its energy and performance tradeoffs. The conclusion was that there is no significant penalty in the throughput of the low-resolution configuration of CESM when it is configured to use less energy. Joel Venzke, a senior from Drake University, helped prototype the post-processing of CESM simulation output using the Bellerophon framework. This mechanism will enable ACME to integrate some of the analysis during the model execution phase of the ACME experiment workflow.
- ORNL staff members were guest lecturers for a first-time offering of a Data Center Design and Management class at the University of Tennessee–Knoxville (UTK). The senior-level course was developed in collaboration with professor Mark Dean of the UTK College of Engineering. It was pursued as a strategic offering to address an emerging gap in the workforce of skilled managers with sufficient understanding of both computing and facilities. Eleven ORNL staff members participated as guest lecturers, providing a diverse offering of lecture topics, including requirements gathering, design, operations, energy efficiency, reliability, security, network systems, and systems management. Jim Rogers, Kathlyn Boudwin, and Stephen McNally served as OLCF guest lecturers for the course.

1.4.8.1 Tiny Titan

Tiny Titan continued to be an effective outreach tool for the OLCF (Figures 1.7 and 1.8). 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.

Tiny Titan has represented the OLCF at regional schools, the DOE National Science Bowl, and the International Supercomputing Conference, and in local and national news outlets. 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. Robert French (left) and Adam Simpson (right) take visitors through the steps of operating Tiny Titan at an exhibit at the American Museum of Science and Energy.



Figure 1.8. Robert French (center) discusses Tiny Titan with Science Bowl participants with Carolyn Lauzon (right) and OLCF's Adam Simpson (left).

1.4.9 Outreach

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 72 highlights—including science, technology, and “people features”—and for more than 151 total outreach products in 2015. In addition, throughout the year, the OLCF provides tours to groups of visitors who range from middle-school students through senior-level government officials. The center conducted tours for 241 groups in 2015.

The team adopted a theme of “Achieving great science on Titan” for outreach efforts in 2015 and carried it out throughout the year.

In all, the Outreach team produced a total of 23 science highlights in 2015 (up from 15 in 2014). Those highlights touched on a variety of science domains, from geology ([Titan Takes on the Big One](#)) to biology ([Researchers Build Bacteria's Photosynthetic Engine](#)); from materials ([Researchers Get Warmer in Understanding High-Temperature Superconductors](#)) to microbes ([Researchers Dig Up Biological Data from Microbes in Mines](#)). In addition, the team completed 19 technology stories and 30 people features.

In 2015, the Outreach team undertook an initiative to be more intentional and proactive in terms of its amplification methods via social media and through a number of electronic dissemination outlets. Further, the team was more methodical in terms of charting these successes. The net result was that Outreach science highlights were picked up by media outlets (science journals, trade publications and, in a few cases, the regular press) a total of 298 times in 2015—including a CNN feature story on Titan.⁹ By comparison, in 2014, the Outreach team tracked a total of 55 media pick-ups.

OLCF Outreach science highlights were published by *Science Daily*, *Phys.Org*, *Newswise*, *R&D Magazine*, *InsideHPC*, and *HPCWire*, and other outlets. Outreach highlights were also published on the DOE Office of Science homepage several times throughout the year.

The story of science on Titan was also the theme for the 2014–15 OLCF annual report, titled “Frontiers of Discovery.” The report, produced in 2015, carried the theme of achieving great science

⁹ <http://money.cnn.com/video/technology/2015/09/10/elon-evolution-titan-americas-most-powerful-supercomputer.cnnmoney/>

throughout the volume. In addition to presenting accomplishments from the previous year, the annual report pointed to the future with an update on the coming of Summit.

The OLCF Outreach Team received an “Award of Quality” in the 40th annual Public Relations Society of America Volunteer Chapter awards competition for the integrated communications category for the media campaign to introduce Summit. That campaign, titled “Summit: Scale New Heights. Discover New Solutions” highlighted the comprehensive communications plan, including a website, news release, fact sheet and video.

Also in 2015, OLCF graphic designer Jason Smith was recognized by the American Inhouse Design Awards competition for the design of the 2013–2014 OLCF Annual Report. That report received a “Certificate of Excellence.”

The OLCF website received 332,321 page views and a mean survey rating of 4.4/5.0 in 2015. The most visited pages in 2015 were the Titan resource overview page and the Titan User Guide. Overall, the OLCF User Guides were the most highly rated aspect of the website with a mean rating of 4.5/5.0, and 94% of users indicated that they were satisfied or highly satisfied with the User Guides. The User Guides received 24,210 page views and 19,637 unique page views in 2015. The tutorials were once again a popular resource for OLCF users; two tutorials were among the top 10 visited pages with a combined total of 11,400 page views. The OLCF launched a new Industrial Partnerships Program webpage located at <https://www.olcf.ornl.gov/accel>. This new page provides companies with information on accessing the OLCF resources and provides a list of current and past OLCF industrial projects.

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. Moreover, because applications have much longer lifespans than do any computer architecture, applications need to be developed for changing architectures. Further, many of the OLCF’s principal investigators have allocations at multiple computing facilities, and having portable applications greatly facilitates their science campaigns.

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 SciComp group; Katie Antypas, Scientific Computing and Data Services Department Head at NERSC; and Timothy Williams, ALCF Principal Project Specialist, are working on an 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 address only the next-generation systems that will be coming to ALCF, NERSC, and OLCF. Using appropriate abstractions to obtain portability and performance on these pre-exascale systems also provides a path to continue toward exascale.

1.5.2 Application Readiness and Early Science

The OLCF’s CAAR is a partnership of the SciComp group, scientific application teams, vendor partners, and tools developers with the goal of readying a set of applications for the Summit architecture. The suite of CAAR applications cover a broad range of scientific disciplines and employ a range of programming models and software designs. In addition to developing highly efficient implementations for Summit, the teams are working toward performance portability across different architectures to serve the user base. The applications that are part of the CAAR program are summarized in Table 1.7. More details can be found on the CAAR webpage at <https://www.olcf.ornl.gov/caar>.

Table 1.7. Applications in the Center for Accelerated Application Readiness

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

An example of the performance improvements obtained in the CAAR program is the RAPTOR application, which allows for combustion simulations in complex geometries, using a finite-volume scheme with so-called body-fitted coordinates. The code includes five main computational kernels that are being ported to a model library to be used in the main integrator code. Special attention is given to developing a portable framework using the KOKKOS abstraction libraries. The kernels are being addressed in order of their computational intensity. Refactoring of the computationally most demanding thermophysics kernel resulted in a reduction of the contribution to the total runtime from the original 50% to 6%. Refactoring is now under way for the second most computationally intensive kernel, which computes the turbulence and scalar sub-grid scale mixing model.

1.5.3 CSEEN Postdoctoral Program

DOE recognizes the need to train and retain computational scientists in a broad range of disciplines that support DOE’s and the nation’s critical mission needs to maintain the US competitive advantage in high performance and data-intensive scientific computing. In light of the ever-increasing capability of high-end computer architectures, there is a continued and increased need to ensure a well-trained computational science workforce in academia and industry and at the national laboratories. In recognition of this need, DOE proposed that ASCR establish a postdoctoral training program at its user facilities OLCF, ALCF, and NERSC for future Computational Scientists for Energy, the Environment, and National Security (CSEEN). The objectives of this program are (1) to help ensure an adequate supply of

scientists and engineers appropriately trained to meet national workforce needs, including those of DOE, for high-end computational science and engineering, with skills relevant to both exascale and data-intensive computing; (2) to make ASCR facilities available, through limited-term appointments, for applied work on authentic problems with highly productive work teams and increasingly cross-disciplinary training, and (3) to raise the visibility of careers in computational science and engineering to build the next generation of leaders in computational science.

The OLCF CSEEN Postdoctoral Associates program seeks to provide opportunities to bridge the experience gap between the need to address domain science challenges and the need for the development of high-performance software development expertise. One of the focus areas is to provide the skills required to port, develop, and use software suites on the leadership computing resources at the OLCF. The software development activities will take place in association with a project in CAAR, but in the context of a challenging scientific research campaign that will allow scientific breakthroughs through computing and provide ample opportunity to publish in the domain scientific literature. This approach will ensure that the Postdoctoral Associates continue to build their reputations in their chosen science communities. Participants in the OLCF Postdoctoral Associates program are encouraged to attend tutorials, training workshops, and training courses in selected computer science topics. One of the most important outcomes for the Postdoctoral Associates is the opportunity to publish and present research accomplishments.

During 2015, the OLCF attracted the following scientists to the Postdoctoral Associates program.

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

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

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

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

Business Results

HIGH PERFORMANCE COMPUTING FACILITY 2015 OPERATIONAL ASSESSMENT OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2016

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 2015 reporting period includes full calendar year production periods for the HPC resources Cray XK7 (Titan), the Cray XC30 (Eos), and the Lustre file systems (Spider II). 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 2015, the OLCF delivered all of the compute hours committed to the three major allocation programs: INCITE, ALCC, and DD. OLCF computational and data resources are critical to scientific research of production simulations across many scientific domains.

2.1 BUSINESS RESULTS SUMMARY

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

2.2 CRAY XK7 (TITAN) RESOURCE SUMMARY

The OLCF upgraded the existing Cray Jaguar from a model XT5 to a model XK7, releasing it to production on May 31, 2013. The resulting system contains 18,688 NVIDIA K20X (Kepler) accelerators, in which each existing AMD Opteron connects to an NVIDIA Kepler to form a CPU-GPU pair. The completed XK7 system, 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 736 Intel Xeon E5-2670 compute nodes and 47.6 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 (SPIDER II) RESOURCE SUMMARY

In October 2013, the OLCF released Spider II, its next-generation Lustre parallel file system, to production. Spider II contains two instantiations of the /atlas file system, with an aggregate capacity of

more than 30 PB and block-level performance of more than 1.3 TB/second. The Spider II file system is the default high-performance file system for all compute resources.

2.5 DATA ANALYSIS AND VISUALIZATION CLUSTER (RHEA) RESOURCE SUMMARY

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

2.6 HIGH PERFORMANCE STORAGE SYSTEM 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 National Laboratories, Lawrence Livermore National Laboratory, Lawrence Berkeley National Laboratory, and ORNL. The ORNL HPSS instance is currently over 50 PB in size and provides up to 200 GB per second of read and write performance. The archive has taken in over 225 TB in a single day several times in the last year; the previous daily maximum was just over 150 TB.

The archive is built from hardware from Dell, Hewlett Packard, Brocade, NetApp, DataDirect Networks, and Oracle. An 18 PB disk cache allows burst rates into the archive at up to 200 GB per second; there is 26 GB per second of read/write bandwidth to the archive via 154 Oracle T10K series tape drives. There are 6 Oracle SL8500 tape libraries for tape archival storage that each contain 10,100 slots; the archive's maximum capacity is over 500 PB, using these libraries.

2.7 VISUALIZATION RESOURCE SUMMARY

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

2.8 OLCF COMPUTATIONAL AND DATA RESOURCE SUMMARY

The OLCF provided the Titan and Eos computational resources and the Spider II and HPSS data resources in 2015 (see Table 2.1). Supporting systems such as EVEREST, Rhea, and data transfer nodes were also offered. Metrics for these supporting systems are not provided.

2.8.1 OLCF HPC Resource Production Schedule

The OLCF production computational systems entered production according to the schedule in Table 2.2. This includes historical data associated with the Cray XT5, the very small overlap in December

2011 beginning with the introduction of the Cray XK6, and the series of Cray XK systems available in 2012 and 2013.

Table 2.1. OLCF production computer systems, 2015

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 symmetric multiprocessing (SMP) CPU + 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	736	2×8-core SMP	64 GB DDR3—1600 per node; 47,104 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
Spider II	Lustre parallel file system	October 3, 2013	Null	Delivered as two separate file systems, /atlas0 and /atlas1. 30+ PB capacity
Eos	Cray XC30	October 3, 2013	Null	Production with 736 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).

Table 2.3. OLCF business results summary for HPC systems

Measurement		2014 target	2014 actual	2015 target	2015 actual
Cray XK7 (Titan)	Scheduled availability	90%	99.59%	95%	99.41%
	Overall availability	85%	95.80%	90%	97.01%
	MTTI (hours)	NAM	310.83	NAM	326.86
	MTTF (hours)	NAM	1,246.54	NAM	1,088.67
	Total usage	NAM	89.63%	NAM	90%
	Core-hours used*	NAM	4,217,292,935	NAM	4,287,795,259
	Core-hours available	NAM	4,705,171,200	NAM	4,764,524,288
	Capability usage				
	INCITE projects	NAM	69.44%	NAM	71.31%
	All projects	35%	62.58%	35%	64.19%

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		2014 target	2014 actual	2015 target	2015 actual
Cray XC30 (Eos)	Scheduled availability	NAM	99.78%	NAM	99.67%
	Overall availability	NAM	97.06%	NAM	97.90%
	MTTI (hours)	NAM	340.09	NAM	476.47
	MTTF (hours)	NAM	1,748.20	NAM	2,182.90

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 2.5. OLCF business results summary for HPSS

Measurement		2014 target	2014 actual	2015 target	2015 actual
HPSS	Scheduled availability	95%	99.83%	95%	99.94%
	Overall availability	90%	98.56%	90%	98.00%
	MTTI (hours)	NAM	297.71	NAM	343.40
	MTTF (hours)	NAM	546.56	NAM	1,459.13

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 system—Spider II

	Measurement	2014 target	2014 actual	2015 target	2015 actual
atlas1	Scheduled availability	95%	99.73%	95%	99.77%
	Overall availability	90%	99.15%	90%	98.90%
	MTTI (hours)	NAM	377.62	NAM	618.84
	MTTF (hours)	NAM	624.02	NAM	1,456.69
atlas2	Scheduled availability	95%	99.50%	95%	99.76%
	Overall availability	90%	98.59%	90%	98.88%
	MTTI (hours)	NAM	411.26	NAM	481.23
	MTTF (hours)	NAM	792.41	NAM	1,092.44

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.

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

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

The Spider II, Titan, and Eos systems all celebrated their 2 year production anniversaries in 2015. The reported results for each system measure are for the full 2015 calendar year and intentionally do not reflect the partial results to their respective production anniversaries. In all cases, the OLCF result exceeded the most stringent year 3 targets for the accompanying metric.

Because an outage that could define the SA, OA, mean time to interrupt (MTTI), or mean time to failure (MTTF) may occur outside the reporting period, the data reflected here artificially assume calculation boundaries of January 1, 2015 0:00 and January 1, 2016 0:00.

2.9 RESOURCE AVAILABILITY

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

2.9.1 Scheduled Availability

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

As shown in Table 2.7, the OLCF has exceeded the SA targets for the facility's computational resources for 2014 and 2015.

Table 2.7. OLCF business results summary: Scheduled availability

	System	2014 target	2014 actual	2015 target	2015 actual
Scheduled Availability	Cray XK7	90%	99.59%	95%	99.41%
	Cray XC30	NAM	99.78%	NAM	99.67%
	HPSS	95%	99.83%	95%	99.94%
	/atlas0	95%	99.73%	95%	99.77%
	/atlas1	95%	99.50%	95%	99.76%

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 OLCF HPC Operations group. Typical PM activities include software updates, application of field notices, and hardware maintenance to replace failed components. Without concurrence, the systems are allowed to continue in operation.

In 2015, OLCF staff executed scheduled maintenance on the Cray XK7 a total of 18 times, associated with hardware maintenance, SMW and CLE firmware/software upgrades, field notices, Lustre software stack testing, and integration with the Spider II file system. Seven unscheduled outages were reported in 2015, including external power quality events and a small number of hardware or software failures that could not be recovered. Similarly, OLCF performed scheduled maintenance on Eos 14 times in 2015, with three unscheduled outages.

2.9.2 Overall Availability

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

As shown in Table 2.8, the OLCF has exceeded the OA targets for the facility's computational resources for 2014 and 2015.

Table 2.8. OLCF business results summary: Overall availability

	System	2014 target	2014 actual	2015 target	2015 actual
Overall Availability	Cray XK7	85%	95.80%	90%	97.01%
	Cray XC30	NAM	97.06%	NAM	97.90%
	HPSS	90%	98.56%	90%	98.00%
	/atlas0	90%	99.15%	90%	98.90%
	/atlas1	90%	98.59%	90%	98.88%

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

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

The MTTI summary is shown in Table 2.9.

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

	System	2014 target	2014 actual	2015 target	2015 actual
MTTI (hours)	Cray XK7	NAM	310.83	NAM	326.86
	Cray XC30	NAM	340.09	NAM	476.47
	HPSS	NAM	297.71	NAM	343.40
	/atlas0	NAM	377.62	NAM	618.84
	/atlas1	NAM	411.26	NAM	481.23

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

2.9.4 Mean Time to Failure

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

The MTTF summary is shown in Table 2.10.

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

	System	2014 target	2014 actual	2014 target	2014 actual
MTTF (hours)	Cray XK7	NAM	1,246.54	NAM	1,088.67
	Cray XC30	NAM	1,748.20	NAM	2,182.90
	HPSS	NAM	546.56	NAM	1,459.13
	/atlas0	NAM	624.02	NAM	1,456.69
	/atlas1	NAM	792.41	NAM	1,092.44

NAM = Not a metric. No defined metric nor target exists for this system. Data provided as reference only.

2.10 RESOURCE UTILIZATION

2015 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 streaming multiprocessor (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. 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, 2015, 4,287,795,259 Titan core-hours were used outside of outage periods from an available 4,764,524,288 Titan core-hours. The total system utilization for the Cray XK7 was 90%.

2.10.2 Total System Utilization

2015 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.

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

The measurement period is for 2015, 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 2015 was 90%.

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 2015.




Figure 2.1. 2015 XK7 resource utilization—Titan core-hours by program.

2.10.2.1 Performance of the Allocated Programs

All allocation programs, including INCITE, ALCC, and DD, are aggressively monitored to ensure that projects within these allocation groups maintain appropriate consumption rates. The 2015 INCITE allocation program was the largest program in 2015, with a commitment for 2.25B Titan core-hours. The consumption of these allocation programs is shown in Table 2.11.

Non-renewed INCITE projects from 2014 continued running through January 2015. 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. Similarly, 2015 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, 2016. ALCC projects from the 2014 allocation period (ending June 30, 2015) were also granted extensions where appropriate.

Table 2.11. The 2015 allocated program performance on the OLCF resources

Program	Allocation	Hours consumed	Percent of total
INCITE ^a	2,250,000,000	2,660,189,558	61.24%
Titan		2,577,269,402	
Eos		82,920,156	
ALCC	Not applicable	1,226,357,633	28.23%
ALCC_2015 ^b	1,153,400,000	794,336,907	
ALCC_2016 ^c	1,074,300,000	432,020,726	
DD	Not applicable	407,500,170	9.38%
NOAA ¹⁰	87,500,000	50,017,705	1.15%
Total		4,344,065,066	100.00%

^a Includes 13th month usage from January 2016

^b Number of hours used within CY 2015 for the ALCC_2015 program

^c Number of hours used within CY 2015 for the ALCC_2016 program

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. The metric for CY 2015 was 35%, and this metric will hold until Titan is retired. 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 wall clock times and increased priority.

The OLCF continues to exceed expectations for capability usage of its HPC resources (Table 2.12). Keys to the growth of capability usage include the liaison role provided by 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.

Table 2.12. OLCF capability usage on the Cray XK7 system

	Leadership usage	CY 2014 target	CY 2014 actual	CY 2015 target	CY 2015 actual
Cray XK7	INCITE	NAM	69.44%	NAM	71.31%
	ALCC	NAM	52.53%	NAM	60.34%
	All Projects	30/35%	62.58%	35%	64.19%

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, 64.19%, was well above the 2015 target of 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

¹⁰ NOAA = National Oceanic and Atmospheric Administration; period of performance reported: January 1, 2015–December 31, 2015.

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) of the nodes through aging boosts. The OLCF assesses job data in 10% “bins” to understand the job size distribution. Further, by assessing the aggregate bins 20–60%, and >60%, the OLCF can assess the impact of queue policy on delivered node-hours.

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




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

2.12 GPU USAGE

The heterogeneous architecture of Titan provides a key capability to users, allowing them to exploit a 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 has chosen to measure GPU enablement with Cray’s Resource Utilization Reporting (RUR) tool, explained in more detail in Section 2.12.1.

Other methods for measuring GPU usage (Automatic Library Tracking Database [ALTD] and OVIS) were evaluated. In those cases, it became clear that RUR was the utility of preference for the OLCF. ALTD required significant diligence to maintain the library index for GPU-enabled libraries. Given the amount of effort required, and that ALTD provides a very rough estimate of GPU usage, it was deemed non-essential for understanding how the GPUs were being used. Rather, ALTD provides valuable information as to which libraries are being linked during compilation. Simply measuring linked libraries is merely a loose correlation to the actual use of the GPUs. Given this, ALTD continues to be supported and maintained within the OLCF, but not for the purpose of understanding how the GPUs are used. The OVIS utility was evaluated in CY 2015. While it can provide GPU usage numbers, it is an unnecessary piece of software to run on the compute nodes given the advances in the native RUR and NVIDIA Management Library tools. The OLCF decided that running yet another monitoring tool to view data that is already available by other means was not an efficient operational practice.

Work to refine GPU statistics reporting progressed through 2015 and continues today. In 2015, the need to reduce time-to-delivery for new GPU usage reports drove OLCF developers to design and deploy an all-new data warehouse centered around industry-standard Online Analytical Processing (OLAP) technology. Paired with lightweight commercial GUI software and open-source OLAP server software, the new warehouse allows OLCF staff to quickly search and aggregate GPU statistics by award/allocation program, science category, research area, and individual project. This novel, innovative technique for

handling business intelligence reporting has proved so flexible that many preexisting reports are being ported to the new OLAP system. The GPU usage statistics reported in Sections 2.12.1 and 2.12.2 were provided by this new infrastructure.

2.12.1 Measuring GPU-Enablement and GPU-Activity

To accurately measure and report GPU usage, the OLCF defines two primary metrics: GPU-enablement and GPU-activity. GPU-enablement, recorded on a per job basis, is defined as the total batch runtime for any job that used a GPU at any point between job start and end. GPU-enablement is a binary metric and is primarily used to delineate between jobs that used the GPUs (GPU-enabled jobs) and jobs that did not. GPU-activity is defined as the runtime in which the GPU was active within GPU-enabled jobs. GPU-activity shows the usage of GPUs on Titan. Specifically, GPU-activity is calculated as follows:

$$\%gpu_activity = \frac{\sum_{i=1}^n \left(\frac{gpu_seconds_i}{gpu_contexts_i} \right)}{\sum_{j=1}^p (end_time_j - start_time_j)} \quad (6)$$

where

<i>%gpu_activity</i>	Percentage of batch job runtime during which a GPU was in use (unit: none).
<i>j</i>	Iterator for batch jobs ended within a queried timeframe that employed a GPU via associated apruns (unit: none).
<i>p</i>	The number of batch jobs ended within a queried timeframe that employed a GPU via associated apruns (unit: none).
<i>i</i>	Iterator for apruns associated with batch jobs in the queried timeframe.
<i>n</i>	The number of apruns associated with batch jobs in the queried timeframe.
<i>gpu_seconds</i>	Integer sum of GPU seconds across all GPU contexts within an aprun in which the GPU was active. Obtained directly from RUR log records. (unit: context-second).
<i>gpu_contexts</i>	Number of GPU contexts within an aprun. A GPU context is the execution of an application process with a unique process identifier. Typically equals 1. (unit: context).
<i>end_time</i>	Wall-time at which a batch job finished. (unit: second).
<i>start_time</i>	Wall-time at which a batch job started. (unit: second).

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 CY 2015, approximately 53% of all delivered compute time on Titan was by GPU-enabled applications. Table 2.13 shows the GPU-enabled and CPU-only hours used and percentage breakdowns of each of the three primary allocation programs at the OLCF (INCITE, ALCC, and DD). As shown, the INCITE program makes the most use of GPU-enabled time on Titan. Compared with the ALCC and DD allocation programs, the INCITE program shows roughly a 2× increase in the number of GPU-enabled and GPU-active hours consumed. Of the hours consumed, the INCITE program ran at a 3:2 GPU-enabled to CPU-only computing ratio. The ALCC program virtually splits time using both the GPUs and CPUs, with a slight edge toward CPU computing. The DD program showed the least preference for GPU-enabled computing with an almost 3:1 CPU-to-GPU computing ratio. In general, these usage patterns match what is expected with each of the allocation programs. The INCITE computational readiness review criteria factor in the proposed use of GPUs when allocating time and projects. The ALCC program does not require computational readiness reviews for GPU usage, and the DD program supports projects that may be in the beginning phases of porting code to GPUs. For most of the program year, GPU-enabled INCITE applications were consistently responsible for over 50% of the delivered hours to those projects and averaged 61.81% for 2015. Figure 2.3 shows the percentage of GPU-enabled compute time by month and Figure 2.4 shows GPU-activity measured in hours for the INCITE, ALCC, and DD programs.

Table 2.13. 2015 GPU-enabled and CPU-only usage by program

Program	Percentage	Hours
INCITE (GPU-enabled)	61.81	1,539,060,854
INCITE (CPU-only)	38.19	950,978,009
ALCC (GPU-enabled)	46.88	574,884,972
ALCC (CPU-only)	53.12	651,472,692
DD (GPU-enabled)	27.21	110,882,123
DD (CPU-only)	72.79	296,618,016




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




Figure 2.4. GPU-activity for the DD, ALCC, and INCITE user programs.

2.12.2 Case Study: Transitioning from Discretionary Program to INCITE

Principal investigator Vittorio Michelassi conducted direct numerical simulation (DNS) of turbines using HiPSTAR (High-Performance Solver for Turbulence and Aeroacoustics Research) under INCITE project TUR105, a continuation of the ARD106 DD project from 2014. HiPSTAR was developed to conduct cutting-edge DNS: the solution of fully nonlinear, time-dependent, 3-D Navier-Stokes equations on traditional HPC systems utilizing hybrid OpenMP/MPI parallelism. To take full advantage of Titan, GPU parallelism using OpenACC was implemented during the year. Initial results have shown that HiPSTAR can yield strong scaling or weak scaling depending on the complexity of the simulation. The project has used HiPSTAR on more than 8,000 GPUs on Titan for its production simulations. In 2015, there was a substantial increase in the use of GPUs during the year, primarily enabled by the OpenACC implementation. Figure 2.5 shows the percentage of compute time by month for TUR105.




Figure 2.5. Percentage of GPU-enabled compute time by month for TUR105.

2.13 PARALLEL DU TOOL SUITE

On file systems the size of Spider II, the standard “du” utility could take hours—or possibly days—to return an answer. The parallel du tool suite, developed at the OLCF, is a collection of tools that work together to provide a useful replacement for the standard *nix ‘du’ utility. The suite consists of three parts: lester, load_du_data, and lustredu. Lester runs on the Lustre metadata servers to read and store information for the associated file system metadata. The load_du_data tool reads the output from lester, combines it with file size information taken from the Lustre object storage servers, and outputs the results to a backend database. The lustredu program is the user space tool that enables users to perform disk usage queries on specific files and directories of the Spider II file system.

Because the lustredu tool queries a database rather than the metadata servers, it can provide an answer in seconds and does not impose the kinds of loads on the file system that the standard du utility would impose. This tool suite enables OLCF HPC Operations staff and users to perform useful file system queries without negatively impacting the shared resources. This same tool suite has been implemented for the National Oceanic and Atmospheric Administration work-for-others customer of the NCCS facility.

2.14 DDNTOOL

The DDNTool, developed by the OLCF, gathers I/O statistics supporting analysis of the low-level I/O performance of the disk controllers, including bandwidth, I/O operations per second, object storage target (OST) usage data, detection of degraded OST pools, and other measures.¹¹ DDNTool provides an aggregation mechanism that allows OLCF staff to more efficiently monitor the health and performance of a file system containing more than 20,000 disk drives and 72 controllers. DDNTool’s primary author, Ross Miller, developed the tool in response to a need to optimize the seemingly serial approach to leadership class file system resource management. Although commercial tools were provided from DDN, they were not intended for use at the scale at which OLCF operates. DDNTool has allowed OLCF staff to centralize the serial operations of monitoring and management, providing a robust utility that works at scale. By designing the tool to operate with standard log aggregation utilities, such as Splunk, staff can gain new insights into the management and operation of such a large-scale file system. Figure 2.6 shows a visualization provided through a Splunk interface using data collected from DDNTool. This data collected from the tool was presented as part of a paper titled “Comparative I/O Workload Characterization of Two Leadership Class Storage Clusters” in *Petascale Data Storage Workshop*, in conjunction with SC’15 (authors are Raghul Gunasekaran, Sarp Oral, Jason Hill, Ross Miller, Feiyi Wang, and Dustin Leverman). The tool is also being evaluated by a group at Los Alamos National Laboratory, and the OLCF has recently received requests from Pacific Northwest National Laboratory and Sandia National Laboratories to evaluate it. Working with the ORNL technology transfer division and DDN, OLCF has successfully made DDNTool available to the computing community by making it open-source.




Figure 2.6. Splunk-provided sample of performance data obtained from DDNTool.

2.15 BALANCED PLACEMENT I/O

The balanced placement I/O (BPIO) library presents a method to mitigate resource contention with a topology-aware strategy. OLCF staff have integrated BPIO with ADIOS, a popular and portable middleware solution for scientific applications. Integrating BPIO with ADIOS allows ADIOS-enabled scientific applications to obtain immediate benefits from BPIO performance improvements without requiring any further code changes. The OLCF is currently configuring and tuning the BPIO and ADIOS integration using a large number of scientific applications, including XGC. As part of this effort, the OLCF and XGC teams have been working together to improve the integrated stack performance. The analysis shows an improvement in I/O performance of up to 47% for POSIX I/O and 25% for MPI-IO interfaces with ADIOS/BPIO. A poster titled “Improving Large-scale Application Performance with ADIOS and BPIO” won the best poster award at the 2015 Smoky Mountains Computational Sciences and

¹¹ <https://www.olcf.ornl.gov/2015/09/09/new-ddntool-streamlines-olcf-file-system-monitoring/>

Engineering Conference (authors are Sarah Neuwirth, Sarp Oral, Feiyi Wang, Qing Liu, and Sudharshan Vazhkudai).

2.16 PARALLEL DATA TRANSFER

FCP is an MPI-based tool for highly scalable data transfers in a clustered environment. FCP has gone through a comprehensive set of test cases from the OLCF User Assistance and Outreach and HPC Operations groups. It was used to move 5 PB of cosmology data between the Spider II file systems (/atlas1 to /atlas2). This tool is included in the open source suite of tools located at <http://www.github.com/olcf/pcircle>. To date, FCP has been shared with Fermilab, HZDR (Helmholtz Association of German Research Centre), and the Department of Defense High Performance Computing Modernization Program. It is expected that FCP will be critical to the need to transition large datasets from the current Spider II file system to the next-generation Spider III in 2017.

Strategic Results

HIGH PERFORMANCE COMPUTING FACILITY 2015 OPERATIONAL ASSESSMENT OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2016

3. STRATEGIC RESULTS

CHARGE QUESTION 3: Is the facility enabling scientific achievements consistent with the 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 are advancing DOE's mission to ensure America's security and prosperity by addressing its energy, environmental, and nuclear challenges through transformative science and technology solutions. The 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 and 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

2015 Operational Assessment Guidance

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

3.1.1 OLCF Publications Report

In 2015, a total of 341 refereed publications—resulting from the use of OLCF resources—were published, as identified in the collection completed on March 3, 2016.¹² A listing of these publications is available on the OLCF website.¹³ Within this 2015 count, users and OLCF staff jointly authored 31 of these publications, users authored 279 without OLCF co-authorship, and OLCF staff published 31 without user co-authorship. In the 2014 OLCF Operational Assessment Report (OAR), 229 publications were reported.

Sponsor guidance allows the reporting of “publications in press or accepted.” However, the OLCF implementation of the guidance is that only publications appearing in print in the year under review (e.g., 2015) are eligible for tabulation in the current report. The OLCF continues to search for publications after the OAR is communicated to DOE. The number of publications reported within previous OARs will be reevaluated in the light of this implementation, with updates reported annually in the OLCF OAR. Table 3.1 reports the updated, verified, and validated publication count for the period of 2012–2015.

**Table 3.1. Summary of unique OLCF publications over the period
2012–2015**

Year	Publications reported in previous OAR	Unique, confirmed OLCF publications
2015	—	341
2014	229	292
2013	262	339
2012	321	308

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 2015, 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 2015, including two in *Nature*, one in *Nature Materials*, one in *Nature Chemistry*, three in *Nature Physics*, two in *Nature Climate Change*, two in *Nano Letters*, two in *ACS Nano*, three in *Journal of the American Chemical Society*, one in *Nature Geoscience*, six in *Nature Communications*, three in *Astrophysical Journal Supplement Series*, and ten in *Physical Review Letters*. Altogether in 2015, OLCF users published 30 papers in journals with a journal impact factor (JIF) greater than 8 and 44 papers in journals with a JIF greater than 7. Also of interest is the large breadth of journal and conference titles in which OLCF publications appear in 2015. More than 150 unique journal and conference titles are represented in this list.

¹² In this document, “year” refers to the calendar year unless it carries the prefix “FY” indicating the fiscal year.

¹³ <https://www.olcf.ornl.gov/leadership-science/publications/>

3.2.1 Flowing Toward Red Blood Cell Breakthroughs: George Karniadakis, Brown University, INCITE

Objective: To develop high-fidelity blood flow simulations in order to develop better drug delivery methods and predictors for tumor cell formation and sickle cell anemia (SCA). A team led by George Karniadakis is using dissipative particle dynamics to simulate microfluidic devices capable of separating cancer cells from blood and to study how SCA cells interact with healthy blood cells.

Impact: Simulating blood dynamics is not as simple as simulating fluid dynamics. Researchers must take into account the multiscale nature of blood—a fluid containing many individual particles within it. The Karniadakis team uses multiscale simulations at the OLCF to focus on two major medical issues—understanding SCA and designing microfluidic devices. Around 8% of the African-American population carries the SCA trait, and more than 180,000 babies are born with the disorder every year. Despite its prevalence, little is known about how this red blood cell related disorder interacts with human blood vessels.

Simulations of microfluidic devices, or devices capable of manipulating extremely small amounts of fluids, could help medical professionals more quickly diagnose and treat tumors. Microfluidic devices would allow a doctor to take a very small sample of blood and quickly identify whether a person had, for example, a malignant tumor. These devices have “lab on a chip” properties that could help doctors test for a broad range of diseases in the least invasive way possible.

Accomplishments: The Karniadakis team was able to run some of its largest dissipative particle dynamics simulations for studying SCA as part of its INCITE work. The team used its multiscale simulations to observe rigid SCA cells interacting with healthy cells and blood vessels. In addition, the team ran successful simulations of microfluidic devices—a stretch goal—using Titan and computational expertise to make the process more efficient (Figure 3.1).

Post-doctoral researcher and project collaborator Yu-Hang Tang of Brown University, and his teammates, exploited Titan’s GPU accelerators and developed uDeviceX, a GPU-accelerated particle solver—an important part of the team’s code that helps plot individual particles in the simulation. Tang’s new solver showed a 45-fold decrease in time-to-solution compared with competing state-of-the-art methods.

Tang’s extensive work with GPUs led to the team’s newest computational tool to help with the integration of the various codes—the multiscale universal interface (MUI). MUI allows the team to quickly compose its multiple applications addressing specific phenomena into one larger multiscale code, significantly reducing the computational costs of running simulations by focusing on the strengths of different hardware configurations.

OLCF contributions: OLCF computational scientist Wayne Joubert helped the Karniadakis team scale its MUI code to make the most efficient use of Titan’s large node count. In addition, the team generated 100s of gigabytes of data during its production runs. It relied on the OLCF’s High-Performance Storage System and the FTP file transfer service to both store its data and efficiently move it to Brown University. During its INCITE project, the Karniadakis team performed 41% of its runs using less than 20% of Titan, 28% using 20–60% of Titan, and the final 31% using more than 60% of the system.




Figure 3.1. Red blood cells (red) and circulating tumor cells (green) traveling through a microfluidic cell sorting device as simulated by uDeviceX. Image credit: Yu-Hang Tang, Brown University

Related publication:

D. Rossinelli, Y.-H. Tang, K. Lykov, D. Alexeev, M. Bernaschi, P. Hadjidoukas, M. Bisson, W. Joubert, C. Conti, G. Karniadakis, M. Fatica, I. Pivkin, P. Koumoutsakos, “The in-silico lab-on-a-chip: Petascale and high-throughput simulations of microfluidics at cell resolution,” ACM 2015 Gordon Bell Award Finalist, in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis* (SC ’15), 2015. <http://dx.doi.org/10.1145/2807591.2807677>

Online story: “[Flowing Toward Red Blood Cell Breakthroughs](#),” *OLCF News* (October 13, 2015).

3.2.2 Researchers Mine Information from Next-Generation Subsurface Flow Simulations: James McClure, Virginia Polytechnic Institute and State University, INCITE

Objective: To improve the understanding of transport phenomena in multiphase systems, or systems in which different phases (solids, liquids, or gases) or different chemical makeups (e.g., oil and water) flow together. The team uses a combination of experiment and computation to improve models for capturing subsurface flow phenomena.

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

Accomplishments: The team’s simulations on Titan have elucidated the role played by interfaces—surface areas where two disparate substances meet in a two-phase system, such as where oil and water meet while flowing through a porous, spongy rock—in determining the behavior of these complex multiscale systems. Analyzing the simulation state in situ enables unprecedented insight into phase interactions that could not otherwise be measured or observed.

Working with the OLCF, James McClure has developed simulation tools that allow researchers to take the information from a 3-D micro-CT image and put it into simulated motion (Figure 3.2). With this approach, the team can study interfacial dynamics and other aspects of a system’s behavior quantitatively. The team can predict rocks’ properties—such as relative permeability—from first principles to measure the resistance to flow that controls the movement of fluids on a larger scale.

These results are vital for situations in which the mobility of trapped, nonwetting phases—in this case, oil or liquid carbon dioxide—is of essential importance, such as carbon sequestration, oil recovery, and contaminant transport. There is much societal benefit to understanding the conditions under which these nonwetting phases can be removed from the subsurface rock or placed to reside permanently in rock.

OLCF contributions: OLCF computational scientist Mark Berrill helped optimize codes so that when running their largest simulations with compute-intensive tasks on Titan’s GPUs, the team could use Titan’s CPUs to analyze the data and then send relevant data sets to the Rhea data-analysis cluster to visualize simulation results as they are generated. This novel approach has allowed the team to churn through hundreds




Figure 3.2. Simulations on Titan enable detailed tracking of two fluid phase systems during flow through porous media. In this case, the behavior of oil ganglia (yellow) can be observed as they are mobilized in experimentally imaged sandstone (solid part of the sandstone is shown in blue). *Image credit: James McClure, Virginia Tech*

of terabytes of data, providing continuous information about the simulation state. OLCF staffer Robert French also helped the team optimize GPU-aware MPI communications on Titan. In the team's INCITE project, 16% of its simulations used less than 20% of Titan and the other 84% used 20–60% of Titan.

Related publication:

W. G. Gray, A. L. Dye, J. E. McClure, L. J. Pyrak-Nolte, and C. T. Miller, “On the dynamics and kinematics of two-fluid-phase flow in porous media,” *Water Resources Research* (2015). DOI: 10.1002/2015WR016921

Online story: “[Researchers Mine Information from Next-Generation Subsurface Flow Simulations](#),” *OLCF News* (September 9, 2015).

3.2.3 Fusion Researchers Use Titan to Burst Helium Bubbles: Brian Wirth, Oak Ridge National Laboratory, ASCR Leadership Computing Challenge

Objective: This project aims to understand the effect of the rate at which helium implants in tungsten. As tungsten is the primary material used to build fusion reactor divertors—such as the divertor being constructed for the ITER experimental fusion reactor in France—understanding its interactions with helium in the plasma state is imperative. The team hopes to develop a database of helium bubble kinetics and validate its models with simulation.

Impact: To realize the promise of limitless, clean energy that fusion reactors offer, researchers must gain a greater understanding of how a plasma made of helium and its ions interacts with the surface of a fusion reactor. Currently, fusion reactor divertors are made of tungsten; and although tungsten is one of the toughest materials known, simulations show that it nonetheless degrades with exposure to high-temperature helium plasmas. By finding a way to mitigate this degradation, researchers can more efficiently plan for ITER operation.

A Los Alamos National Laboratory (LANL) -based team made up of Luis Sandoval, Danny Perez, Blas Uberuaga, and Arthur Voter is working to understand more fully how tungsten behaves when subjected to the harsh conditions inside a fusion reactor.

To study the plasma-material interface over long times, the team is using the Parallel Replica (ParRep) method, developed at LANL in the 1990s, with the support of DOE's Office of Basic Energy Sciences. The ParRep method aims to address the timescale limitation of molecular dynamics for systems that evolve through infrequent or rare state-to-state transitions.

Accomplishments: Using ParRep, Sandoval et al. were able to simulate many microseconds of time rather than pico- or nanoseconds (10^{-12} and 10^{-9} , respectively). This speedup allows researchers to more fully understand how helium molecules impinge, diffuse, and disrupt the divertor wall (Figure 3.3). In particular, the LANL team found that the evolution of helium bubbles, when simulated over these realistic time scales, is qualitatively different from what is observed when traditional molecular dynamics is used over shorter time scales.

A simulation using approximately 8,000 nodes on Titan can be done in 30 minutes using the LANL method, whereas it would have taken 7 months using traditional molecular dynamics simulations.




Figure 3.3. Bubble pressure as a function of the number of helium atoms added to a growing, sub-surface gas bubble in tungsten, revealing significant pressure drops as the bubble expands through loop punching before eventually bursting and releasing helium. Image credit: L. Sandoval et al. *Physical Review Letters* 114 (2015).

OLCF contributions: By gaining access to Titan, the team was able to run its large-scale ParRep dynamics simulations, allowing them to complete its simulations in minutes or hours—something that would be impossible with cluster computing or smaller supercomputers. In addition, the team benefitted greatly from Titan’s hybrid architecture—GPU acceleration helped the team obtain a two- to four-fold speedup in its simulations. Using Titan, 58% of the team’s runs used less than 20% of Titan, 41% used 20–60% of the machine, and 1% used more than 60%.

Related publication:

L. Sandoval et al., “Competing kinetics and He bubble morphology in W,” *Physical Review Letters* **114**, 105502 (2015).

Online story: “[Fusion Researchers Use Titan to Burst Helium Bubbles](#),” *OLCF News* (May 5, 2015).

3.2.4 Titan Takes on the Big One: Thomas Jordan, Southern California Earthquake Center, INCITE

Objective: To develop realistic simulations of earthquakes to better understand the potential seismic hazards from known faults and the impact of strong ground motions on urban areas. A team headed by Thomas Jordan uses the Southern California Earthquake Center (SCEC) CyberShake platform—a physics-based computational framework that integrates many features of an earthquake event—to calculate a probabilistic seismic hazard map for Southern California.

Impact: The last major earthquake—a quake of magnitude 7.5 or bigger—to hit San Francisco occurred in 1906, registering a 7.8 magnitude. The quake took 700 lives and caused \$400 million worth of damage, a figure equivalent to more than \$10 billion today. In the last century, researchers have collected data from smaller quakes throughout California, but such data fails to inform emergency officials and structural engineers for the next big one.

By creating the first physics-based seismic hazard models, Jordan’s team is providing a highly accurate tool for assessing the risks earthquakes pose to critical structures, such as large dams, nuclear power plants, and energy transportation networks.

Accomplishments: Using Titan, Jordan and his collaborators completed their highest-resolution CyberShake map for Southern California to date. With more realistic physics and an increase in the maximum frequency included in the simulation—from 0.5 to 1 hertz—the latest map gives engineers in Los Angeles and the surrounding region more detailed information on ground-motion risk for designing and building critical infrastructure.

The team achieved greater realism by leveraging Titan’s GPU accelerators via their anelastic wave propagation code AWP-ODC. The updated code performed six times faster than its CPU-only implementation, incorporating topography and additional geometrical and attenuation effects (near-fault plasticity, frequency-dependent attenuation, small-scale near-surface heterogeneities, near-surface nonlinearity, and fault roughness) for the first time (Figure 3.4).




Figure 3.4. The CyberShake seismic hazard map shows the magnitude, or level of shaking, for the Los Angeles region, defined by the amount of change of a surface or structure in a 2 second period, with a 2% probability of increasing within the next 50 years.

Image credit: Scott Callaghan, Kevin Milner, and Thomas Jordan, Southern California Earthquake Center

The enhancements pave the way for realistic simulations of higher-frequency earthquakes, which impact smaller structures like homes. The team has performed simulations using Earth models near 4 hertz and is working to push the maximum simulated frequency above 10 hertz.

OLCF contributions: Titan’s hybrid architecture played a pivotal role in this team’s accomplishments, providing the GPU power needed to reduce the necessary compute time by 20% compared with a CPU-only implementation. Sixty-six percent of the team’s runs used GPUs. Additionally, OLCF support staff helped develop and implement a workflow tool that facilitated high-throughput jobs. Seventy-three percent of the team’s runs used more than 60% of Titan, and 27 percent of the runs used less than 20% of Titan.

PhD dissertation:

J. R. Donovan, “Forecasting Directivity in Large Earthquakes in Terms of the Conditional Hypocenter Distribution,” PhD dissertation, University of Southern California, 2015.

Invited talk:

T. Jordan, “Societal impact of earthquake simulations at extreme scale,” SC15 Conference, Austin, Texas, 2015.

Online stories: [“OLCF User Earns NVIDIA Award for GPU-accelerated Earthquake Simulations,”](#) *OLCF News* (April 8, 2015).

[“Titan Takes on the Big One,”](#) *OLCF News* (November 10, 2015).

3.2.5 Researchers Build Bacteria’s Photosynthetic Engine: Klaus Schulten, University of Illinois at Urbana-Champaign, INCITE

Objective: To simulate a complete photosynthetic organelle, or chromatophore, to better understand its atomic-level processes and characterize how hundreds of proteins work together to convert solar energy to food.

Impact: Nearly all life on Earth depends on photosynthesis, the conversion of light energy to chemical energy. Purple bacteria are thought to be the first photosynthetic organisms, possessing light-harvesting organelles whose complexity—at 100 million atoms—is within the range of current leadership-class supercomputers. Klaus Schulten’s team targeted the chromatophore to study the fundamentals of photosynthesis at the atomic level, basic research that could lead to better solar energy technology. Furthermore, the project marks a shift in computational biophysics from analyzing the individual cell part (e.g., a single protein) to analyzing the specialized systems of the cell (e.g., hundreds of proteins working together to carry out an autonomous function), a significant step toward the long-term goal of simulating an entire living organism.

Accomplishments: Schulten’s team used experimental data gathered from atomic force microscopy and the molecular dynamics code NAMD to calculate the forces exerted by the spherical chromatophore’s millions of atoms (Figure 3.5). Synthesized on Titan, the final




Figure 3.5. Researchers used experimental data to create a chromatophore model that contained about 16,000 lipids and 101 proteins, including the five major types of proteins that contribute to the clockwork of processes resulting in the conversion of light energy to ATP. Image credit: Abhi Singharoy and Melih Sener, UIUC

chromatophore model, the largest biomolecular simulation ever, measured 70 nanometers across and contained about 16,000 lipids and 101 proteins, including the five major types of proteins that contribute to the conversion of light energy to adenosine triphosphate (ATP), the common fuel for cellular function across all branches of life.

With a stable model, the team was able to study the processes and properties of the chromatophore at timescales approaching 1 microsecond, enough time to observe key slower-moving steps of the energy-conversion system.

To capture microsecond-scale interactions, Schulten's team offloaded a subset of NAMD's calculation to Titan's GPUs, achieving a two- to three-fold speedup in time to solution.

OLCF contributions: Schulten's team used OLCF resources to manage and analyze the large volume of data produced by this project. Because of the simulation's size, transferring data to the University of Illinois at Urbana-Champaign (UIUC) proved impractical. Instead, the OLCF enabled Titan's GPUs to visualize the team's data remotely, giving UIUC's John Stone a near real-time "window" into calculations and eliminating the need to transfer files. Under its INCITE allocation, 7% of the team's simulation runs required less than 20% of Titan, and 91% of the project required between 20 and 60% of system resources. Two percent of the work required 60 to 100% of Titan.

Related publications:

M. Sener, J. Strümpfer, A. Singharoy, C. N. Hunter, and K. Schulten, "Overall energy conversion efficiency of a photosynthetic vesicle," (submitted and under review).

J. E. Stone, M. Sener, K. L. Vandivora, A. Barragan, A. Singharoy, I. Teo, J. V. Ribeiro, B. Isralewitz, B. Liu, B. C. Goh, J. C. Phillips, C. MacGregor-Chatwin, M. P. Johnson, L. F. Kourkoutis, C. N. Hunter, and K. Schulten, "Atomic detail visualization of photosynthetic membranes with GPU-accelerated ray tracing," *Parallel Computing* (in press, December 12, 2015). DOI:10.1016/j.parco.2015.10.015

D. E. Chandler, J. Strümpfer, M. Sener, S. Scheuring, and K. Schulten, "Light harvesting by lamellar chromatophores in *Rhodospirillum rubrum*," *Biophysical Journal* **106**(11), 2502–2510 (2014). DOI:10.1016/j.bpj.2014.04.030.

Online story: "[Researchers Build Bacteria's Photosynthetic Engine](#)," *OLCF News* (July 29, 2015).

3.3 DIRECTOR'S DISCRETIONARY PROGRAM

2015 Operational Assessment Guidance

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

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

The goals of the DD program are threefold:

- Enable users to 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 by enabling new and nontraditional research topics.

- Support research and development 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.

Research and development partnerships are those aligned with DOE and ORNL strategic agendas. They may be entirely new areas with respect to high-performance computing (HPC), or ones in need of nurturing. Example projects are those associated with the ORNL Laboratory Director's Research and Development program; programmatic science areas (fusion, materials, chemistry, climate, nuclear physics, nuclear engineering, and bioenergy science and technology); and key academic partnerships (e.g., the University of Tennessee–ORNL Joint Institute for Computational Sciences). Examples of strategic partners in the DD program include the Consortium for Advanced Simulation of Light Water Reactors; Critical Materials Institute hub led by Ames National Laboratory; Accelerated Climate Model for Energy program; Center for Nanophase Materials Sciences; and large experimental facilities such as the Spallation Neutron Source (SNS) and the ATLAS (DOE Office of High Energy Physics) and ALICE (DOE Nuclear Physics program) experiments at CERN. A science achievement highlight from the SNS-OLCF user collaboration is summarized in Section 3.3.1. Also included in this broad category are projects that come to the OLCF through the Accelerating Competitiveness through Computational Excellence (ACCEL) Industrial HPC Partnerships Program, providing opportunities for industrial researchers to access the leadership systems to carry out research that would not otherwise be possible. See Section 3.4 for more information about this program.

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 co-design 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 awards can range from tens of thousands to 12 million hours or more.

In 2015, the OLCF DD program participants used 10% of the total allocable resource planned for these DD program goals, consuming 407,500,170 Titan core hours. A full listing of the enabled OLCF DD projects can be found in Appendix E.

3.3.1 Mechanical Properties of Nanoscopic Lipid Domains, Director's Discretionary

In an effort to better understand the biological processes that govern lipid raft formation (processes with broad implications for biological membrane research ranging from how cells regulate proteins to how HIV invades healthy human cells) a team led by ORNL researcher Xiaolin Cheng is comparing multi-million atom molecular dynamics simulations performed on Titan on time scales up to 1 microsecond with neutron spin echo experiments at the SNS. The combination of neutron research at SNS

and HPC on Titan offer the team leading-edge research tools to address questions that have stymied others. Having access to SNS has allowed the collaborative team led by John Katsaras to perform first-of-their-kind neutron spin-echo experiments to observe the bending properties of the lipid patches that populate the membrane. Cheng's team was able to successfully verify its experimental data with large-scale, all-atom molecular simulations (Figure 3.6).

Using Titan's hybrid architecture, the team needed 2 months of time to produce its data, a task that would have taken more than three times longer on a traditional CPU-only machine. This work was supported by the Scientific User Facilities Division of the DOE Office of Basic Energy Sciences.

Related publication:

J. D. Nickels, X. Cheng, B. Mostofian, C. Stanley, B. Lindner, F. A. Heberle, S. Perticaroli, M. Feygenson, T. Egami, R. F. Standaert, J. C. Smith, D. A. A. Myles, M. Ohl, and J. Katsaras, "Mechanical properties of nanoscopic lipid domains," *J. Am. Chem. Soc.* **137**(50), 15772–15780 (2015). DOI: 10.1021/jacs.5b08894. JIF: 12.311




Figure 3.6. The dynamics of contrast-enhanced unilamellar lipid vesicles reveal the local mechanical properties of nanoscopic domains. Image credit: Barmak Mostofian, John Nickels and Renee Manning

3.4 INDUSTRIAL HPC PARTNERSHIPS PROGRAM: ACCEL

ACCEL completed its seventh full year in 2015, helping companies advance in their ability to address complex science and engineering problems that exceed their internal computing capabilities. As these firms are sharing their results, ACCEL is helping DOE meet its goal of lifting the knowledge base of the science and engineering community, as well as increasing the community of researchers able to use next-generation leadership computing resources.

Forty-three industrial projects were under way during 2015, which represented 13.6% of the total number of projects provided to external user programs (INCITE, ALCC, and DD). These projects used 197,939,797 hours, representing approximately 4.6% of the total hours that Titan delivered in 2015.

- In 2015, 64% of the industrial project hours were allocated through INCITE, 19% via ALCC, and 17% through the OLCF DD program.
- Of 43 projects, 24 were new. These firms received awards via INCITE (2 projects), ALCC (4 projects), and DD (18 projects).
- A total of 7 of the 24 new projects were from companies that are first-time users at OLCF: automotive firm Fiat-Chrysler, rocket manufacturers Orbital ATK and SpaceX; industry-leading software firms Ansys, Dassault Systemes Simulia, and Numeca; and small software firm Appentra.
- Appentra represents another small business user at OLCF. Orbital ATK and SpaceX represent a new industry segment for OLCF—rocket manufacturers.

3.4.1 FM Global: Fighting Fires with Modeling and Simulation

To better understand how fires spread and the best methods to suppress them, FM Global, one of the world's largest commercial and industrial insurance companies, used Titan to scale its flagship code FireFOAM and model much larger fires than it could on its internal cluster.

FM Global has the world's largest fire technology laboratory with a fire testing area the size of a football field, and testing rooms with movable ceilings that can be raised from 15 to 60 feet high. Despite these world-class resources, FM Global researchers found they might not be able to replicate fires for the world's newest, largest "mega" warehouses and distribution centers.

The high costs of large-scale fire tests, along with the challenges in generalizing and extrapolating test results for these new mega-warehouses, prompted FM Global research scientists to develop computational models and simulate fires on an internal cluster. They did this using FireFOAM, a code they developed internally (and made open source) for simulating all of the complex physics that occur during an industrial fire. But FM Global researchers knew that gaining access to a larger supercomputer was necessary to improve their simulations. With a DD allocation at OLCF, the team was able to successfully scale FireFOAM from 100 CPUs to thousands of CPUs. And through the larger-scale simulations, they could conduct on Titan, they discovered that stacking storage boxes on wooden pallets significantly impedes the horizontal flame spread in an industrial fire, substantially reducing the fire hazard in the early stages of fire growth. FM Global is sharing these important insights with its clients and the combustion research community.

Related publication:

J. de Vries Ren, K. Meredith, M. Chaos, and Y. Wang, "FireFOAM modeling of Standard Class 2 commodity rack storage fires," published in the *Proceedings of Fire and Materials 2015* (February 2–4, 2015), 340.

Online story: "[Fighting Fire with FireFOAM](#)," *OLCF News* (July 7, 2015).

3.4.2 GE Successfully Develops New Predictive Simulation Methods for Gas Turbines

GE Global Research used more than 11.2 million compute hours on Titan to investigate combustion dynamic instability in gas turbines. Combustion dynamic instabilities have the potential to reduce the performance of new gas turbines, some of which include new DOE-funded technology.

A 2014 test of a gas turbine system discovered a new combustion instability mode that had not been previously observed. The mode was coupled between multiple combustion chambers. (Much like a car engine's cylinders, gas turbine combustors often have multiple combustion chambers or "cans.") Fortunately, the instability levels were acceptable for sustained operation and would not affect gas turbine performance.

GE wanted to understand the cause of the new instability and predict how it could manifest in future new products. The timing was good to examine the predictive capability of Cascade's CHARLES code, because a new gas turbine incorporating technology developed with DOE funds was to be tested in 2015. This presented an opportunity to demonstrate combustion instability prediction at the engine level.

A way to answer their questions with precision within the compressed timeframe was through modeling and simulation, but at a scale that far exceeded GE's own internal HPC systems. The research and development team turned to OLCF and received a DD allocation of time on Titan.

First, the team had to determine whether they could replicate the instability that appeared in the 2014 test. Then they needed to predict through simulation whether that instability would manifest in the new turbine design and at what level.

Using its Titan allocation, the GE team successfully delivered the first-ever GE multi-can dynamic instability simulations of a gas turbine's combustor system. With boundary condition tuning, the models accurately replicated the engine combustion instability mode from the 2014 test. These simulations revealed important new sub-millisecond insights into the instability and its causes that could not be obtained through physical testing.

The team also developed breakthrough new combustion modeling methods. By improving solver numerics and developing a new grid generation/optimization methodology, the team achieved a 30× speedup in the CHARLES code from Cascade compared with a 2014 prototype of the code.

The new insights coupled with the new simulation methods were used to predict the instability in the 2015 new turbine test. The simulations predicted a low instability level acceptable for operation without affecting the performance of the new turbine. The comparison to test results validated these predictions and affirmed the predictive accuracy of the new simulation methods GE developed and used on Titan.

Innovation

HIGH PERFORMANCE COMPUTING FACILITY 2015 OPERATIONAL ASSESSMENT OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2016

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. 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 section will focus on several key strategic areas of operations in 2015: data analysis and workflows, I/O optimization, application development, and computational resource management.

4.1 DATA ANALYSIS AND WORKFLOWS

The OLCF experienced a notable year in 2015 in the areas of data analysis and workflows. The NCCS responded to a growing need in these areas by creating a new group, the Advanced Data and Workflows Group, focusing on data-driven needs.¹⁴

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.

The OLCF followed the pilot thread, as reported in the 2014 Operational Assessment Report (OAR), and continued to partner with OLCF users to evaluate workflow technologies and use cases that are beneficial to users. Although the workflow of each pilot project had unique components, common requirements emerged; many of them could be met only by building upon existing scalable computing and data technologies and practices in operation at the OLCF. Two of these activities—ORiGAMI and BigPanDA—are described in Sections 4.2.2 and 4.2.3, respectively.

4.1.1 OLCF—State of the Workflow Update

Workflows and workflow systems enhance developer and scientist productivity. As the initial step toward understanding the science of workflows, OLCF conducted studies via several pilot projects by collaborating with its users in 2015. The aim of this exercise was to derive a comprehensive understanding of the current state of the theory and practice of workflow systems. Workflow requirements and expectations were documented based on discussions with several INCITE (Cyber Shake, Hardware/Hybrid Accelerated Cosmology Code, and Accelerated Climate Model for Energy

¹⁴ <https://www.olcf.ornl.gov/2015/10/13/group-leader-tapped-for-new-advanced-data-and-workflow-group/>

[ACME]) and strategic discretionary (BEAM, BigPanDA, and Center for Nanophase Materials Science [CNMS]) projects. These discussions concluded with the observation that the current proliferation of workflow systems in response to perceived domain-specific needs of scientific workflows makes it difficult to choose a site-wide operational workflow manager, particularly for the leadership-class machines. However, there are opportunities where facilities can centralize workflow technology offerings to reduce anticipated fragmentation. This is especially true if a facility attempts to develop, deploy, and operate each and every workflow solution requested by the user community. Through these evaluations, the OLCF seeks to identify interesting intersections that are of the most value to OLCF stakeholders.

Several existing workflow systems such as Pegasus, Galaxy, Kepler, Taverna, Bellerophon, Fireworks, and Swift were installed, evaluated, and supported for specific pilot projects. Pegasus was instantiated on a virtual machine to enable job submissions for the ACME and CyberShake workflows that ran jobs on multiple geographically distributed computing resources; a material science use-case from CNMS leveraged Bellerophon¹⁵; BigPanDA was deployed to use backfill cycles in support of the DOE Office of High Energy Physics (HEP) ATLAS and ALICE projects; and BEAM continued its success with image analysis for user-generated microscopy supporting collaborations at CNMS, the Institute for Functional Imaging of Materials, and the Spallation Neutron Source.¹⁶ The experience gained from these pilots has proved valuable in the planning and design of future hardware and software at the OLCF.

4.1.2 OLCF and CADES collaborate in Health Sciences

Director's Discretionary project CSC106 supported a collaboration between ORNL and the US National Library of Medicine (NLM), a division of the National Institutes of Health. The goal of the collaboration was to plow through millions of subject-verb-object groupings pulled from 23 million published articles and identify the strongest relationships, or better still the strongest potential relationships or the weakest yet most relevant associations, a specialized task that requires a specialized super-computer. Apollo, a Cray Urika supercomputer (XMT architecture), possesses massive multithreaded processors and 2 TB of shared memory, attributes that allow it to host the entire MEDLINE database and compute multiple pathways on multiple graphs simultaneously. Combined with Helios, the Compute and Data Environment for Science Cray Urika extreme analytics platform, the team had the cutting-edge hardware needed to process large datasets quickly—about 1,000 times faster than a workstation—and at scale. The result of this ORNL–NLM collaboration was featured in the Historical Clinical Pathological Conference in Baltimore as a demonstration of artificial intelligence for medicine, was a data demo at the IEEE/ACM Supercomputing conference, and was recognized at the Centennial Showcase of the Radiological Society of North America Conference in Chicago in December 2015.

These results were featured in a story in the *Los Angeles Times*¹⁷ newspaper; the following is an excerpt from that story.

While [Dr. Sanjay] Saint hit his books, computer specialists from the U.S. Department of Energy's Oak Ridge National Laboratory in Tennessee fed details of Cromwell's medical history into a network of supercomputers. A software program designed by Oak Ridge engineers churned through 27 million medical articles and texts, analyzing 70 million data associations relevant to Cromwell's case. In 4.5 seconds, ORIGAMI — short for Oak Ridge Graph Analytics for Medical Innovation — converged on virtually the same conclusion drawn after weeks of research and deliberation by Saint: Cromwell was done

¹⁵ <https://www.olcf.ornl.gov/2015/11/10/ornl-researchers-hail-promise-of-data-driven-materials-discovery-and-design/>

¹⁶ <https://www.ornl.gov/events/beam-computational-workflow-system-enabling-scalable-silico-and-empirical-exploration>

¹⁷ <http://www.latimes.com/science/la-sci-medical-mysteries-20160123-story.html>

in by malaria. For Saint, the unexpected competition was a humbling reminder of his limitations. But for Dr. Eliot Siegel, a University of Maryland radiologist and nuclear medicine specialist who worked on the ORIGAMI project, the exercise underscored the astonishingly efficient brainpower physicians bring to the task of diagnosis. Years of education and practice have endowed physicians such as Saint with a trove of medical data and an arsenal of cognitive shortcuts.

“Who would I rather have making a diagnosis? It would be hands-down Dr. Saint,” Siegel said. On the other hand, “if you told me there was a mystery disease and no one had any ideas about it and we needed some new insight ... I like the computer.” Oak Ridge engineers are preparing to reopen several of the cases pondered by Mackowiak and his colleagues through the years. The results may upend some of the conclusions so painstakingly drawn — and deepen the medical mysteries thought to be solved.

4.1.3 BigPanDA Workflow Management on Titan

The largest scientific instrument in the world—the Large Hadron Collider (LHC)—operates at the CERN Laboratory in Geneva, Switzerland. The ATLAS and ALICE experiments at the LHC explore the fundamental nature of matter and the basic forces that shape our universe. The BigPanDA project¹⁸ has provided the first important demonstration of the capabilities that a workload management system (WMS) can provide for improving the uptake and utilization of leadership computing facilities from both the application and systems points of view. Support from DOE ASCR and DOE HEP has led to the successful deployment of the BigPanDA workflow management tools on Titan and, as a result, Titan has been used by the ATLAS collaboration for Monte-Carlo Geant4 simulations. Working with collaborators from Brookhaven National Laboratory (BNL) and the University of Texas–Arlington, OLCF adapted PanDA for Titan and the OLCF environment, reusing much of the existing PanDA components and workflow.¹⁹

The project team developed and implemented a new capability in PanDA to collect information about unused worker nodes on Titan and, based on that information, adjust workload parameters to fill free resources. Proof-of-concept tests of this mechanism, executed over a few days, achieved increased system utilization levels from 90 to 93% (14.3% of free cycles) during the performance period of this test²⁰ and provided short wait times to ATLAS and ALICE for jobs submitted to Titan via PanDA. All of this was accomplished with no negative impact on OLCF’s ability to schedule large, leadership-class jobs.

Perhaps most important, Titan was fully integrated with the ATLAS PanDA-based Production and Analysis system, and now the ATLAS experiment routinely runs Monte-Carlo simulation tasks there. All operations, including data transfers to and from Titan, are transparent to the ATLAS Computing Operations team and physicists. Titan can contribute a noticeable fraction of computing resources for simulations. In November of 2015, utilization rose to 3 million core hours on Titan to run ATLAS detector simulation jobs, using only opportunistic, backfill resources.²¹ Work on the integration of the PanDA WMS with Titan was highlighted in communications by OLCF and BNL.²²

¹⁸ BigPanDA: DOE ASCR and HEP -funded project (2012–2015) to extend the ATLAS workload management system (aka PanDA) beyond the Grid, in particular to clouds and supercomputers.

¹⁹ A. Klimentov et al. “Next generation workload management system for big data on heterogeneous distributed computing,” *J. Phys. Conf. Ser.* **608**(1), 012040 (2015).

²⁰ Sergey Panitkin, Danila Oleynik, Kaushik De, Alexei Klimentov, Alexandre Vaniachine, Artem Petrosyan, Torre Wenaus, and Jaroslava Schovancova, *Integration of PanDA Workload Management System with Titan Supercomputer at OLCF*, Technical report, ATL-COM-SOFT-2015-021, 2015.

²¹ Backfill resources are those compute resources available for immediate use, e.g., because no resource requests of acceptable size or duration are made.

²² “World’s most powerful accelerator comes to Titan with a high-tech scheduler,” BNL Newsroom, May 7, 2014. <http://www.bnl.gov/newsroom/news.php?a=24864>.

4.2 I/O EFFICIENCY

4.2.1 Parallel File System Profiling with fprof

Fprof is a utility designed to traverse and characterize the usage of large-scale parallel file systems by gathering file count, file size, and associated distribution information, among other measures. It is built upon the principles of “work stealing pattern” and “distributed termination algorithm” to achieve scalability and performance. Work stealing pattern provides each worker process with the ability to handle its local work queue independently and can randomly query other workers for more jobs once its local work queue depletes. This independence is the key to achieving self-stabilization in the face of unbalanced work queues and varying processing speeds in a clustered environment. The distributed termination algorithm is a classic Dijkstra-proposed algorithm to allow a set of independent workers to collaboratively determine termination conditions without a single global coordinator.

As part of the suite of file system utilities developed by the OLCF, the fprof tool enabled staff to gain a deeper understanding of the file system and plot a file size distribution for the large-scale Spider II parallel file system, which consists of approximately half a billion files, in under 20 hours for a single run. Figure 4.1 illustrates the file size distribution on the Atlas1 partition. Using fprof, it has been determined that 70% of the stored files on the Spider II file system are smaller than 256 KiB. This information will be useful in helping OLCF staff gain deeper performance insights into current I/O workloads, and will assist in building next-generation file and storage systems in a more informed manner. The tool will also be open-sourced and released to the HPC and parallel file system community in the near future.




Figure 4.1. File size distribution of the Atlas1 partition of the Spider II file system.

4.2.2 Checkpoint Advisory Tool

In 2014, OLCF staff members Devesh Tiwari, Saurabh Gupta, and Sudharshan Vazhkudai presented the Lazy Checkpointing²³ concept with a goal toward optimal checkpointing intervals for improved scientific productivity and system efficiency. The team took the Lazy Checkpointing concept and created the Checkpoint Advisory tool in 2015.

System failures can impede the progress of scientific applications. Therefore, applications have traditionally relied on “checkpointing” as a resilience mechanism against failures. However, checkpointing and restoring the application state after a failure may exert severe pressure on the I/O

²³ <https://www.olcf.ornl.gov/2015/01/13/computer-scientists-increase-efficiency-through-lazy-checkpointing/>

subsystem, and it takes away the precious core hours from useful scientific computation. By not checkpointing often enough, the application assumes a higher risk of losing scientific work and may not make appropriate progress within the provided allocation. To address this dilemma, the OLCF set out to answer the critical but difficult question of preferred application checkpoint frequency. The resulting checkpoint advisory estimation is based on multiple factors, such as job size, system reliability characteristics, checkpointing overhead of the job, failure inter-arrival time, and time to write a single checkpoint. Overall, the Checkpoint Advisory tool has built-in intelligence based on mathematical grounds to provide such guidance. The tool is available as a module that users can easily load to incorporate this approach during application runtime.

Most recently, the Checkpoint Advisory tool helped the XGC turbulence application. The eventual aim of these simulations is to predict the divertor heat-flux width for ITER. In that process, multiple large-scale runs are required to make the plasma profile consistent with blobby turbulence. The implications for ITER's design and physics of these application runs are very significant. The runs used more than 16,000 nodes for approximately 24 hours. Initially, users set the checkpointing interval at one hour, which was conservative. Each checkpoint is approximately 4 TB in size. This amounts to close to 96 TB of checkpoint data for a 24 hour run. The Checkpoint Advisory tool suggested that the optimal checkpointing interval might be closer to 2 hours. This immediately provided a checkpointing overhead reduction of 50%, which has resulted in multiple direct and indirect benefits. First, the amount of space required to store checkpoint data shrank by 50% (close to 50 TB for a single run). This helps to improve OLCF daily processing and management tasks. Second, the stress on the storage subsystem was reduced, hence potentially avoiding performance degradation. Third, the change enabled precious core hours to be used for computation instead of waiting on I/O operation. Each checkpointing process took approximately 5 minutes, on average. Therefore, for a long 24 hour run, the longer checkpointing interval results in more than 490,000 core hours devoted back to useful computation. This is a significant step toward improving scientific productivity. Interestingly, increasing the checkpoint interval did not result in increased lost work. As expected, Titan was fairly reliable, so a job could often continue for more than 8 hours without any hardware interruptions. Overall, the Checkpoint Advisory tool helped these large-scale runs by expediting the discovery process while reducing the amount of I/O or data movement significantly.

4.2.3 HPSS Operational Investments and Improvements

In the second quarter of CY 2015, the High Performance Storage System (HPSS) administrative team deployed the Redundant Array of Independent Tapes (RAIT) feature in the HPSS instance of the OLCF. This feature provides operational redundancy if a tape hosting data is rendered inoperable. Previously, the only data protection method available within HPSS was data duplication. The RAIT feature allows a lower overall tape cartridge utilization and has recovery built in, just like the Redundant Array of Independent Disks (RAID). At present, OLCF is one of only two production HPSS sites that have this feature enabled. The OLCF first deployed a 4+1 RAIT configuration, meaning that files are split into four data chunks and a single parity chunk is written at each iteration. This feature has proved beneficial for both data integrity and performance. Where prior configurations allowed striping, having stripe widths greater than 2 saw sublinear scaling in real-world performance. With RAIT, the OLCF observes four times the performance of a single drive write.

In CY 2015, the HPSS team also deployed 18 PB of new disk cache. This addition was the culmination of 3 years of infrastructure investments that have allowed for a maximum daily ingestion increase from tens of TB per day to well over 200 TB per day. The increase has allowed several users of the OLCF systems to archive petabytes of data over the period of just a few weeks instead of months. The projects Cosmological Simulations for Large-scale Sky Surveys (HEP100), and Simulating Reionization of the Local Universe: Witnessing our own Cosmic Dawn (AST031) have both taken advantage of the increased ingestion performance within HPSS during CY 2015.

Additionally, the HPSS team took on the task of fully understanding the implications of putting the HPSS disk mover software onto an embedded virtual machine inside a disk storage controller. This effort

was undertaken to explore options related to the future architecture and ongoing operation of HPSS within the OLCF, potentially reducing complexity in deployment, points of failure, and the total cost to operate the HPSS archive. The evaluation showed that the technology to achieve this function is readily available, and that it is a trivial matter to integrate it into the OLCF's operational infrastructure for managing systems. Calculating a 5 year total cost of ownership showed approximately a 10% increase in overall system cost for equivalent bandwidth, but significantly reduced operational complexity. This information will be used to inform decisions regarding future disk cache acquisitions.²⁴

4.2.4 I/O Test Harness

The I/O test harness effort²⁵, first introduced in the 2014 OAR, was continued and improved upon in CY 2015. This approach was so successful in 2014 that the OLCF decided to further the initiative by continuing to invest in the development and maintenance of the test suite.

The harness was used to evaluate the effectiveness of the Lustre 2.8 client. During these evaluations, OLCF staff discovered an issue with LNET routing. This issue would have been difficult to identify lacking a full-scale system test or a dedicated I/O test harness to specifically stress components of the system. OLCF staff were also able to effectively gather baseline numbers for the I/O test harness as changes, enhancements, or additions were implemented to the application suite.

In November of 2015, Titan was upgraded to CLE 5.2 UP 04 with CUDA 7.0. This upgrade was a fairly major software change and required diligence to adequately test the system before releasing it to production. Both the acceptance and I/O harnesses were used in the increasingly larger software upgrades and tests leading up to the November Titan production upgrade. The harnesses were able to verify that the new CUDA version worked as advertised and that no regressions were introduced to the new Lustre client that was provided as a part of the new operating system.

4.3 APPLICATION DEVELOPMENT

4.3.1 GUIDO: Grand Unified Information Directory for OLCF Operations

The OLCF has developed the GUIDO software system to aggregate the logs from several subsystems into a single, unified infrastructure, which can then be queried for a variety of purposes. These include basic presentation of metrics, performance debugging of applications, application hot spot analysis, visual analytics, and use of the log information by higher-level tools and libraries. Previously, these system logs were too distributed, meaning that custom scripts or processes were required to extract and analyze them. Further, this distribution reduced the possibility of correlating the logs to infer deeper insights. In some cases, it was discovered that log data may not have ever received proper analysis, or that the analysis was not as comprehensive as it could have been.

To this end, the OLCF set up a log aggregation and analysis infrastructure using the Splunk commercial product. The team, consisting of members from the Technical Integration and HPC Operations teams, imported logs from several subsystems, such as the scheduling environment; block storage system; Lustre file system; and Titan/GPU reliability, DDNTool, LustreBRW Stats, and Titan system logs. The team is also working to integrate other data sources, such as HPSS usage statistics and Titan I/O router congestion data.

By plotting the I/O data, OLCF staff were able to perform visual analytics and performance debugging of large-scale runs of the XGC fusion application simulation in June 2015. These runs used the entirety of Titan and thus provided clear I/O patterns, as there were no other jobs running. It was initially observed that the application was only using ~800 of the 1,008 object storage targets (OSTs) available on

²⁴ <https://www.olcf.ornl.gov/2015/05/05/storage-system-infrastructure-upgrades-improve-performance-data-security/>

²⁵ <https://www.olcf.ornl.gov/2015/02/10/new-test-harness-keeps-olcf-resources-running-smoothly/>

the Atlas 2 partition of the Spider II file system. The users were made aware of this and they increased their OST usage, which resulted in a 15% improvement in the application's I/O performance.

Further, the log data from Splunk also serves as a source for other higher-level tools, such as the Checkpoint Advisory and file system selection tools. In fact, the Checkpoint Advisory tool combines two log data streams, the Titan/GPU reliability and Spider II I/O bandwidth logs, to determine an optimal checkpoint interval. As discussed in Section 4.2.2, the Checkpoint Advisory tool provided a recommendation to increase the checkpoint interval of the XGC application from 1 hour to 2 hours. This interval was observed and verified in analyzing the backend I/O operations per second of the Spider II block storage when an XGC job was running.

The health of the Spider II file system partitions are now displayed on the OLCF user support page next to the status indicators for Titan, Rhea, EOS and HPSS.²⁶ These indicators take the raw data from Splunk and distill it into simple red, yellow or green indicators, which provide an easily referenced point-in-time performance indicator for OLCF stakeholders.

4.3.2 GLEIPNIR: Understanding Application Memory Usage Patterns and Scalability

Gleipnir is a memory profiling and tracing tool, originally developed by Tomislav Janjusic during his PhD research at the University of North Texas and extended during a postdoctoral research fellowship at ORNL under the direction of Christos Kartsaklis. Gleipnir provides the ability to collect detailed traces of memory allocation, deallocation, and read/write operations. In a recent study,²⁷ Janjusic and Kartsaklis used the tool to understand the potential scalability of the locally self-consistent multiple scattering (LSMS) materials modeling code with respect to its memory usage. Using this tool, Janjusic and Kartsaklis were able to develop a detailed understanding of the memory usage within the LSMS application, including the ability to differentiate between memory used in the application itself versus in supporting libraries, and to attribute the memory usage to specific routines. They were also able to track how memory usage changes based on the number of MPI processes used for each data structure in the application.

The researchers found that memory usage in LSMS has relatively weak dependence on process counts: a test case that currently requires approximately 20 MB per process on a Titan-scale system with 300,000 processes (one per core) would require only 372 MB per process when scaled up to 100 million processes. However, the MPI library used by the application might be a greater concern. Measuring the Open MPI implementation of the MPI standard, the researchers projected that a case using 1,356 MB per process at today's scale of 300,000 cores would grow to 454 GB per process on a future system with 100 million processes. Although it may be some time before users are faced with systems with 100 million cores, this kind of analysis can help guide application developers as to when and where (within their code base) they might need to consider revising their approaches, as well as provide more detailed guidance to system designers on the potential impacts of node and system architectures and byte-to-FLOP ratios for future systems.

In a similar spirit, but in a separate study²⁸, the researchers used Gleipnir to create a new tool, Glprof. The new tool significantly reduces the complexity in understanding how intricate data structures are being accessed and how a given memory system responds to them, providing a report structured similarly to other widely-used tools such as Gprof and PAPI performance tools. Specifically, Glprof helps to connect the performance hot spots in the code with the data structures used most intensively in those regions.

²⁶ The status indicators can be found at the top of the OLCF Support page located at <https://www.olcf.ornl.gov/support/>.

²⁷ T. Janjusic and C. Kartsaklis, "Memory Scalability and Efficiency Analysis of Parallel Codes," Cray User Group (CUG 2015), April 2015, Chicago, Illinois.

²⁸ T. Janjusic and C. Kartsaklis, "Glprof: A Gprof inspired, Callgraph-oriented Per-Object Disseminating Memory Access Multi-Cache Profiler," *Workshop on Tools for Program Development and Analysis in Computational Science (TOOLS 2015)*, in conjunction with ICCS, June, 2015, Reykjavik, Iceland.

4.4 RESOURCE MANAGEMENT

4.4.1 Dual-Ended Scheduling

Resource contention on HPC systems can lead to significant performance variability (i.e. longer runtimes for some jobs). This is exacerbated when large-scale HPC systems are highly utilized and are multiplexed with uncoordinated workloads, as is done on Titan. For example, in a recent test on Titan, the team measured up to 24% variations in performance between two separate invocations of the same application. The implications of variability can lead to jobs being terminated before simulation convergence, ultimately resulting in wasted computing hours. Therefore, it is critical to reduce runtime variability through smart techniques.

Cray's Application Level Placement Scheduler (ALPS), Adaptive Computing's MOAB scheduler, and the TORQUE resource manager determine which nodes to assign to a given job on Titan. ALPS generates the layout of the machine by enumerating the set of compute nodes and creates a list of nodes optimized for communication locality and latency. MOAB is responsible for the mapping of free resources to job reservations. Using a first-fit strategy, MOAB performs a linear scan of the list of resources provided by ALPS until it finds enough free resources to satisfy the requirements of a reservation. For example, if a reservation requests five nodes, MOAB will scan the list and allocate the first five free nodes that it passes during the scan. These nodes need not be consecutive nor close in physical or logical proximity; the only criterion is that they are available. In a heavily utilized machine, this will result in highly fragmented allocations, which can lead to unmitigated consequences such as longer runtimes as simulations are forced to communicate over and through other compute jobs and I/O patterns.

To address this problem, a team (Chris Zimmer, Scott Atchley, Sudharshan Vazhkudai, Don Maxwell, Saurabh Gupta, and Carl Albing, formerly lead designer at Cray for ALPS) devised a novel algorithm to segregate small, short-lived jobs from larger, longer-running jobs in hopes that this would reduce fragmentation and variability. To accomplish this, the team proposed scheduling from both "ends" of the node list used by MOAB. Jobs smaller than a certain threshold select available nodes from the "end" of the list, whereas larger jobs continue the existing approach of selecting available nodes from the "beginning" of the list. The team calls this algorithm "Dual-Ended Scheduling."

To determine the threshold, the team analyzed the actual Titan workload from January to July 2015. In this time period, Titan had ~180,000 jobs, of which 98,000 required 16 nodes or fewer. The team used the MOAB simulator to implement the Dual-Ended Scheduling algorithm and analyzed its merits. The team fed the simulator a production job trace of 976 jobs extracted from Titan and confirmed that it allocated nodes from both ends of the node list based upon this threshold. To evaluate the effectiveness of the algorithm, the team used an all-to-all distance metric, which calculates the mean distance (i.e. hop count) between nodes within the allocation. Comparing the default scheduling with Dual-Ended Scheduling, the team observed lower all-to-all distances for the job allocations.

Starting on July 15, Titan began using Dual-Ended Scheduling. Since then, two demarcation points have been evaluated with thresholds of 16 nodes and 125 nodes. The 16-node evaluation showed a decrease in average hop-count across all of the evaluated job sizes, representing the most common job sizes that use above 500 nodes. The 125-node evaluation showed some reduced fragmentation but appears to have created a hot-spot around 2,048 nodes; this is consistent with the simulation workloads and results. Based on these results, the team has reverted to a threshold of only 16 nodes. One further evaluation is planned using a time-based demarcation strategy that has shown promising results in the simulation environment.

The investigation into the scheduling interactions showed that MOAB uses a first-fit algorithm, which favors nodes on one end of the node list, meaning that it schedules these nodes more frequently than those found on the opposite end of the node list. It was discovered that approximately the last 100 nodes on Titan were rarely scheduled. These nodes typically were scheduled only when Titan ran at or near 100% utilization. This approach can lead to imbalanced resource usage patterns, which can have an effect on

component lifecycles, resulting in decreased mean time to interrupt, mean time to failure and, in some cases, an increase in the duration and frequency of system downtimes. Duel-Ended Scheduling seeks to use the computational resource more efficiently and proved that there were efficiencies to be gained.

Risk Management

HIGH PERFORMANCE COMPUTING FACILITY 2015 OPERATIONAL ASSESSMENT OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2016

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, reclassified, or mitigated. A change history is maintained for historical reference.

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

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 a change summary of affected 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 issues 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 47 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 2015

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

Risk Owner	Jack C. Wells, NCCS Director of Science		
Probability	Medium		
Impact	<i>Cost:</i> Low	<i>Schedule:</i> Medium	<i>Scope/Tech:</i> 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 (CR) score greater than or equal to 4 out of 5. Likewise, of 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 experience in porting to the heterogeneous architecture. Many applications teams appear to be porting their codes.

OLCF has also seen an increase in the number of ALCC submissions. In addition, 29 CAAR proposals (out of 8 available slots) were submitted for work on Summit, which is also a heterogeneous architecture. The heterogeneity of Summit requires that the OLCF continue to monitor this risk.

5.2.2 ID# 407—Loss of Key Personnel

Risk Owner	Arthur S. Bland, OLCF Project Director		
Probability	Medium		
Impact	<i>Cost:</i> Low	<i>Schedule:</i> Low	<i>Scope/Tech:</i> 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.

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

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

Metadata performance is critical to a wide variety of leadership applications. It 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 the 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 has deployed Lustre 2.5, which has the Distributed Namespace (DNE) Phase 1 feature meant to address part of the metadata bottleneck. As part of the collaboration with OpenSFS, the DNE Phase 2 feature will be released in Lustre 2.8, targeted for mid-CY 2016. Because potential migration tools are not available for a DNE Phase 1 to Phase 2 migration, the OLCF will postpone the DNE Phase 2 implementation to provide the greatest impact for the facility's computational users. In addition, the OLCF is upgrading the metadata servers with solid-state hard drives and faster processors, which will improve overall metadata performance.

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 risk continues to be monitored.

5.2.4 ID# 1063—Programming Environment Tools May Be Insufficient

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

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 NCCS initiated contracts with vendors to accelerate development and to add needed key functionality to the tools. The NCCS has developed portable programming models (through vendor partners) such as the directive-based OpenACC standard and the OpenMP directives for accelerators. 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, and has contracted with Allinea and Dresden to make 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	<i>Cost:</i> Medium	<i>Schedule:</i> Low	<i>Scope/Tech:</i> 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 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 MITIGATION

5.3.1 ID #408 – System Outages from External Causes

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

Systems may be negatively impacted by external causes, including adverse power quality, events affecting mechanical systems availability, wide area network outages, and other external events that are largely out of the immediate control of the OLCF. Outages may cause existing work to fail, introduce unscheduled outages and maintenance events, disrupt time-sensitive workflows, introduce unanticipated delays in user job completion, or otherwise hinder system performance.

Mitigations include the following actions. For electrical risks, constantly evaluate the risks in this area. Identify and implement diverse sources of power to the NCCS. Implement cost-effective backup capabilities (e.g., generators, uninterruptable power sources, dual core cabinet designs) where applicable. For network risks, implement diverse physical and light paths for wide-area-network connections to diverse remote locations (Atlanta, Chicago). Continue to harden local metropolitan and wide -area - network infrastructure where applicable. For mechanical risks, further harden the infrastructure through additional capacity and control systems.

On December 24, 2015, the Tennessee Valley Authority (TVA) experienced three major power quality events across its 161 kV distribution system, caused by a significant number of lightning strikes in the eastern Tennessee area. Every single-fed system not connected to an uninterruptable power source went down during the first power quality event. However, not a single component was lost during the recovery process. An after -action meeting held in early 2016 concluded there were no material corrective actions.

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

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

The OLCF Building 5600 central electrical plant was built and is maintained to support the total capacity of the Building 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 costs 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 seek to reduce maintenance costs by reducing the amount of 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 interests of ORNL.

During 2015, OLCF signed a new 5year contract with the National Oceanic and Atmospheric Administration to provide infrastructure support for its replacement for Gaea. This agreement has helped lower the probability of this risk.

5.4 RISKS RETIRED DURING THE CURRENT YEAR

5.4.1 ID# 124—Storage System Reliability and Performance Problems

Risk Owner	Sudharshan S. Vazhkudai, Group Leader, Technology Integration		
Probability	Low		
Impact	<i>Cost:</i> Low	<i>Schedule:</i> Low	<i>Scope/Tech:</i> Medium
Rating	LOW		
Status	RETIRED: Performance and reliability have met requirements		

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 FY 2013, quarter 4, at no cost to the OLCF. This 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 remains; therefore, this risk item remains active.

Retired: Acceptance testing was in October and the system was accepted. Performance and reliability have met stated requirements.

5.4.2 ID# 1093—Impact of Load Swings on Power Grid

Risk Owner	Bart A. Hammontree, OLCF Project Manager		
Probability	Low		
Impact	Cost: Low	Schedule: Low	Scope/Tech: Low
Rating	LOW		
Status	RETIRED: The load swings, even with Summit's anticipated larger range, are well within tolerable parameters for TVA.		

HPC system startup or shutdown can result in sudden power increases or decreases to the electrical grid. These large power swings can result in transient impacts to power grids, which can cause power quality problems for the electrical utility.

To date, no transient power impacts have been observed or reported at ORNL as a result of large HPC load swings. To confirm that future load swings will not have a negative impact on the power grid, ORNL will model the local power grid and simulate the impact of sudden load swings of up to 20 MW in magnitude. If the model indicates negative impacts, we will work with utility providers to develop an engineered solution to prevent negative impacts.

Trigger: Results from the new TVA model or other intelligence on potential power issues or actual grid problems.

Retired: Discussions with TVA and output from the model showed the power system can tolerate large load swings (even as large as the ± 10 MW predicted with Summit).

5.4.3 ID# 1143—ORNL-Furnished equipment

Risk Owner	Bart A. Hammontree, OLCF Project Manager		
Probability	Low		
Impact	Cost: Medium	Schedule: Medium	Scope/Tech: Medium
Rating	Low		
Status	RETIRED: The OLCF found ways to get the data center tax breaks extended to vendors.		

To guarantee data center sales tax incentives, ORNL may procure equipment required for operations. However, ORNL might procure equipment that does not meet the project requirements, procure it late, or damage the equipment during delivery.

ORNL will work with the design-build contractors to ensure that the equipment to be procured meets the design requirements and is delivered in a time frame that does not impact schedule performance. ORNL will accept the risk of vendor defects and try to resolve any defect issues with the vendors as quickly as possible.

Retired: The OLCF found ways to get the data center tax breaks extended to vendors so ORNL would not need to buy the equipment.

5.5 NEW OR RECHARACTERIZED RISKS SINCE LAST REVIEW

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

Risk Owner	James Rogers, NCCS Director of Computing and Facilities		
Probability	Low		
Impact	<i>Cost:</i> Low	<i>Schedule:</i> Low	<i>Scope/Tech:</i> 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 reclassified 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 field-replaceable unit 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 ten 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 and might require OLCF to reevaluate the probability as Titan ages.

5.5.2 ID# 1006—Inability to Acquire Sufficient Staff

Risk Owner	Arthur S. Bland, OLCF Project Director		
Probability	High → Medium		
Impact	<i>Cost:</i> Low	<i>Schedule:</i> Low	<i>Scope/Tech:</i> Low
Rating	MEDIUM		
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.

The probability of this risk occurring was decreased from High to Medium in 2015 because the OLCF is currently fully staffed. We will continue to monitor staffing levels.

5.5.3 New Risk: ID #1145 —Changes from External Project Managers Cause Development Impacts to HPSS

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

The software used for tape storage at the OLCF is called the High Performance Storage System (HPSS). HPSS is a distributed development project involving a commercial entity (IBM) and multiple government laboratories. Feature requests and reprioritizations from one entity can cause development disruptions. Features requested or required by OLCF users might not be implemented in a timely fashion or at all. Features requested or required by other laboratories/IBM might have adverse side effects (e.g., impaired performance, increased bugs) for OLCF users.

Each laboratory has a representative to the technical council that helps set the direction for HPSS development. OLCF works closely with the council to ensure users’ needs are met. However, IBM has trumped established development plans before, causing disruptions in schedules.

Trigger: Input from the technical council changing project direction

5.5.4 New Risk: ID #1146 —Suboptimal Resource Utilization in HPSS Causes Lack of Storage and Greater Expenses

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

Suboptimal resource utilization causes excessive costs and performance issues for HPSS users. Storing small files is inefficient. The disk bandwidth is much greater (and growing) compared with tape bandwidth. Tapes, tape robots, and silos are expensive in terms of dollars and physical space. Poor space utilization leads to inadequate space for users to store archival data.

Trigger: Increase in resource utilization or costs

To respond to this risk, the OLCF is increasing the bandwidth from disks to tape libraries, deprecating older tape formats as soon as is practical, providing user education regarding small files, and supporting HPSS Tape Archiver (htar). In addition, there is now a larger disk cache for the tape drives.

5.5.5 New Risk: ID #1147 —Data Integrity Issues in Archival Data in HPSS

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

As increasing amounts of data are stored on tape, the probability of a data integrity failure increases. The risk is not being able to identify corrupt data and manage it appropriately. An HPSS software bug, degradation of data on tapes, or inadvertent administrator action can cause the archived data to become corrupted or unreadable. This might lead to incorrect application results or loss of data altogether.

Trigger: Reported loss of data by users or automated tools

To respond to the risk, OLCF employs redundant array of independent tapes (RAIT) technology when writing data to tape. Newer tape formats are also more robust. In addition, a “data crawler” utility was developed to verify data on tape periodically.

5.6 MAJOR RISKS FOR NEXT YEAR

The major risks for 2016 will be similar to the major risks tracked during 2015. However, as Titan has now completed the second full calendar year in production, and users continue to effectively use the system, many of the risks of Titan’s architecture are lower than the level at which the center rated them in years past. As Titan ages, the OLCF will continue to track and monitor component failures to ensure high availability and usability to the OLCF stakeholders.

Site Office Safety Metrics

HIGH PERFORMANCE COMPUTING FACILITY
2015 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2016

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 CFR 851,²⁹ “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, safety talks, 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 it partners with the ORNL Offices of Institutional Planning, Integrated Performance Management, and Safety Service Division on some independent safety management system assessments.

The culture of safety at ORNL is reflected in these 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 Dan Hoag, the Federal Project Director, who

²⁹ 10 CFR 851 outlines the requirements for a worker safety/health program to ensure that DOE contractors and their workers operate a safe workplace. Additionally, 10 CFR 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 2015 Performance Evaluation and Measurement Plan. A review of the monthly Safety Charts and the TRC and DART summary documents submitted to OSO indicated that overall FY 2015 TRC and DART rates have decreased from FY 2014. Operations of the OLCF in the NCCS remained safe, efficient, and effective. They have resulted in zero TRCs and zero DARTs in FY 2015.
- UT-B placed emphasis on the subjects of Material Handling and Electrical Safety because of incidents and recommendations from DOE. UT-B has lowered the incident rate of material handling injuries in 2015 by incorporating some interim rules that provided lifting restrictions in some scenarios and is in the final stages of finalizing procedures that will specifically address this subject matter and provide guidance to all employees. The Electrical Safety Program continues to be enhanced by implementing an arc flash analysis procedure, creating R2A2s (roles, responsibilities, authorities, and accountabilities) for division electrical safety officers, enhancing the lockout-tagout and electrical training, and reviewing electrical equipment across the site to evaluate the risk of using equipment that is not inspected by a National Recognized Testing Laboratory.
- The OLCF-4 project Health and Safety Plan and the Hazard Analysis were reviewed in preparation for the DOE Office of Science Independent Project Review conducted in October 2015. No health and safety issues were cited during the out-briefing.
- The OLCF project strives to go above and beyond the minimum requirements by requiring limited scope hazard analyses and by employing award-winning professionals throughout the project. Each contractor is required to submit an Activity Hazard Analysis to UT-B personnel for review so UT-B can be aware of the scope of work, ensure hazards and controls have been appropriately addressed, and understand the plan for completing the job scope. This process keeps UT-B aware of the projects being completed while at the same time holding contractors accountable for recognizing hazards and implementing mitigating controls. The OLCF hired John Campbell as project safety officer in 2015. John has been awarded the 2015 Rising Star Award by the National Safety Council, which recognizes 40 health and safety professionals under the age of 40 who have demonstrated their commitment to safety and continuous improvement. John, who is a Certified Safety Professional, brings a broad safety and health background related to demolition, construction, and utilities.³⁰

³⁰ <https://www.olcf.ornl.gov/2015/09/29/campbell-to-receive-40-under-40-safety-award/>

Cyber Security

HIGH PERFORMANCE COMPUTING FACILITY 2015 OPERATIONAL ASSESSMENT OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2016

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 15, 2015. The current ATO expires on June 15, 2016.

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 implementation of a continuous monitoring program 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 15, 2015. 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 of controls for protecting the confidentiality and integrity of user and system information (Federal Information Processing Standards Publication 199), which authorizes the facility to process sensitive, proprietary, and export-controlled data.

In the future, 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 to which it will respond and make modifications as the needs of the facility change to provide an appropriately secure environment.

7.2 SAFEGUARDS AND SECURITY SURVEY

Representatives from the DOE Office of Enterprise Audits were onsite July 13–17, 2015, to conduct the periodic Safeguards and Security Survey of ORNL, specifically related to cyber security and OLCF policies. The onsite review included a tour of the facilities, personal interviews with cyber security staff, and clear demonstration of the implementation of various technical security controls. The supercomputing sections of the cyber audit did not reveal any findings or observations. In particular, OLCF HPC Operations was recognized as having a good mentorship program for bringing in new system administrators.

Summary of the Proposed Metric Values

HIGH PERFORMANCE COMPUTING FACILITY
2015 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

March 2016

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 2015 and proposed metrics and targets for 2016 and 2017

2015 metric and target	2015 actual	2016 target	2017 target
Are the processes for supporting the customers, resolving problems, and outreach effective?			
<i>Customer Metric 1: Customer Satisfaction</i>			
Overall score on the OLCF user survey. Target: Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample. (Annual)	The OLCF exceeded the metric target: 4.6/5.0.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.
Improvement on results that scored below satisfactory in the previous period. Target: Results will show improvement in at least one-half of questions that scored below satisfactory (3.5) in the previous period. (Annual)	The OLCF exceeded the metric target: No question scored below satisfactory (3.5/5.0) on the 2015 survey.	Results will show improvement 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.
<i>Customer Metric 2: Problem Resolution</i>			
OLCF survey results related to problem resolution. Target: Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample. (Annual)	The OLCF exceeded the metric target: 4.6/5.0.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.
OLCF user problem resolution time period. Target: 80% of OLCF user problems will be addressed within 3 business days, by either resolving the problem or informing the user how the problem will be resolved. (Monthly/Annual)	The OLCF exceeded the metric target: 92%.	Eighty percent of OLCF user problems will be addressed within 3 business days, by either resolving the problem or informing the user how the problem will be resolved.	Eighty percent of OLCF user problems will be addressed within 3 business days, by either resolving the problem or informing the user how the problem will be resolved.
<i>Customer Metric 3: User Support</i>			
Average of user support ratings. Target: Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample. (Annual)	The OLCF exceeded the metric target: 4.6/5.0.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.
Is the facility maximizing the use of its HPC systems and other resources consistent with its mission?			
<i>Business Metric 1: System Availability^a</i>			
Scheduled Availability. 2015 Targets: Titan: 95%; HPSS: 95%; external file systems: 95%. (Monthly/Annual)	The OLCF exceeded the metric target. Titan: 99.41%; HPSS: 99.94%; /atlas0: 99.77%; /atlas1: 99.76%.	Titan: 95%; HPSS 95%; external file systems 95%.	Titan: 95%; HPSS 95%; external file systems 95%.

Table 8.1. (continued)

2015 metric and target	2015 actual	2016 target	2017 target
Overall Availability. 2015 Targets: Titan: 90%; HPSS 90%; external file systems 90%. (Monthly/Annual)	The OLCF exceeded the metric target: Titan: 97.01%; HPSS: 98%; /atlas0: 98.9%; /atlas1: 98.88%.	Titan: 90%; HPSS 90%; external file systems: existing, 90%.	Titan: 90%; HPSS 90%; external file systems: existing, 90%.
<i>Business Metric 2: Capability Usage</i>			
OLCF will report on capability usage. Target: In the first year of production, at least 30% of the consumed node-hours will be from jobs requesting 20% or more of the available compute nodes. In subsequent years, this increases from 30% to 35%. (Monthly/Annual)	The OLCF exceeded the metric target. The capability usage was 64.19%.	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.
OLCF will report GPU usage (reference only, no target)	53.13%	N/A ^b	N/A ^b

^a The Cray XK7, Titan, went into production on May 31, 2013. The external file system, Spider II, went into production on October 3, 2013.

^b Not applicable.

APPENDIX A. RESPONSES TO RECOMMENDATIONS FROM THE 2014 OPERATIONAL ASSESSMENT REVIEW

In March 2015, the OLCF presented the 2014 operational activities of the center to the DOE sponsor. During this review, no recommendations were presented. The overall assessment of the OLCF for CY 2014 is as follows:

The OLCF is an exceedingly well-managed organization that is fulfilling its mission to enable world-class science with leadership computing systems. OLCF has across-the-board high results on availability, utilization, reliability, and its user survey, exceeding all performance targets. OLCF achieves efficient and effective use of its resources through efforts all across the organization. OLCF is a leader in applying GPU technology to large-scale HPC. OLCF continually seeks to innovate in areas that improve operational performance and advance the state-of-the-art. OLCF has a strong and mature safety and risk management program, with no safety incidents in CY 2015. OLCF has management and staff that are both deeply committed to its mission. Keep up the good work!

APPENDIX B. TRAINING, WORKSHOPS, AND SEMINARS

Event type	Description	Date	Participants
Workshop/Training	CAAR: Call for Proposals Webinar	01/06/15	39
User Con Call	OLCF Users Conference Call	01/14/15	30
Workshop/Training	Portability Meeting	1/27/15 – 1/29/2015	70
User Con Call	OLCF Users Conference Call: Score-P	02/18/15	43
Workshop/Training	Advances in Nuclear Fuel Management Workshop	03/06/15	20
User Con Call	OLCF Users Conference Call	03/25/15	33
Workshop/Training	2016 INCITE Proposal Writing Webinar	04/09/15	24
User Con Call	OLCF Users Conference Call	04/15/15	35
Seminar Series	Software Engineering for HPC—Experiences in Developing Software Tools	04/16/15	—
Workshop/Training	Allinea DDT Training	4/16/2015 – 4/17/2015	9
Workshop/Training	2015 GPU Hackathons—NCSA	4/20/15 – 4/24/2015	36
Workshop/Training	Center for Accelerated Readiness (CAAR) Workshop 1	04/22/15 – 04/23/2015	37
Workshop/Training	OpenACC Training F2F Meeting and User Feedback	05/12/2015 – 05/14/2015	23
User Con Call	OLCF Users Conference Call	05/20/15	27
Workshop/Training	2016 INCITE Proposal Writing Webinar	05/27/15	60
Workshop/Training	Center for Accelerated Readiness (CAAR) Workshop 2	06/02/15	42
Seminar Series	Using Splunk as a Data Source	06/12/15	—
Workshop/Training	2015 OLCF Users Meeting	06/23/2015 – 06/25/2015	106
Workshop/Training	New Users Training	06/22/15	18
Workshop/Training	2015 GPU Hackathons—CSCS	7/6/2015 – 7/10/2015	65
Workshop/Training	Supercomputing in a Nutshell (Track 1) at the 2015 DOE CSGF Annual Program Review	07/29/15	—
Workshop/Training	Virtual School of Computational Science and Engineering	08/03/15	7
Seminar Series	Performance Analysis for GPU Computing	08/07/15	37
Seminar Series	On the Implications of Large-Scale Manycores and NoCs for Exascale Prof. Frank Mueller, North Carolina State University	08/11/15	—
Seminar Series	Code Integration Between the BISON Fuel Performance and PROTEUS Neutronics Applications	08/12/15	—
Seminar Series	“Oh Rats!” — The NCCS Resource and Allocation Tracking System	08/21/15	—
Seminar Series	Integrated Modeling and Simulation with Eclipse ICE and Its Applicability to Neutron Science	09/17/15	—
User Con Call	OLCF Users Conference Call	09/23/15	41

Event type	Description	Date	Participants
Seminar Series	LSMS and WL-LSMS: Codes for First-Principle Calculation of the Ground State and Statistical Physics of Materials	09/30/15	—
Workshop/Training	2015 GPU Hackathon—ORNL	10/19/2015 – 10/23/2015	72
Seminar Series	Build and Test Automation at Livermore Computing	10/20/15	—
User Con Call	OLCF Users Conference Call	10/28/15	43
User Con Call	OLCF Users Conference Call	11/15/15	14
Workshop/Training	Parallel Programming in Modern Fortran	11/15/2015 – 11/20/2015	25
Conference	BigNeuron Data Analytics and Visualization Design Meeting	11/16/2015 – 11/20/2015	126
Workshop/Training	Scientific Software Development Design Meeting	12/1/2015 – 12/3/2015	19
Seminar Series	The Ceph Storage System: New Features and Future Directions	12/02/15	—
Seminar Series	Computational Understanding and Design of Transition Metal Catalysts	12/11/15	—
User Con Call	OLCF Users Conference Call	12/16/15	43
Seminar Series	Using Software Engineering Methodologies to Port a Scientific Code to GPUs	12/21/15	—

APPENDIX C. OUTREACH PRODUCTS

Date	Type of product	Title
1/15/15	Highlight	Computer Scientists Increase Efficiency Through Lazy Checkpointing
1/15/15	PPT Slide	Computer Scientists Increase Efficiency Through Lazy Checkpointing
1/15/15	Highlight	New Test Bed Helps OLCF Scale Up for Next-Generation Supercomputer
1/15/15	PPT Slide	New Test Bed Helps OLCF Scale Up for Next-Generation Supercomputer
1/15/15	Highlight	OLCF Emphasizes Importance, Furthers Development of OpenACC Standard
1/15/15	PPT Slide	OLCF Emphasizes Importance, Furthers Development of OpenACC Standard
1/29/15	Poster	ALCC
2/3/15	Chart	Accounts Flow Chart
2/10/15	Highlight	New Test Harness Keeps OLCF Resources Running Smoothly
2/10/15	PPT Slide	New Test Harness Keeps OLCF Resources Running Smoothly
2/10/15	Highlight	Wigner Distinguished Lecturer Visits OLCF
2/10/15	PPT Slide	Wigner Distinguished Lecturer Visits OLCF
2/10/15	Highlight	Buddy Bland Named as one of HPCwire's 2015 People to Watch
2/10/15	PPT Slide	Buddy Bland Named as one of HPCwire's 2015 People to Watch
3/1/15	Publication	White Paper Version of Annual Report
3/10/15	Highlight	Researchers Get Warmer in Understanding High-Temperature Superconductors
3/10/15	Quad Chart	Researchers Get Warmer in Understanding High-Temperature Superconductors
3/10/15	Highlight	Researchers Dig Up Biological Data from Microbes in Mines
3/10/15	Quad Chart	Researchers Dig Up Biological Data from Microbes in Mines
3/10/15	Highlight	Test Systems Pike and Crest Pave the Way for Next OLCF Supercomputer
3/10/15	PPT Slide	Test Systems Pike and Crest Pave the Way for Next OLCF Supercomputer
3/10/15	Files	ALCC Exec Summary Compile
3/17/15	Graphics	Social Media Icons for ORNL and OLCF
4/1/15	Graphics	CUG swag artwork
4/8/15	Highlight	Titan Vaporizes R&D Time and Cost for Liquid Spray Studies
4/8/15	Quad Chart	Titan Vaporizes R&D Time and Cost for Liquid Spray Studies
4/8/15	Highlight	OLCF Staff Study Possible Temperature Changes to Chilled Water Systems
4/8/15	PPT Slide	OLCF Staff Study Possible Temperature Changes to Chilled Water Systems
4/8/15	Highlight	OLCF Well Represented at NVIDIA GPU Technology Conference
4/8/15	PPT Slide	OLCF Well Represented at NVIDIA GPU Technology Conference
4/8/15	Highlight	OLCF User Earns NVIDIA Award for GPU-accelerated Earthquake Simulations
4/8/15	PPT Slide	OLCF User Earns NVIDIA Award for GPU-accelerated Earthquake Simulations
4/15/15	Branding	OLCF Identity Tweaks
4/15/15	Highlights/Art	OLCF Selects Application Readiness Projects to Prepare for Summit Supercomputer
5/4/15	Highlight	Fusion Researchers Use Titan to Burst Helium Bubbles
5/4/15	Quad Chart	Fusion Researchers Use Titan to Burst Helium Bubbles
5/4/15	Highlight	New RAIT System Keeps Large Data Sets Safe
5/4/15	PPT Slide	New RAIT System Keeps Large Data Sets Safe
5/4/15	Highlight	Storage System Infrastructure Upgrades Improve Performance, Data Security
5/4/15	PPT Slide	Storage System Infrastructure Upgrades Improve Performance, Data Security

Date	Type of product	Title
5/4/15	Highlight	OLCF Staff, Users Make Strong Showing at APS
5/4/15	PPT Slide	OLCF Staff, Users Make Strong Showing at APS
5/26/15	Highlight	Using the Forest to See the Trees
5/26/15	Quad Chart	Using the Forest to See the Trees
5/26/15	Highlight	OLCF Names CAAR Projects at ICD HPC User Forum
5/26/15	PPT Slide	OLCF Names CAAR Projects at ICD HPC User Forum
5/26/15	Highlight	OLCF Shares Lustre Knowledge at International Workshop
5/26/15	PPT Slide	OLCF Shares Lustre Knowledge at International Workshop
5/26/15	Highlight	The First Round of 2015 Hackathons Gets Underway
5/26/15	PPT Slide	The First Round of 2015 Hackathons Gets Underway
6/3/15	Publication	OLCF Annual Report
6/16/15	Highlight	The Protein Problem
6/16/15	Quad Chart	The Protein Problem
6/16/15	Highlight	Mastering Magnetic Reconnection
6/16/15	Quad Chart	Mastering Magnetic Reconnection
6/16/15	Highlight	Tiny Titan Is Now Accepting Visitors
6/16/15	PPT Slide	Tiny Titan Is Now Accepting Visitors
6/16/15	Highlight	DOE User Facilities Unite at ORNL for Annual Meeting
6/16/15	PPT Slide	DOE User Facilities Unite at ORNL for Annual Meeting
6/30/15	Poster	SciDAC Poster 1
6/30/15	Poster	SciDAC Poster 2
7/7/15	Graphics	Updates to Welcome Pack and Slides
7/7/15	Highlight	Unlocking Lignin for Sustainable Biofuel
7/7/15	Quad Chart	Unlocking Lignin for Sustainable Biofuel
7/7/15	Highlight	The Ins and Outs of QCD
7/7/15	Quad Chart	The Ins and Outs of QCD
7/7/15	Highlight	22 Projects Earn 1 Billion Processor Hour Allocations through ALCC
7/7/15	PPT Slide	22 Projects Earn 1 Billion Processor Hour Allocations through ALCC
7/7/15	Highlight	Science DMZ Set to Expedite Data In and Out of ORNL
7/7/15	PPT Slide	Science DMZ Set to Expedite Data In and Out of ORNL
7/7/15	Highlight	Students Study Computational Nuclear Reactor Modeling at CASL Summer Workshop
7/7/15	PPT Slide	Students Study Computational Nuclear Reactor Modeling at CASL Summer Workshop
7/28/15	Highlight	Researchers Build Bacteria's Photosynthetic Engine
7/28/15	Quad Chart	Researchers Build Bacteria's Photosynthetic Engine
7/28/15	Highlight	OLCF Postdoc Collaborates with Researchers to Simulate Magnesium Carbonate
7/28/15	PPT Slide	OLCF Postdoc Collaborates with Researchers to Simulate Magnesium Carbonate
7/28/15	Highlight	OLCF Users Talk Latest Science, Elect OUG Board Members during Annual Meeting
7/28/15	PPT Slide	OLCF Users Talk Latest Science, Elect OUG Board Members during Annual Meeting
7/28/15	Highlight	OLCF Science-Writing Internship Benefits Students, ORNL
7/28/15	PPT Slide	OLCF Science-Writing Internship Benefits Students, ORNL

Date	Type of product	Title
7/28/15	Highlight	Banishing Bad Cache Behavior
7/28/15	PPT Slide	Banishing Bad Cache Behavior
7/28/15	Highlight	OLCF Upgrades System to CUDA 6.5
7/28/15	PPT Slide	OLCF Upgrades System to CUDA 6.5
8/17/15	Highlight	The Future of Forecasting
8/17/15	Quad Chart	The Future of Forecasting
8/17/15	Highlight	Code Speedup Strengthens Researchers' Grasp of Neutrons
8/17/15	Quad Chart	Code Speedup Strengthens Researchers' Grasp of Neutrons
8/17/15	Highlight	OLCF Welcomes Whitt, McNally into Leadership Roles
8/17/15	PPT Slide	OLCF Welcomes Whitt, McNally into Leadership Roles
8/17/15	Highlight	Oak Ridge Staff Lead OpenACC, OpenMP Meetings
8/17/15	PPT Slide	Oak Ridge Staff Lead OpenACC, OpenMP Meetings
8/17/15	Highlight	OLCF Team Members Play Key Roles at DOE CSGF Annual Program Review
8/17/15	PPT Slide	OLCF Team Members Play Key Roles at DOE CSGF Annual Program Review
8/17/15	Highlight	Wells Talks Titan Big and Small at ISC '15
8/17/15	PPT Slide	Wells Talks Titan Big and Small at ISC '15
9/7/15	Highlight	Researchers Mine Information from Next-Generation Subsurface Flow Simulations
9/7/15	Quad Chart	Researchers Mine Information from Next-Generation Subsurface Flow Simulations
9/7/15	Highlight	Tiny Titan on Stage at DOE National Science Bowl
9/7/15	PPT Slide	Tiny Titan on Stage at DOE National Science Bowl
9/7/15	Highlight	New DDNTool Streamlines OLCF File System Monitoring
9/7/15	PPT Slide	New DDNTool Streamlines OLCF File System Monitoring
9/7/15	Highlight	Next-Generation Scientists Work on Next-Generation Supercomputers
9/7/15	PPT Slide	Next-Generation Scientists Work on Next-Generation Supercomputers
9/28/15	Highlight	Titan Helps Unpuzzle Decades-Old Plutonium Perplexities
9/28/15	Quad Chart	Titan Helps Unpuzzle Decades-Old Plutonium Perplexities
9/28/15	Highlight	Campbell to Receive "40 Under 40" Safety Award
9/28/15	PPT Slide	Campbell to Receive "40 Under 40" Safety Award
9/28/15	Highlight	On the Road to Summit
9/28/15	PPT Slide	On the Road to Summit
9/28/15	Highlight	Oak Ridge Computational Scientist Joins Argonne Reps to Discuss Current, Future LCF Resources
9/28/15	PPT Slide	Oak Ridge Computational Scientist Joins Argonne Reps to Discuss Current, Future LCF Resources
10/2/15	Poster	SC15 Poster 1
10/2/15	Poster	SC15 Poster 2
10/13/15	Highlight	Flowing Toward Red Blood Cell Breakthroughs
10/13/15	Quad Chart	Flowing Toward Red Blood Cell Breakthroughs
10/13/15	Highlight	Nanoelectronics Researchers Employ Titan for an Electrifying Simulation Speedup
10/13/15	Quad Chart	Nanoelectronics Researchers Employ Titan for an Electrifying Simulation Speedup
10/13/15	Highlight	Retrofitting Rhea
10/13/15	PPT Slide	Retrofitting Rhea
10/13/15	Highlight	Moab Scheduling Tweak Tightens Titan's Workload
10/13/15	PPT Slide	Moab Scheduling Tweak Tightens Titan's Workload

Date	Type of product	Title
10/13/15	Highlight	Group Leader Tapped for New Advanced Data and Workflows Group
10/13/15	PPT Slide	Group Leader Tapped for New Advanced Data and Workflows Group
11/1/15	Fact Sheets	INCITE Fact Sheets
11/10/15	Highlight	Titan Takes on the Big One
11/10/15	Quad Chart	Titan Takes on the Big One
11/10/15	Highlight	Warmer Water to Cool Future OLCF Supercomputers
11/10/15	PPT Slide	Warmer Water to Cool Future OLCF Supercomputers
11/10/15	Highlight	ORNL Researchers Hail Promise of Data-Driven Materials Discovery and Design
11/10/15	PPT Slide	ORNL Researchers Hail Promise of Data-Driven Materials Discovery and Design
11/10/15	Highlight	ORNL Engages UT Students in Data Center Design and Management
11/10/15	PPT Slide	ORNL Engages UT Students in Data Center Design and Management
11/10/15	Highlight	OLCF's Straatsma Educates Graduate Students in Summer School Program
11/10/15	PPT Slide	OLCF's Straatsma Educates Graduate Students in Summer School Program
11/13/15	Website	SC15
11/13/15	Web Article	SC15 Event Details
11/13/15	Website	INCITE 2016
11/13/15	News Release	INCITE Grants Awarded to 56 Computational Research Projects
11/18/15	Web Article	HPCwire Recognizes OLCF at SC15
11/18/15	Highlight	Wells to Serve as Instructor at 2015 MRS Fall Meeting and Exhibit
11/18/15	PPT Slide	Wells to Serve as Instructor at 2015 MRS Fall Meeting and Exhibit
11/25/15	Poster	INCITE 2016 Projects Poster
12/8/15	Highlight	ORNL Celebrates Women at Grace Hopper Conference
12/8/15	PPT Slide	ORNL Celebrates Women at Grace Hopper Conference
12/8/15	Highlight	ORNL Helps Transform High-Performance Computing at SC15
12/8/15	PPT Slide	ORNL Helps Transform High-Performance Computing at SC15
12/8/15	Highlight	Titan Helps Researchers Explore Explosive Star Scenarios
12/8/15	Quad Chart	Titan Helps Researchers Explore Explosive Star Scenarios
12/8/15	Highlight	2015 Hackathons Round Out in Knoxville
12/8/15	PPT Slide	2015 Hackathons Round Out in Knoxville

APPENDIX D. BUSINESS RESULTS FORMULAS

2015 Operational Assessment Guidance

Scheduled Availability

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

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

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

Overall Availability

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

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

Mean Time to Interrupt

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

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

Mean Time to Failure

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

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

System Utilization

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

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

**APPENDIX E. DIRECTOR'S DISCRETIONARY PROJECTS ENABLED
(AT ANY POINT) IN CY 2015**

PI	Institution	Allocation	Utilization	Project Name
Michael Bussmann	Helmholtz-Zentrum Dresden-Rossendorf	750,000	944,736	Laser-Wakefield Simulations Using PICONGPU
Vahid Ranjbar	Brookhaven National Laboratory	1,000,000	2,706,212	Spin Tracking for RHIC
Ilya Pogorelov	Tech-X Corporation	1,000,000	46,691	GPU-Accelerated Beam Dynamics Simulations
John Edward Bussoletti	Boeing Company	3,000,000	0	Exploration Of New Computing Technology For Modeling High Lift Systems Of Commercial Aircraft
Luigi Capone	Texas A&M University–Commerce	1,000,000	106,693	HYDRA GPUs Architecture Migration Task I
Ravichandra Srinivasan	Dresser-Rand Company	1,000,000	967,512	Visualization of Tip Injection Phenomena in the Near Stall Regime of a Transonic Fan Stage
Ravichandra Srinivasan	Dresser-Rand Company	2,000,000	23,063,073	Compressible Flow Turbomachinery Optimization: Numerical Tools Advancement
Brian E. Mitchell	GE Global Research	1,500,000	10,770	TACOMA GPU Port
Vittorio Michelassi	General Electric Nuovo Pignone	10,000,000	9,504,564	HIPSTAR-G-01
Gecheng Zha	University of Miami	600,000	3,535,589	Large Eddy Simulation of Base Drag Reduction Using a novel Passive Jet Flow control method
Gabriel Staffelbach	CERFACS	4,000,000	1,470,611	Tolentucoco: Towards LES in Turbines Using Code Compiling
Brian E. Mitchell	GE Global Research	200,000	0	TACOMA GPU Port Readiness
Kumar Srinivasan	Fiat Chrysler Automobiles	1,000,000	90,561	Advanced Shape Optimization Studies to Improve Automotive Aerodynamics
Shailendra Kaushik	General Motors	4,000,000	0	Multidisciplinary Optimization (MDO) of An Automobiles (1) Aerodynamic and thermal Performance and (2) Aerodynamics and Aero-Acoustics
Spencer Sherwin	Imperial College London	400,000	191,432	Pioneering Scale-Resolving Simulations of Flow over Complex Automotive Geometries
Paul Galpin	Ansys Inc.	4,000,000	43,594	ANSYS AIM Fluid Solver Scaling
Jaime Peraire	Massachusetts Institute of Technology	4,000,000	0	Scalable Discontinuous Galerkin Solvers for the Implicit Large-Eddy Simulation of Transitional Flows
Peter Edward Vincent	Imperial College London	2,000,000	153,762	Development and Benchmarking of an In Situ Visualization Pipeline for Next Generation CFD Tools

PI	Institution	Allocation	Utilization	Project Name
Bronson Messer	Oak Ridge National Laboratory	15,000,000	20,410,573	Explosive Nucleosynthesis and Deflagration to Detonation in Type Ia Supernovae
Simon Portegies Zwart	Leiden University	2,000,000	3,174,533	The Fine Structure of the Milky Way Galaxy
Michael Andrew Clark	NVIDIA Corporation	2,000,000	1,663,604	Petascale Cross Correlation
Brant Edward Robertson	University of Arizona	3,000,000	99,234	Scaling the GPU-Enabled Hydrodynamics Code Cholla to the Power of Titan
Eric J Lentz	The University of Tennessee	35,000,000	2,546,309	Examination of Spatial and Network Resolution Core-Collapse Supernova Simulations
Kelly Holley-Bockelmann	Vanderbilt University	1,000,000	958,703	Binary Black Hole Coalescence in Galactic Nuclei
Zhihui Du	Tsinghua University	10,000,000	118	GPU Accelerated Gravitational Wave Source Modeling
Tomasz Plewa	Florida State University	10,000	0	Thermonuclear Turbulent Combustion
Pierre Ocvirk	Universite de Strasbourg	2,000,000	0	Preparing Cosmic Dawn II
Kiran Alapaty	United States Environmental Protection Agency	1,800,000	1,858,392	MPAS–US EPA
Jack L. Ziegler	National Renewable Energy Laboratory	3,000,000	108	Determining FCC Catalyst Residence Time Distributions in Circulating Biofuel Reactors
Bhanu Prasad Rekepalli	The University of Tennessee	1,000,000	0	Developing Highly Scalable Parallel Bioinformatics Applications on Titan
Daniel Jacobson	Oak Ridge National Laboratory	1,000,000	238,011	Scaling Up of Parallelized Ortholog Detection Algorithms for Comparative Genomics of Bacterial Genomes
Manjunath Gorentla Venkata	Oak Ridge National Laboratory	5,100,000	9,700,674	Achieving Near Real-time Imputation in Thousands of Samples using HPC Systems
Hwee Kuan Lee	Oak Ridge National Laboratory	100,000	11	Development of Probabilistic Neural Network for Computer Aided Diagnosis of Cancer Specimens
Jerome Baudry	The University of Tennessee	12,350,000	20,555,799	Massive Ensemble Docking for Drug Toxicity Prediction
Xiaolin Cheng	Oak Ridge National Laboratory	2,000,000	18,703,710	Computational Study of Cellulose Synthase via Enhanced Sampling in High Performance Computing
Nikolay Dokholyan	University of North Carolina at Chapel Hill	500,000	159	Characterization of Structure and Dynamics of Clinically Relevant Membrane Proteins by Means of Molecular Dynamics Simulations
Pratul K Agarwal	Oak Ridge National Laboratory	1,000,000	18,831,139	Conformational Substates in Enzyme Catalytic Efficiency
George Karniadakis	Brown University	1,000,000	727,668	Development of GPU-Accelerated Mesoscale Simulations
Frank Noe	Freie Universtat Berlin	11,000,000	12,485,493	Adaptive Molecular Simulation of the Immunological Synapse

PI	Institution	Allocation	Utilization	Project Name
Jerome Baudry	The University of Tennessee	10,000,000	655,666	Drugging the “Undruggable”
B. Montgomery Pettitt	University of Texas Medical Branch at Galveston	100,000	43,497	Solubility of Peptides and Proteins
Antti-Pekka Hynninen	National Renewable Energy Laboratory	500,000	9,618	High Performance GPU Molecular Dynamics Engine in CHARMM
Martin Zacharias	Technical University of Munich	9,000,000	3,200,211	Molecular Dynamics of Amyloid Formation at Atomic Resolution: Test of Simulation and Boundary Conditions
Ruth Nussinov	National Institutes of Health	1,000,000	2,435,248	Large Scale Simulations of Oncogenic Ras Protein Its Complexes with Effectors and Regulators
Xiaolin Cheng	Oak Ridge National Laboratory	23,500,000	27,287,350	Functional Domains in Asymmetric Membranes Probed with High-performance Simulation and Neutron Scattering
Edward Lyman	University of Delaware	6,000,000	4,695,632	Subdiffusive Signatures of Nanoscale Lipid Bilayer Domains
Jens Meiler	Vanderbilt University	5,000,000	67,050	Structure Determination of Cellulose Synthase Using Neutron Scattering and High Performance Computing
Gerald W. Feigenson	Cornell University	2,000,000	627,871	Proteins Stiffen Membranes: Bending Moduli of Multicomponent Lipid Bilayers and the Effects of Peptides
Jun Fan	City University of Hong Kong	6,000,000	5,320,203	Multi-scale Simulations on Membrane Remodeling by BAR Domain Proteins
Greg Voth	University of Chicago	5,000,000	921,836	Highly Scalable Ultra-Coarse-Grained Molecular Simulations
D. Peter Tieleman	University of Calgary	650,000	37,070	Lipid-lipid Interactions Underlying Lateral Organizations in Cell Membranes
Aleksei Aksimentiev	University of Illinois at Urbana-Champaign	1,000,000	868,348	DNA Nanotechnology
David N. Beratan	Duke University	1,000,000	0	Computer Assisted Diversity Oriented Molecular Library Design
Rajkumar Prabhu	Mississippi State University	3,000,000	3,010,828	Nano-Microscale Neuronal Membrane Simulations for Understanding Traumatic Brain Injury
Joseph Curtis	National Institute of Standards and Technology	500,000	0	Atomistic Modeling of Small-Angle Scattering Experimental Data on HPC Resources via SASSIE-Web
Mitchel John Doktycz	Oak Ridge National Laboratory	5,000,000	0	Multiscale Simulations in Adaptive Biosystems Imaging
Theresa Windus	Iowa State University	500,000	6,421,846	Critical Materials Institute: Separations Science

PI	Institution	Allocation	Utilization	Project Name
William David Laidig	Proctor & Gamble Company	3,000,000	1,200,279	Can Supercomputing Help Mechanistic Understanding of a Novel Catalyst for a Strategic Raw Material?
Rodney J. Bartlett	University of Florida	20,000	9,156,413	Coupled-cluster Studies of the Active Site of the Cytochrome P450 Enzyme
Remco W.A. Havenith	University of Groningen	7,000,000	7,869,092	Computational Modeling of Organic Photovoltaics
Sriraj Srinivasan	Arkema Inc.	2,000,000	1,991,935	High Performance Catalysts in Fluorination of Climate Impacting Fluorogases
Ariana Beste	The University of Tennessee	1,500,000	631,403	Predictive Computational Catalysis
David N. Beratan	Duke University	1,000,000	1,048,209	Toward Infrared Controlled Charge Transfer Reactions in Chemistry and Biology
Jacek Jakowski	The University of Tennessee	3,000,000	1,996,337	Quantum Dynamics of Nano-scale Materials
Karl Andrew Wilkinson	University of Cape Town	4,000,000	16,095	Dynamic Properties of Porous Frameworks upon the Absorption of Gas Molecules
Giriprakash Palanisamy	Oak Ridge National Laboratory	0	0	ARM Archive Large Scale Data Processing System
Galen Mark Shipman	Los Alamos National Laboratory	0	0	ESG
James Joseph Hack	Oak Ridge National Laboratory	23,520,000	0	Ultra High Resolution Global Climate Simulation to Explore and Quantify Predictive Skill for Climate Means, Variability and Extremes
Moetasim Ashfaq	Oak Ridge National Laboratory	4,000,000	4,542,926	Towards the Development of an Integrated Energy-Water Risk Assessment Tool for Probable Maximum Precipitation and Flood
Katherine J. Evans	Oak Ridge National Laboratory	11,200,000	11,788,195	The Computational Climate Science Integrated Allocation
Forrest M Hoffman	Oak Ridge National Laboratory	6,000,000	5,315,314	DOE SciDAC-3 ACES4BGC Partnership Project Gen2ESM Foundry
David Lawrence Hart	University Corporation of Atmospheric Research	200,000	0	NCAR Benchmarking for Climate and Weather Models
Feng He	University of Wisconsin	300,000	39,177	Implications of the Early Anthropogenic Hypothesis
Xiaobiao Xu	Florida State University	3,000,000	490,286	Simulating the Circulation of North Atlantic Ocean at 1/50° resolution
Andrew Salinger	Sandia National Laboratories	1,000,000	1,048,888	Albany/FELIX
Steven John Ghan	Pacific Northwest National Laboratory	20,000	0	ACME Calibration
Francois William Primeau	University of California–Irvine	150,000	0	Implementation and Testing of a Tracer Transport Matrix for the MPAS-O Ocean Model

PI	Institution	Allocation	Utilization	Project Name
William I Gustafson	Pacific Northwest National Laboratory	0	0	The DOE Atmospheric Radiation Measurement Programs LES ARM Symbiotic Simulation and Observation (LASSO) Initialization, Forcing and Multiscale Data Assimilation Program
Valentine Anatharaj	Oak Ridge National Laboratory	10,000	697,234	Provisioning of Climate Data
Roy J. Primus	GE Global Research	500,000	5,805,869	Application of High Performance Computing for Simulating Cycle to Cycle Variation in Dual Fuel Combustion Engines
Vaidyanathan Sankaran	United Technologies Research Center	1,500,000	395,948	Towards Combustor Simulation Using Large Eddy Simulation and Graphical Processing Units
Xiaoyi Li	United Technologies Research Center	1,000,000	29,950	High Fidelity Direct Numerical Simulation of Sprays for Realistic Injection Applications in Industry
Ramanan Sankaran	Oak Ridge National Laboratory	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
Ning Ren	FM Global	0	0	FireFOAM Modeling of Roll Paper Fires and VOF Modeling of Sprinkler Atomization
Venkat E Tangirala	GE Global Research	5,000,000	4,575,971	Unsteady Combustion Processes in a Gas Turbine Combustor
Guillaume Blanquart	California Institute of Technology	3,500,000	3,857,932	PTF High Karlovitz
Jin Yan	GE Global Research	47,000,000	11,289,117	Modelling of Combustion Dynamics in a Gas Turbine
Vaidyanathan Sankaran	United Technologies Research Center	1,000,000	263,884	Higher Fidelity Fire Simulations Using Fire Dynamics Simulator
Stephen Jones	SpaceX Space Exploration Technologies Corp	5,000,000	0	Scaling and Validation of Adaptive-Grid Turbulent Mixing Simulations
Xiaoguang Zhang	Oak Ridge National Laboratory	300,000	1,506,447	A Comprehensive Theoretical/Numerical Tool for Electron Transport in Mesoscale-Heterostructures
Lucas R. Lindsay	US Naval Research Laboratory	300,000	191,312	First Principles Thermal Transport and Thermoelectric Properties of Materials
Yevgeniy Puzyrev	Vanderbilt University	5,000,000	1,631,322	Flexural Phonons and Mechanical Properties of Two-dimensional Materials
Shiwei Zhang	College of William & Mary	6,000,000	9,734,373	Quantum Many-Body Computations of Strongly Correlated Systems
Panchapakesan Ganesh	Oak Ridge National Laboratory	5,000,000	2,828,448	Data Driven Discovery by Design of Energy Materials

PI	Institution	Allocation	Utilization	Project Name
Timo Thonhauser	Wake Forest University	1,000,000	1,276,835	Catalysis and Diffusion in Metal Organic Frameworks
Andreas Glatz	Argonne National Laboratory	1,000,000	4,575,971	OSCon
Jamison Daniel	Oak Ridge National Laboratory	1,000,000	21,558	Cluster-based Visualization of Tera- and Peta-Scale Datasets
Barbara Chapman	University of Houston	219,449	0	A Similarity-Based Analysis Tool for Pattern Derivation and Large Scale Program Restructuring
David Ronald Pugmire	Oak Ridge National Laboratory	8,500,000	9,264,432	SDAV
Adam Simpson	Oak Ridge National Laboratory	100,000	1,336	Supercomputing in the Classroom
Sergey Panitkin	Brookhaven National Laboratory	500,000	13,259,969	Next Generation Workload Management System
Fernanda Schafer Foertter	Oak Ridge National Laboratory	5,000,000	1,338	Developing Scalable Heterogeneous Computing Training Code Examples
Mark Richard Fahey	Oak Ridge National Laboratory	1,000	0	Automatic Library Tracking Database
Suleyman Kocak	Dassault Systemes Simulia Corp	5,000,000	61,859	Performance, Capacity and Scalability Assessment and Improvements for Realistic Simulation Software—Abaqus/Standard and Abaqus/Explicit
John Franklin Harney	Oak Ridge National Laboratory	50,000	7,841,817	Data Intensive Science Incubators
Cyril Zeller	NVIDIA Corporation	1,000,000	13,372,301	CoDesign
Judith C Hill	Oak Ridge National Laboratory	5,000,000	27,541,155	Computational Partnerships
Jason Micah Cope	DataDirect Networks	200,000	1,876	Assessing the Scalability of DataDirect Networks “Iron Monkey” Burst Buffer on Titan
Judith C Hill	Oak Ridge National Laboratory	100,000	125,646	Computational Science Graduate Fellowship Program
Yuji Shinano	Zuse Institute Berlin	5,000,000	6,220,421	ParaSCIP
Mohammed Sourouri	Simula Research Laboratory	405,000	403,647	User-Friendly Programming of GPU Clusters
Terry Ray Jones	Oak Ridge National Laboratory	3,000,000	1,916,790	Colony
Bruce David D’Amora	IBM	500,000	8,533	CORAL Benchmarking
Erik Deumens	University of Florida	5,000,000	226,597	Predicting and Improving the Performance in the Super Instruction Architecture
John A Turner	Oak Ridge National Laboratory	7,500,000	6,686,753	Computational Engineering and Energy Sciences (CEES) Group Projects
Scott A Klasky	Oak Ridge National Laboratory	5,000,000	22,334,693	ADIOS
Terry Ray Jones	Oak Ridge National Laboratory	3,000,000	0	Hobbes: Operating System and Runtime Research for Extreme Scale

PI	Institution	Allocation	Utilization	Project Name
Jakub Kurzak	The University of Tennessee	200,000	150,657	Bench Testing Environment for Automated Software Tuning
Travis Selby Humble	Oak Ridge National Laboratory	3,000,000	33	Quantum Computing Simulations
Rong Tian	Chinese Academy of Sciences Institute of Computing Technology	1,000,000	670,469	Scalable and Fault Tolerant Meshfree/Particle Simulation
Sreenivas Rangan Sukumar	Oak Ridge National Laboratory	0	0	Hadoop on Rhea
Michael Joseph Brim	Oak Ridge National Laboratory	1,500,000	921,105	Lustre-Vision
Manjunath Gorentla Venkata	Oak Ridge National Laboratory	2,000,000	4,678,542	Enhancing Application Performance and Resiliency of Extreme-Scale Systems
Suzanne Theresa Parete-Koon	Oak Ridge National Laboratory	0	0	Data Transfer Working Group
Robert M Patton	Oak Ridge National Laboratory	2,000,000	1,129,897	Scalable Deep Learning Systems for Exascale Data Analysis
Joshua Ryan New	Oak Ridge National Laboratory	2,000,000	0	Big Data Mining for Building Analytics
Dominic Meiser	Tech-X Corporation	3,000,000	1,530,473	TxHPCG
Todd Coleman	University of California–San Diego	500,000	0	Efficient Bayesian Inference Methods with Distributed Convex Optimization
Shantenu Jha	Rutgers University	500,000	156,734	AIMES
Daniel James Quinlan	Lawrence Livermore National Laboratory	500,000	0	D-TEC: Domain-Specific Language Technology
Katherine Yelick	Lawrence Berkeley National Laboratory	30,000,000	0	DEGAS: Dynamic Exascale Global Address Space
Ron Brightwell	Sandia National Laboratories	26,922,896	47	XPRESS Program Environment Testing at Scale
Beverly Ann Sanders	University of Florida	3,000,000	2,833,822	Predicting and Improving the Performance of the Super Instruction Architecture
Mark Christopher Miller	Lawrence Livermore National Laboratory	100,000	0	FASTMath Installation and Portable Performance Testing
Lubomir Riha	IT4Innovations National Supercomputing Center	2,700,000	717	ESPRESO—ExaScale PaRallel FETI SOLver
Kalyan Perumalla	Oak Ridge National Laboratory	3,000,000	331	Cloning for Exascale
Oded Schwartz	The Hebrew University of Jerusalem	60,000	0	Algorithmic Linear Algebra
Manuel Arenaz	Appentra Solutions S. L.	50,000	196	Porting Parallware Tool to Large HPC Installations Including Titan
Matthew Dearing Wolf	Georgia Institute of Technology	2,000,000	1,612	In Situ Analytics Infrastructures
Jing Gong	Kungliga Tekniska Högskolan	1,000,000	39,213	NekBone with GPU-Direct Communication and Optimized OpenACC Directives

PI	Institution	Allocation	Utilization	Project Name
Eric T. Phipps	Sandia National Laboratories	3,000,000	440,606	EQUINOX Embedded Uncertainty Quantification
Catherine Schuman	Oak Ridge National Laboratory	2,000,000	63,681	Scalable Evolutionary Optimization for Designing Networks
David Gutzwiller	Numecca USA Inc.	2,000,000	0	Preparation of the NUMECA FINE/Open CFD Solver for Leadership Computing: Parallel I/O, Profiling, and Acceleration
Eduardo Francisco D'Azevedo	Oak Ridge National Laboratory	50,000	0	Investigation in Task-based Run-time frameworks for Geosciences Application
Vivek Sarkar	Rice University	1,000,000	0	Unified Portable Programming System for Integrating Task Parallelism, Accelerator Parallelism, and Message Passing on Exascale Systems
Alan Gray	University of Edinburgh	3,000,000	0	Performance Portability for Large-scale Scientific Applications
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
William M Tang	Princeton University	8,000,000	11,209,766	Extreme Scale PIC Research on Advanced Architectures
Judy Hill	Oak Ridge National Laboratory	0	0	Workflow Optimization and Processing of Complex Datasets for Off-site Fusion Energy Research
Frank Jenko	University of California–Los Angeles	5,000,000	1,198,676	GENE
Zhihong Lin	University of California–Irvine	5,000,000	0	Fusion SciDAC GSEP Center: Gyrokinetic Simulation of Energetic Particle Turbulence and Transport
Jeff Candy	General Atomics	3,000,000	0	CGYRO Multicore and Accelerator Optimization
David Lindsay Green	Oak Ridge National Laboratory	1,000,000	400,346	Proto-MPEX Simulation
Charlotte Barbier	Oak Ridge National Laboratory	1,000,000	314,813	Large Scale Hydraulic Fracture Simulations
Henri Calandra	Total S.A.	5,000,000	10,363	Advance Computing for Geoscience Applications
Dilip Reddy Patlolla	Oak Ridge National Laboratory	2,000,000	79,849	GPU Accelerated Settlement Detection
Terry Jones	Oak Ridge National Laboratory	3,000,000	288,815	Foresight
David Trebotich	Lawrence Berkeley National Laboratory	1,000,000	0	Mesoscale Simulation of Subsurface Fractured Materials
George Serban Constantinescu	University of Iowa	2,000,000	0	A Critical Assessment of the Performance of a Fully 3-D Nonhydrostatic Flow Solver to Simulate Flood Propagation in Natural Environments and Improve Efficiency of Flood Protection Measures
Thomas J. LeCompte	Argonne National Laboratory	1,000,000	1,546	Extending Highly Parallel HEP Codes to Titan

PI	Institution	Allocation	Utilization	Project Name
Gerald Tuskan	Oak Ridge National Laboratory	500,000	0	Poplar Genome Wide Association Study—BioEnergy Science Center
Mark Sansom	University of Oxford	1,000,000	549,795	The Influenza A Frontier: Simulating the Interactions of a Full-scale Enveloped Virus with a Complex Host Cell Membrane
Arvind Ramanathan	Oak Ridge National Laboratory	5,000,000	1,189,139	Integrating Neutron Scattering Experiment with Atomistic Simulations to Characterize Biophysical Mechanisms of Intrinsically Disordered Proteins
Predrag Krstic	State University of New York at Stony Brook	1,470,000	23	Science of the Plasma–Material Interface at Extreme Conditions
Leonid Zhigilei	University of Virginia	8,820,000	59,577,943	Atomistic Simulations of Laser Interactions with Metals
Byung Hoon Park	Oak Ridge National Laboratory	8,000,000	18,624,010	Accelerating Materials Modeling Loop of Leadership Computing and Spallation Neutron Source
Joshua D Webb	Caterpillar, Inc.	6,000,000	0	GPU Enhancement of Weld Distortion Prediction
Bobby Sumpster	Oak Ridge National Laboratory	8,000,000	15,199,341	Understanding and Manipulating Surface Mediated Interactions
Tao Wei	Lamar University	1,000,000	21,547	Anti-Biofouling Material Design
Joost VandeVondele	ETH Zurich	6,000,000	4,101,873	5th Rung Functionals with CP2K
Xiaoye Sherry Li	Lawrence Berkeley National Laboratory	15,500,000	10,659,830	Enabling Next-Generation Light Source Data Analysis through Massive Parallelism
Simon Billinge	Brookhaven National Laboratory	1,000,000	0	Nanostructure Complex Modeling
Jerzy Bernholc	North Carolina State University	3,148,800	2,934,104	High Performance Simulations of Electron Spin Distribution and Dynamics in Low-Dimensional Materials
Yong Han	University of Utah	1,000,000	3,256,780	Catalysis and Diffusion in Metal Organic Frameworks
Michael J. Demkowicz	Massachusetts Institute of Technology	3,000,000	0	First-Principles Molecular Dynamics Modeling of Knock-on Damage in Amorphous Silicon Oxycarbide
Peter Coveney	Yale University	150,000	18,245	Benchmarking LAMMPS for Composite Layered Materials
Maarten De Jong	University of California–Berkeley	2,000,000	0	Charting the Complete Elastic Properties of All Known Inorganic Compounds: An Exploratory Study
Sungkwang Mun	Mississippi State University	100,000	50,484	Molecular Dynamic Simulation for Polymer Deformation
Dongwon Shin	Northwestern University	3,000,000	792,218	High Performance Cast Aluminum Alloys for Next Generation Passenger Vehicle Engines

PI	Institution	Allocation	Utilization	Project Name
Trung Dac Nguyen	Vietnam Academy of Science and Technology	1,000,000	28,024	Implementing and Optimizing GPU-Accelerated Molecular Dynamics Models
Balasubramaniam Radhakrishnan	Oak Ridge National Laboratory	1,000,000	0	Micromagnetics Simulations for Magnetic Materials in Electric Motors
David Michael Eike	Procter & Gamble, Inc	2,000,000	1,160,317	Large Size and Time Scale Investigation of Surfactant Aggregation and Structure Formation
Efthimios Kaxiras	Harvard University	0	0	Catalyst Screening and Machine Learning for Catalytic Oxidation of Methanol on Metal Substrates
Georgia Tourassi	Oak Ridge National Laboratory	1,000,000	279,881	High Performance Infodemiology Architecture for Knowledge Extraction from Unstructured Health Text
Scott Christley	University of Chicago	5,000,000	12,128	HPC Lung Model of Pulmonary Acinus
Robert Chase Cockrell	Iowa State University	4,000,000	0	Examining the Pathogenesis of Pouchitis Using a Multiscale Model of the Intestinal Mucosa: Spatially Explicit General-purpose Model of Enteric Tissue
Predrag S Krstic	State University of New York at Stony Brook	3,000,000	2,390,518	Gas-Liquid-Solid Interfaces for Energy Applications
David J Keffer	The University of Tennessee	1,000,000	173,903	Computational Study of Novel Lignin-Derived Carbon Composite Li-ion anodes
Uwe Thumm	Kansas State University	500,000	0	Ab-initio Calculations for Laser-Assisted XUV Double Ionization of Helium
Klaus Bartschat	Drake University	5,000,000	1,598,577	Computational Studies of the Interaction of Time-Dependent Electromagnetic Fields and Charged Particles with Atoms and Molecules
Bamin Khomami	University of Tennessee	1,000,000	0	Multiscale Modeling
Mathieu Luisier	ETH Zurich	17,000,000	18,462,623	Breaking the 10 PFlop/s Barrier with ab-initio Quantum Transport Simulations
Sefa Dag	GlobalFoundries	2,000,000	499,275	Computational Design and Optimization of Nano Device Architectures Based on Ab Initio Quantum Approaches
Benoit Forget	Massachusetts Institute of Technology	6,000,000	3,229,888	OpenMC/TH
John A Turner, Ashley D Barker	Oak Ridge National Laboratory	0	308	CASL Student Workshop
Christophe Calvin	Commissariat à L'énergie Atomique et aux énergies Alternatives	1,500,000	453,754	NMC: New Monte Carlo

PI	Institution	Allocation	Utilization	Project Name
William Bryant Bird	Tennessee Valley Authority	1,000,000	1,021,100	TVA CASL Test Stand—Evaluation of Lower Plenum Flow Anomaly using VERA/Hydra-TH
Guillermo Ivan Maldonado	University of Tennessee	1,200,000	17,229,449	Fuel/Core Heavy Metal Design Optimization to Improve Performance/Safety of LSCR
Kenneth Read	The University of Tennessee	347,000	32,762	Probing Fluctuating Initial Conditions of Heavy-Ion Collisions
Robert Varner	Oak Ridge National Laboratory	0	0	Majorama Demonstrator Secondary Data Archive
Jirina Rikovska Stone	The University of Tennessee	8,000,000	8,319,971	Phase Transitions in High Density Matter in Neutron Stars and Supernovae
Aurel Bulgac	University of Washington	25,000,000	16,113,248	Fission Fragment Yields within Time-Dependent Density Functional Theory
Charles Horowitz	Indiana University	2,000,000	922,787	Nuclear Pasta
Andrey Vladimirovich Zakirov	Keldysh Institute of Applied Mathematics	1,000,000	0	High-Performance Wave Modeling in Nanooptics
Hunchuan Peng	Allen Institute	8,000,000	2,975,631	BigNeuron
Cevdet Noyan	Columbia University	5,000,000	3,717,893	Analysis of Powder X-ray Diffraction (PXRD) Profiles of Polycrystalline Nanopowders in Kinematic Regime
Ping Liu	Brookhaven National Laboratory	800,000	600,486	CO2 Activation through Heterogeneous Catalysis
Misun Min	Argonne National Laboratory	2,500,000	749,947	Nek-HOM (Codes for High Order Methods)
Mark Alexander Jack	Florida Agricultural and Mechanical University	2,000,000	40,491	QRing: A Scalable Parallel Software Tool for Quantum Transport Simulations in Carbon Nanoring Devices Based on NEGF Formalism and a Parallel C++ / MPI / PETSc Algorithm
Cory Hauck	Oak Ridge National Laboratory	2,000,000	1,708,471	Moment Methods for Linear Kinetic Equations
Amitava Bhattacharjee	Princeton University	10,000,000	20,770,485	Magnetic Reconnection and Laboratory Astrophysics with Laser-Produced Plasmas
Eirik Endeve	Oak Ridge National Laboratory	5,000,000	1,185	Fast Algorithms for Multiphasic Modeling of Non-Equilibrium Transport
Wanjian Yin	University of Toledo	3,000,000	803,127	Defect Properties of Solar Cell Materials Using Advanced Quasi-Particle GW Approach
Gabriel Kotliar	Rutgers University	1,000,000	1,085,229	Application of MQSGW+DMFT to Delta Plutonium
Peter Zaspel	Heidelberg University, Ruperto Carola	2,000,000	0	Scalable Numerical Algebra for UQ and CFD
Homer Dewey	Orbital ATK	240,000	0	HERO Scaling Study

PI	Institution	Allocation	Utilization	Project Name
Carlos Federico Lopez	Vanderbilt University	5,000,000	4,335,816	Reducing Unidentifiability in Cell-Signaling Network Models through Multidimensional Analysis and Molecular Simulation
Philip F Locascio	University of Oxford	2,750,000	46,189,269	Molecular Action of the Membrane Bound Dimer Osmolarity Sensor of E.coli, EnvZ
Pui-kuen Yeung	Georgia Institute of Technology	8,000,000	2,031,658	Scale-Similarity and Turbulence Mixing: Schmidt Number Effects and New Algorithmic Developments
Andrew Corrigan	US Department of Defence	100,000	450,399	Benchmarking the Jet Engine Noise Reduction (JENRE) Code on Titan
Jacopo Buongiorno	Massachusetts Institute of Technology	1,000,000	0	Development of a New, High-Fidelity, CFD-Informed Closure Relation for Taylor Bubble Velocity in Slug Flow in Pipes with Horizontal to Vertical Inclinations
Gregory Michael Laskowski	GE Global Research	3,000,000	3,070,797	Adjoint-based Techniques for LES
Duane Lee Rosenberg	SciTec Inc.	2,000,000	5,902,643	Small Scale Statistics and Intermittency in Rotating Strongly Stratified Turbulence: Verification and Connection to Bolgiano-Obokhov Phenomenology
Olivier Desjardins	Cornell University	5,000,000	3,264,361	Large-scale Computation of Particle-Laden and Multiphase Turbulent Flows
Sumanta Acharya	University of Memphis	4,000,000	29,247	Thermally Effective and Efficient Cooling Technologies for Advanced Gas Turbine Systems