

US Department of Energy, Office of Science

**High Performance Computing
Facility Operational Assessment
2012
Oak Ridge Leadership
Computing Facility**

February 2013

Prepared by

**Ashley D. Barker
Arthur S. Bland
David E. Bernholdt
James J. Hack
Douglas L. Hudson
Bronson Messer
James H. Rogers
Kevin G. Thach
Sudharshan S. Vazhkudai
Jack C. Wells
Julia C. White**

US Department of Energy, Office of Science

**HIGH PERFORMANCE COMPUTING FACILITY
OPERATIONAL ASSESSMENT 2012
OAK RIDGE LEADERSHIP COMPUTING FACILITY**

Ashley D. Barker	Bronson Messer
Arthur S. Bland	James H. Rogers
David E. Bernholdt	Kevin G. Thach
James J. Hack	Sudharshan S. Vazhkudai
Douglas L. Hudson	Jack C. Wells
	Julia C. White

February 2013

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

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.

CONTENTS

	Page
LIST OF FIGURES	vi
LIST OF TABLES	vii
EXECUTIVE SUMMARY	ix
Communications with Key Stakeholders	x
Summary of 2012 Metrics	xi
Responses to Recommendations from the 2011 Operational Assessment Review	xiii
1. USER RESULTS	1
1.1 User Results Summary	1
1.2 User Support Metrics: User Satisfaction	3
1.2.1 Overall Satisfaction Rating for the Facility	3
1.2.2 Average Rating across All User Support Questions	4
1.2.3 Improvement on Past Year Unsatisfactory Ratings	5
1.3 Assessing the Effectiveness of the OLCF User Survey	6
1.4 Problem Resolution Metrics	8
1.4.1 Problem Resolution Metric Summary	8
1.5 User Support and Outreach	10
1.5.1 User Support	10
1.5.2 Scientific Computing Liaisons	11
1.5.3 UAO Analysts	13
1.5.4 Training and User Engagement	16
1.5.5 Outreach	19
2. BUSINESS RESULTS	20
2.1 Business Results Summary	20
2.2 Cray XT Compute Partition Summary	20
2.3 Resource Availability	24
2.3.1 Scheduled Availability	24
2.3.2 Overall Availability	25
2.3.3 Mean Time to Interrupt	26
2.3.4 Mean Time to Failure	26
2.4 Resource Utilization	27
2.4.1 Total System Utilization	28
2.5 Capability Utilization	29
2.6 Management of INCITE Selection Process (LCFs only)	31
2.6.1 Process	31
2.6.2 Peer Reviewers	32
2.6.3 Proposal Allocation	33
2.7 Assessing GPU Utilization	34
2.7.1 Identifying the Use of GPU-Specific Libraries with ALTD	34
2.7.2 Accessing Limited NVML Data through Cray's Resource Utilization Reporting	35
2.8 Safety	36

2.9	Internal ORNL Operational Performance Review	36
3.	STRATEGIC RESULTS	37
3.1	Science Output	37
3.2	Scientific Accomplishments.....	38
3.2.1	Understanding Solar Storm/Magnetosphere Interaction	38
3.2.2	How Many Isotopes Can There Be? Simulations Provide an Answer	40
3.2.3	Simulating Paleoclimate Shows Carbon Dioxide Drove Warming at End of Last Ice Age.....	41
3.2.4	Showing Sea-Level Rise Will Continue Even with Aggressive Emission Mitigation	42
3.2.5	Simulations Map Nuclear Pasta in a Core-Collapse Supernova	44
3.2.6	Simulations Lead to Experimental Verification of Bose Glass.....	45
3.2.7	Accelerating Turbomachinery	46
3.2.8	Cutting Costs and Risks in Drug Discovery and Development	48
3.2.9	Providing Modeling and Simulation for CASL: Large-Eddy Simulation Investigation of Grid-to-Rod Fretting.....	49
3.3	Allocation of Facility Director's Reserve	50
3.3.1	Director's Discretionary Program	51
3.3.2	Industrial HPC Partnerships Program	52
3.3.3	2012 Director's Discretionary Allocations	54
4.	INNOVATION	64
4.1	Innovation Summary	64
4.2	Innovations in Publication Tracking	64
4.3	Innovations in Scheduler Technology Transfer	65
4.4	Innovations in Streaming Data Management: Technology Export to the Spallation Neutron Source	65
4.5	Innovations in the Lustre Parallel File System.....	66
4.5.1	Shaping the Lustre Community through Leadership in OpenSFS	66
4.6	Innovations in the HPSS Archival Storage	67
4.7	Innovations in Networking: The Common Communication Interface.....	68
4.8	Innovations in Application Support: Hierarchical Collective Communications Library to Improve Application Performance.....	68
4.9	Innovations in Application Support: On-Site Support for the DDT Debugger and HMPP Compiler	69
4.10	Innovations in Management Knowledge Transfer	69
5.	RISK MANAGEMENT.....	71
5.1	Risk Management Summary	71
5.2	Current Risk Status.....	72
5.3	Major Risks Tracked in the Current Year (2012).....	72
5.3.1	ID# 979 – Insufficient Funding to Meet DOE Commitments (FY2013).....	73
5.3.2	ID# 1006 – Inability to Acquire Sufficient Staff.....	73
5.3.3	ID# 361 – Scientists Decline to Port to Heterogeneous Architecture.....	73
5.3.4	ID# 917 – Robust Support Will Not Be Available to Ensure Portability of Restructured Applications.....	74

5.3.5	ID# 906 – Programming Environment Tools May Be Insufficient	75
5.3.6	ID# 948 – Lack of Infrastructure for an Exascale System	76
5.3.7	ID# 721 – Lustre Metadata Performance Continues to Impact Applications	76
5.3.8	ID# 412 – Inadequate System Availability	77
5.3.9	ID# 992 – Leadership Computing Is Not Achieved (CY2012).....	77
5.3.10	ID# 974 – Insufficient Funding to Meet DOE Commitments (FY2012).....	78
5.3.11	ID# 973/975 – Supply Chain Issues (i.e., Flooding in Thailand) May Impact Disk Drive Acquisitions.....	78
5.4	Risks That Occurred During the Current Year and the Effectiveness of Their Mitigations.....	78
5.4.1	ID# 990 – Titan Could Be Destabilized When Kepler Upgrade Installed	79
5.5	Risks Retired during the Current Year	80
5.5.1	ID# 912 – Lack of Hardware Availability Impacts Porting to Heterogeneous Architecture.....	81
5.5.2	ID# 915 – Upgrade from XT5 to XK6 Takes Too Long, Causing Users to Seek Other Alternatives	81
5.6	Major New or Recharacterized Risks Since Last Review	82
5.7	Major Risks for Next Year	82
6.	SUMMARY OF THE PROPOSED METRIC VALUES	83

LIST OF FIGURES

Figure	Page
1.1 Number of Helpdesk Tickets Issued per Month.	9
1.2 Categorization of Helpdesk Tickets.....	10
2.1 Commissioning Activities Affecting System Availability in 2012.....	21
2.2 2012 XT5 Resource Utilization – Opteron Core Hours by Program.	29
2.3 Effective Scheduling Policy Enables Leadership-Class Usage.	30
2.4 Capability Usage by Project Type.	31
2.5 INCITE Award Decisions Workflow.	32
2.6 2013 INCITE Reviewer Affiliation.	33
2.7 Tracking GPU Usage on Titan.....	35
3.1 Isosurface of Electron Vorticity Showing the Development of Vortex Tubes Along the Top (Magnetospheric Side of the Layer.....	39
3.2 Map of Bound Even-Even Nuclei as a Function of Z and N.	40
3.3 Global Paleoclimate Reconstructure of Temperature and CO ₂ and Jaguar Simulation at OLCF.....	42
3.4 Sea Level Rise through 2300.	43
3.5 Nuclear Matter Density Distributions for Several Different "Pasta Phases."	44
3.6 Phase Diagram in the Magnetic Field (H) vs. Temperature (T) Plane.....	46
3.7 Ramgen Applies Aerodynamic Design Optimization Techniques on Jaguar.....	47
3.8 Experimental and Computer Axial (y-Direction) Time-Averaged Velocities.	49

LIST OF TABLES

Table	Page
1.1 2012 User Result Metrics Summary	2
1.2 OLCF User Support Summary: Metric Targets and Calendar Year Results.....	3
1.3 Satisfaction Rates by Program Type for Key Indicators.....	4
1.4 Overall Satisfaction Results.....	5
1.5 Timeline of Survey Communications and Responses.....	6
1.6 User Survey Participation	7
1.7 User Survey Responders by Program Type	8
1.8 Statistical Analysis of Key Results	8
1.9 Problem Resolution Metric Summary	9
1.10 Application Performance Benchmarking.....	11
1.11 Training Event Summary	17
2.1 OLCF Production Computer Systems, 2012	22
2.2 OLCF HPC System Production Dates, 2008–Present	22
2.3 OLCF Business Results Summary for HPC Systems	23
2.4 OLCF Business Results Summary for HPSS.....	24
2.5 OLCF Business Results Summary for the External Lustre File Systems	24
2.6 OLCF Business Results Summary: Scheduled Availability	25
2.7 OLCF Business Results Summary: Overall Availability.....	26
2.8 OLCF Business Results Summary: Mean Time to Interrupt.....	26
2.9 OLCF Business Results Summary: Mean Time to Failure.....	27
2.10 2012 OLCF System Utilization	28
2.11 OLCF Capability Utilization Definition	30
2.12 OLCF Leadership Usage on the Cray XT and XK Systems.....	30
2.13 2013 Call for Proposals and Awards Activities	32

2.14	INCITE Reviewer Survey Results (>80% Response Rate)	33
2.15	INCITE Submittal and Award Statistics	34
3.1	OLCF Publications	38
3.2	Director's Discretionary Program: Domain Allocation Distribution.....	52
3.3	Director's Discretionary Program: Awards and User Demographics.....	52
3.4	Industry Projects at the OLCF	55
3.5	2012 Director's Discretionary Allocations	56

Executive Summary

HIGH PERFORMANCE COMPUTING FACILITY 2012 OPERATIONAL ASSESSMENT OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2013

EXECUTIVE SUMMARY

Oak Ridge National Laboratory's Leadership Computing Facility (OLCF) continues to deliver the most powerful resources in the United States for open science. At 2.33 petaflops peak performance, the Cray XT Jaguar delivered more than 1.4 billion core hours in calendar year (CY) 2012 to researchers around the world for performing computational simulations relevant to national and energy security; advancing the frontiers of knowledge in physical sciences and areas of biological, medical, environmental, and computer sciences; and providing world-class research facilities for the nation's science enterprise.

OLCF users achieved numerous wide-ranging research accomplishments and technological innovations in 2012. It is not possible to fully summarize the breadth and impact of their productivity within this brief report—OLCF researchers published more than 300 research articles in 2012 alone. Here we identify selected accomplishments that advance the state of the art in science and engineering research and development. The work was recognized through peer-review publications in notable journals such as *Science*, *Nature*, and *The Proceedings of the National Academy of Sciences (PNAS)*, among the many noted in this report.

Exploration of the nuclear landscape carried out by Innovative and Novel Computational Impact on Theory and Experiment (INCITE) researchers and their theoretical prediction of isotopes were featured in *Nature* in 2012 (see Section 3.2.2). On the other end of the physical-scale spectrum, researchers executed a 22,000-year transient Earth System Model simulation that explained the misleading lag of carbon dioxide behind Antarctic temperature records and pointed to carbon dioxide's role in driving global climate change over glacial cycles. This deglaciation study and related work by OLCF users were featured in *Nature* and *PNAS* in 2012 (see Section 3.2.3). Other representative achievements by OLCF users include generation of spectroscopic and photometric data for more accurate distance measurements, necessary for planning future National Science Foundation and Department of Energy (DOE) missions such as the Large Synoptic Survey Telescope and the Palomar Transient Factory; calculations to enable the experimental verification of Bose glass; and screening of 2 million compounds against a targeted receptor in a matter of days, as opposed to the months that would be required for computing clusters, creating a vast library of molecular compounds that can be used for future screenings of potential drug candidates. These and other accomplishments are described in Section 3.

Effective operations of the OLCF play a key role in the scientific missions and accomplishments of its users. This Operational Assessment Report (OAR) delineates the policies, procedures, and innovations implemented by the OLCF to continue delivering a petaflop-scale resource for cutting-edge research. This report covers

CY2012, which, unless otherwise specified, denotes January 1, 2012, through December 31, 2012.

Highlights of OLCF operational activities for 2012 include the following:

- User Support remains a key element of OLCF operations. Due to frequent and detailed communications with users throughout the year, the 2012 user survey showed overall customer satisfaction improved to 4.4/5.0 in 2012 (from 4.1/5.0 in 2011). The center also increased the percentage of tickets that were resolved within three business days, from 89.8% in 2011 to 92.3% in 2012. The OLCF continues to aggressively pursue outreach and training activities to promote awareness—and effective use—of US leadership-class resources (see Section 1), delivering more than 25 unique events, many focused on the new Titan system.
- The OLCF continues to exceed DOE metrics for capability usage (30% target in CY2012, 50.67% delivered). The schedule availability and overall availability target metrics of 85% and 80%, respectively, for Jaguar were exceeded in CY2012 (98.11% and 91.45%, respectively) (see Section 2).

Communications with Key Stakeholders

Communication with the Program Office

The OLCF regularly communicates with the Advanced Scientific Computing Research (ASCR) Program Office through a series of established events. These include weekly Integrated Project Team calls with the local DOE Oak Ridge Office (DOE-ORO) and the Program Office, monthly highlight reports, quarterly reports, the annual Operational Assessment, an annual Budget Deep Dive, and the OLCF annual report. In addition, the DOE-ORO and Program Office have access to tailored web pages that provide system status and other reporting information at any time. Through a team of communications specialists and writers, the OLCF produces a steady flow of reports and highlights for sponsoring agencies, potential users, and the public.

Communication with the User Community

OLCF communications with users take a wide variety of forms and are tailored to the objective, ranging from relating science results to the larger community or instructing users on more efficiently and effectively using OLCF systems. The OLCF offers many training and educational opportunities throughout the year for both current facility users and the next generation of high-performance computing (HPC) users (see Section 1.5).

The impact of OLCF communications is assessed through an annual user survey. Eighty-one percent of respondents rated their overall level of satisfaction with communications from the OLCF as satisfied or very satisfied and submitted an overall rating for communications as 4.0/5.0. The OLCF uses various methods to communicate with users including the following:

- Weekly e-mail message
- Welcome packet
- General e-mail announcements
- Opt-in e-mail notification lists
- Message of the Day

- OLCF website
- Smartphone applications
- Conference calls
- OLCF User Council
- One-on-one interactions through liaisons and analysts
- Social networking vehicles

Survey respondents indicated that the weekly e-mail message was the most useful form of communication.

Communication with the Vendors

The OLCF conducts formal quarterly reviews of projects and operations with Cray, Inc., and NVIDIA. This process includes specific meetings with the product and program managers, correlation of development schedules across hardware and software products, and field demonstrations of emerging equipment. Early involvement is the key to driving design considerations that positively affect emerging products. Supplementing these formal events, the OLCF meets weekly with its Cray Site Advocate and Cray Hardware and Systems Analysts to ensure that there is frequent and consistent communication about known issues, bug tracking, and near-term product development.

The OLCF maintains a robust vendor briefing schedule with other product manufacturers as well, making certain that emerging products targeted to this program are well suited to the high-performance, high-capability, and high-capacity needs of the center.

Communication with Advisory Groups

The OLCF User Council provides a forum for the exchange of ideas and development of recommendations to the OLCF regarding the center's current and future operation and usage policies. The council consists of researchers who have active accounts on the leadership computing facility resources and meets via teleconference on a monthly basis. The current User Council is chaired by Balint Joo of the Thomas Jefferson National Accelerator Facility. The council has been engaged and provided valuable input to OLCF management this past year through activities such as reviewing queuing policy changes and other operational policies affecting users, assisting in gathering survey results, and participating in outreach activities.

Summary of 2012 Metrics

In consultation with the DOE program sponsor, a series of metrics and targets were identified to assess the operational performance of the OLCF in CY2012. The metrics are associated with a series of questions posed to reviewers of the center. The 2012 metrics, target values, and actual results as of December 31, 2012, are summarized below.

Summary of the 2012 Metrics

2012 Metric	2012 Target	2012 Actual
<i>Are the processes for supporting the customers, resolving problems, and communicating with key stakeholders and Outreach effective?</i>		
<i>Customer Metric 1: Customer Satisfaction</i>		
Overall score on the OLCF user survey.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	The OLCF exceeded the metric target: 4.2/5.0.
Improvement on results that scored below satisfactory in the previous period.	Results will show improvement in at least one-half of questions that scored below satisfactory (3.5) in the previous period.	The OLCF exceeded the metric target: No question scored below satisfactory (3.5/5.0) on the 2012 survey.
<i>Customer Metric 2: Problem Resolution</i>		
OLCF survey results related to problem resolution.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	The OLCF exceeded the metric target: 4.5/5.0.
OLCF user problem resolution time period.	Eighty percent of OLCF user problems will be addressed within three business days by either resolving the problem or informing the user how the problem will be resolved.	The OLCF exceeded the metric target: 92.3%.
<i>Customer Metric 3: User Support</i>		
OLCF survey results related to Overall User Assistance and Outreach.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	The OLCF exceeded the metric target: 4.4/5.0.
The OLCF will provide a summary of training events including number of attendees.	Center will hold at least 4 training events.	More than 25 distinct events were held.
<i>Is the facility maximizing the use of its HPC systems and other resources consistent with its mission?</i>		
<i>Business Metric 1: System Availability (For a period of 1 year following a major system upgrade, the targeted scheduled availability is 85% and overall availability is 80%.)</i>		
Scheduled Availability.	Jaguar: 85%; HPSS 95%; External File Systems 95%.	The OLCF exceeded the metric target. Jaguar: 98.11%; HPSS: 98.46%; Widow1: 99.88%; Widow2: 99.81%; Widow3: 98.95%.
Overall Availability.	Jaguar: 80%; HPSS 90%; External File Systems 90%.	The OLCF exceeded the metric target: 91.45% (Jaguar); 98.46% (HPSS); Widow1: 98.25%; Widow2: 98.69%; Widow3: 98.95%.
<i>Business Metric 2: Resource Utilization</i>		
The facility reports total system utilization for each HPC computational system.	Report only; no target.	The OLCF reported 84.39%.
<i>Business Metric 3: Capability Usage</i>		
The OLCF will report on capability usage.	At least 30% of the consumed node hours will be from jobs requesting 20% or more of the available Opteron nodes.	Capability usage was 50.67%. The OLCF exceeded the metric target.

2012 Metric	2012 Target	2012 Actual
<i>Is the facility enabling scientific achievements consistent with the Department of Energy strategic goals 3.1 and/or 3.2?</i>		
<i>Strategic Metric 1: Scientific Output</i>		
The OLCF will report numbers of refereed publications resulting from work done in whole or part on the OLCF systems.	Report only; no target.	The center reported 321 such publications.
<i>Strategic Metric 2: Scientific Accomplishments</i>		
The OLCF will provide a written description of major accomplishments from the users over the previous year.	Target: Descriptions of at least five major accomplishments.	See Section 3 of this OAR.
<i>Strategic Metric 3: Allocation of Facility Director's Reserve Computer Time</i>		
The OLCF will report on how the Facility Director's Discretionary time was allocated, including project title, PI, PI's home organization, processor hours allocated, and usage to date.	Report only; no target.	See Section 3 of this OAR.
<i>What innovations have been implemented that have improved the facility's operations?</i>		
<i>Innovation Metric 1: Infusing Best Practices</i>		
The OLCF will report on new technologies that it has developed and best practices it has implemented and shared.	Report only; no target.	See Section 4 of this OAR.
<i>Innovation Metric 2: Technology Transfer</i>		
The OLCF will report on technologies it has developed that have been adopted by other centers or industry.	Report only; no target.	See Section 4 of this OAR.
<i>Is the facility effectively managing risk?</i>		
<i>Risk Management</i>		
The OLCF will provide a description of major operational risks, including realized or retired risks.	Report only; no target.	See Section 5 of this OAR.

HPSS = High-Performance Storage System.

PI = Principal investigator.

Responses to Recommendations from the 2011 Operational Assessment Review

In February 2012 the OLCF presented the 2011 operational activities of the center to the DOE sponsor and an on-site review committee. Recommendations provided by reviewers, ORNL actions, and DOE ASCR comments and actions are given in the tables below.

1. Are the processes for supporting the user/customers, resolving problems, and outreach effective?

Recommendation	ORNL Action/Comments	HQ Action/Comments
No recommendation.		

2. Is OLCF maximizing the use of its resources consistent with its mission? Is the proposed definition of an allocable unit reasonable for a heterogeneous architecture? Was the 2011 OLCF operations budget reasonable? Is the projected budget for FY [fiscal year] 2012 reasonable?

Recommendation	ORNL Action/Comments	HQ Action/Comments
No recommendation.		

4. Is the OLCF enabling scientific achievements consistent with the Department of Energy Strategic Goal 2 which is to “maintain a vibrant U. S. effort in science and engineering as a cornerstone of our economic prosperity and clear leadership in strategic areas? Is the OLCF’s communication of its scientific impact adequate? Are processes for managing INCITE allocations reasonable? Is the OLCF use of external advisory committees and user groups effective?

Recommendation	ORNL Action/Comments	HQ Action/Comments
Solicit “important Invited talks” from users along with publications and when they submit progress reports.	The OLCF will begin asking principal investigators to self-report invited talks as part of the quarterly report. 2/28/2013: Done.	
Consider adding to the list http://www.olcf.ornl.gov/titan/early-science/ of general (easy access) codes that are known to scale well on the local machine.	The OLCF will expand the list at the cited URL to include a number of additional application codes. This expansion will include both of the following: <ul style="list-style-type: none"> Codes (many of them publicly available) that have already been ported to the center’s TitanDev partition in collaboration with NVIDIA and other colleagues at ORNL and Cray, including NAMD, QMCPACK, CP2K, GTC, CHIMERA, VH-1, SPECFEM-3D, and Chroma. Codes that are being developed on TitanDev by the current user community. The center has made TitanDev generally available to users in the past month (February 2012) and requires monthly reporting of progress and techniques from all those granted access. This reporting mechanism will allow the center to track application maturation on the machine and to quickly determine, in concert with the users doing the work, when an individual code should be listed on the early-science list. <p>Though relatively few truly Open Source codes have been ported to graphics processing unit (GPU)–based clusters, The center will closely monitor the available literature, TitanDev reports, and user quarterly reports to discover any such ports.</p> 2/28/2013: Done.	
Consider collaborating with ALCF [Argonne Leadership Computing Facility] and	The OLCF is collaborating with ORNL’s Computational Data Analytics (CDA) Group in the research and development of text-	

Recommendation	ORNL Action/Comments	HQ Action/Comments
NERSC [National Energy Research Scientific Computing Center] on the 5-year publication tracking tools.	based search tools for the purpose of finding OLCF user papers publicly available online in distributed databases because this task is well aligned with the expertise and capabilities of CDA and will leverage its existing tools and techniques. Once the center demonstrates the basic utility and functionality of these tools, it will share them with ALCF and NERSC and other Office of Science user facilities. 2/28/2013: The search capability was successfully launched in 2012; in 2013 the OLCF plans to continue its development and share the technology with other user facilities.	
We encourage OLCF to include more of their largest-scale examples in future reviews to highlight scientific utilization of the unique capabilities of their resources.	The OLCF will make every effort to include in the Strategic Results sections of the OAR scientific achievements that make effective use of a significant fraction, in most cases 20% or more, of its HPC system, for example, in terms of processor count, aggregate memory, disk storage, or other resources found in more limited fashion on lower-capability-class platforms and/or require very large computer allocations not available at other HPC centers. 2/28/2013: Done.	

4. What innovations have been implemented that have improved OLCF's operations?

Recommendation	ORNL Action/Comments	HQ Action/Comments
No recommendation.		

5. Is OLCF effectively managing risk?

Recommendation	ORNL Action/Comments	HQ Action/Comments
No recommendation.		

6. Are the performance metrics used for the review year and proposed future years sufficient and reasonable for assessing OLCF's Operational performance?

Recommendation	ORNL Action/Comments	HQ Action/Comments
No recommendation.		

7. Other comments, Hybrid Computing?

Recommendation	ORNL Action/Comments	HQ Action/Comments
OLCF should quickly develop one or more methods to measure GPU utilization, deploy on a trial basis to evaluate the efficacy of the measurement approaches,	The OLCF is working directly with Cray and NVIDIA to determine effective methods for assessing use of the Fermi and Kepler GPUs. NVIDIA has historically provided an application programming interface that provides modest information about GPU and GDDR memory usage. Near term, the OLCF	

Recommendation	ORNL Action/Comments	HQ Action/Comments
and implement the most useful measures for reporting in the future, also leading to a capability metric for the use of GPUs.	<p>has arranged a technical interchange as a Birds-of-a-Feather (BoF) at CUG2012 for Cray XK6 customers, Cray, and NVIDIA to discuss the needs of the sites, potential solutions, and timeframes for increasing the fidelity of those measurements.</p> <p>2/28/2013: The OLCF has implemented a rudimentary method for determining, on a per-job basis, whether an application executing on the hybrid compute node has taken advantage of the GPU. The center has worked with Cray, NVIDIA, and the larger user community to define the requirements for accessing the NVIDIA Management Library (NVML) data from the GPU and aggregating this information on a per-job basis through a revised resource utilization software framework. Changes to the NVIDIA device driver, NVML library, and Cray accounting software will be released in 2013. Dissemination of these methods to the community continues, with a scheduled BoF and two presentations at CUG2013.</p>	
Work with NVIDIA and Cray to provide automated tools to measure and collect GPU utilization data.	<p>The OLCF will provide an initial description of the minimum reporting requirements for measuring and collecting GPU utilization data to Cray and NVIDIA and work with these companies to develop, test, and field automated tools that meet these requirements.</p> <p>2/28/2013: Done. Changes to the NVIDIA device driver, NVML library, and Cray accounting software will be released in 2013.</p>	
Implement a standard, defensible way to compare the performance of hybrid systems versus the best conventional systems (e.g., 2xCPU versus CPU+GPU) with equivalently optimized code.	<p>The OLCF believes that the most preferred, defensible approach to comparing the performance of its hybrid architecture is to compare the Cray XK (CPU+GPU) node with the Cray XE6 node (2xCPU), all other characteristics being equal. However, the OLCF does not have a Cray XE6 system, so such comparisons are dependent upon access to another HPC center's resources. Typically, such access is available, and when this is the case, the OLCF will make every effort to produce this preferred benchmark. If the final Titan system has both GPU-accelerated nodes and nonaccelerated ones, the comparison of performance on Cray XK nodes with acceleration to those without acceleration will also be appropriate and necessary for users to judge on which partition of Titan to run a particular application.</p>	

User Results

HIGH PERFORMANCE COMPUTING FACILITY 2012 OPERATIONAL ASSESSMENT OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2013

1. USER RESULTS

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

OLCF RESPONSE: Yes. The Oak Ridge Leadership Computing Facility (OLCF) has a dynamic user support model that is based on continuous improvement, regular assessment, and a strong customer focus. One key element of internal assessment is the annual user survey, developed with input from qualified survey specialists, the OLCF User Council, and the Department of Energy (DOE) Program Manager. In the 2012 survey, OLCF users stated that they are very satisfied with the facility and its services. As part of the survey, users are asked to rate their overall satisfaction with the OLCF and 2012 results once again showed that users are satisfied as evidenced by a mean rating of 4.2. The survey asks users to rate their satisfaction with User Assistance and Outreach and this rating improved from 4.1/5.0 in 2011 to 4.4/5.0 in 2012. 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 met or exceeded all metric targets for user satisfaction in 2012. The center also increased the percentage of tickets that were resolved within three business days, from 89.8% in 2011 to 92.3% in 2012. 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 has developed and implemented a dynamic, integrated customer support model. The 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 comprehensive 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.

The 2012 OLCF user survey was launched on September 4, 2012, and remained open for participation through November 9, 2012. The survey was sent electronically to individuals with active accounts on Innovative and Novel Computational Impact on Theory and Experiment (INCITE), Advanced Scientific Computing Research (ASCR) Leadership Computing Challenge (ALCC), or Director's Discretionary (DD) projects. Three hundred eighty-six users completed the survey out of 1,029 possible

respondents, with an overall response rate of 38%. More active solicitation of responses improved the response rate from 31% in 2011.

Information was collected about the various users, user experience with the OLCF, and OLCF's support capabilities. Attitudes and opinions on the performance, availability, and possible improvements for the OLCF and its staff were also solicited. Data collected from the user survey was analyzed by 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 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 2012 survey can be found on the OLCF website. *

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

Table 1.1. 2012 User Result Metrics Summary

2011 Metric	2011 Actual	2012 Metric	2012 Target	2012 Actual
<i>Customer Metric 1: Customer Satisfaction</i>				
Overall OLCF score on the user survey will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	4.2/5.0	Overall score on the OLCF user survey.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	4.2/5.0
Annual user survey results will show improvement in at least one-half of the questions that scored below satisfactory (3.5) in the previous period.	No question scored below satisfactory (3.5/5.0) on the 2011 survey.	Improvement on results that scored below satisfactory in the previous period.	Results will show improvement in at least one-half of the questions that scored below satisfactory (3.5) in the previous period.	No question scored below satisfactory (3.5/5.0) on the 2012 survey.
<i>Customer Metric 2: Problem Resolution</i>				
N/A	N/A	OLCF survey results related to problem resolution.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	4.5/5.0
Eighty percent of OLCF user problems will be addressed within three working days by either resolving the problem or informing the user how the problem will be resolved.	89.8%	OLCF user problem resolution time period.	Eighty percent of OLCF user problems will be addressed within three business days by either resolving the problem or informing the user how the problem will be resolved.	92.3%

* <http://www.olcf.ornl.gov/media-center/center-reports/2012-outreach-survey/>

Table 1.1. 2012 User Result Metrics Summary (continued)

2011 Metric	2011 Actual	2012 Metric	2012 Target	2012 Actual
<i>Customer Metric 3: User Support</i>				
OLCF will report on survey results related to user support.	The 2011 survey solicits an overall user satisfaction rating and comments about support, services, and resources.	OLCF survey results related to overall user assistance and outreach.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	4.4/5.0

N/A = Not applicable.

1.2 User Support Metrics: User Satisfaction

The operational assessment (OA) metrics for the HPC facility's user support as assessed by the annual user survey are the following:

- Overall satisfaction rating for the facility is satisfactory.
- Average of all user support questions on user surveys is satisfactory.
- Improvement on past year's unsatisfactory ratings occurs as agreed upon with the facility's DOE Program Manager.

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	CY2011		CY2012	
	Target	Actual	Target	Actual
Overall satisfaction rating	3.5/5.0	4.2/5.0	3.5/5.0	4.2/5.0
Average of all user support ratings	3.5/5.0	4.1/5.0	3.5/5.0	4.5/5.0

1.2.1 Overall Satisfaction Rating for the Facility

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

Overall ratings for the OLCF were positive, as 91% reported being satisfied or very satisfied with the OLCF overall. With regard to the degree of overall satisfaction with the center, the percent of satisfied (satisfied and very satisfied) respondents has steadily increased from 2007 (86%) to 2012 (91%). Mean responses were between 4.0/5.0 and 4.2/5.0, in 2012 showing a high degree of satisfaction with the OLCF across project classifications (Table 1.3). The calculated mean in 2012 was 4.2/5.0, which is well above the stated metric of 3.5.

Key indicators from the survey, including overall satisfaction, are shown in Table 1.3. These are summarized and broken out by program.

Table 1.3. Satisfaction Rates by Program Type for Key Indicators

Indicator	Mean	Program		
		INCITE	ALCC	DD
Overall satisfaction with the OLCF	4.2/5.0	4.2/5.0	4.0/5.0	4.0/5.0
Overall experience with User Assistance and Outreach	4.4/5.0	4.5/5.0	4.3/5.0	4.4/5.0
Overall experience with Account Management staff	4.5/5.0	4.6/5.0	4.5/5.0	4.5/5.0
Overall system performance of the Cray XT5	4.2/5.0	4.2/5.0	4.4/5.0	4.2/5.0

1.2.2 Average Rating across All User Support Questions

The calculated mean of all answers to all user support questions on the 2012 survey was 4.5/5.0, indicating that the OLCF exceeded the 2012 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 46% of the survey respondents.

After reviewing the results of the 2011 survey in detail, the OLCF determined that additional questions were warranted in the 2012 survey to solicit better feedback regarding user training. With input from the ORISE survey specialist, the OLCF added questions to solicit improved feedback regarding the quality of the live training classes conducted in 2012, as well as the top choices for the material to be covered in 2013. This information will be used to refine the center's training curriculum in 2013. A brief summary of some of the findings from the 2012 survey is listed below. See Table 1.4 for overall satisfaction results from each of these areas.

User Assistance Evaluation

- For support services used, 69% of the 386 respondents reported using the User Assistance Center (UAC), followed by 29% using the Scientific Computing/Liaison service, and 8% using Visualization.
- Overall satisfaction with the user support services (UAC, Scientific Liaisons, Account Management, and Visualization) provided by the OLCF was high with a mean response of 4.5/5.0.
- Respondents with at least one interaction with the UAC and its staff were asked about the speed of initial contact and quality of the response; a large percentage were satisfied or very satisfied with the initial contact (89%) and with the quality of the response (86%).

Training and Education

- The majority of OLCF users said “yes” or “maybe” (91%) to the prospect of attending future OLCF training events based on their previous experience.
- The number one reason users gave for not participating in any in-person training events was that they do not have the time to attend, followed by budget.
- Documentation was listed as the top choice (75%) for training preference, followed by online training (52%).

- The Crash Course in Supercomputing held June 12–13, 2012, was the highest rated OLCF training event in 2012 with a mean rating of 4.3/5.0. The mean rating across all training classes was 4.0/5.0.
- When presented with a list of training topics, respondents' most frequently requested topic was Graphics Processing Unit (GPU) Programming (62%) followed by Tuning and Optimization (53%), and Advanced Message Passing Interface (MPI) (49%).

OLCF Communications

- Eighty-one percent of respondents (290) rated their overall satisfaction with communications from the OLCF as satisfied or very satisfied.
- Respondents indicated the e-mail message of the week was most useful; Twitter was found to be the least useful communication mechanism.
- A few users commented that it would be useful to have more information on the OLCF website regarding current and upcoming outages, especially the expected duration. User Assistance and Outreach (UAO) will investigate what is needed to meet this request.

OLCF Website

- Ninety-eight percent of respondents indicated that they had visited the <http://www.olcf.ornl.gov> website. Of these users, 33% indicated that they visit the site once a week or more. Only six respondents indicated they had never visited the site.
- The greatest number of respondents indicated being satisfied with the accuracy of information provided. One aspect that users noted could be improved is navigation. The OLCF made changes to the navigation of the support materials at the end of CY2012 and will follow up to see if users are better able to find information on the redesigned site.

Table 1.4. Overall Satisfaction Results

	Overall Satisfaction with User Assistance and Outreach	Overall Satisfaction with OLCF Training Events	Overall Satisfaction with OLCF Communications	Overall Satisfaction with the OLCF Website
Number of Users Who Responded to the Survey	386	386	386	386
Number of Users Who Responded to the Particular Questions	267	65	359	330
Rating	4.4/5.0	4.0 /5.0	4.0/5.0	4.0/5.0

1.2.3 Improvement on Past Year Unsatisfactory Ratings

Each year the OLCF works to show improvement in no less than half of any questions that scored below satisfactory (3.5/5.0) in the previous year's survey. All questions scored above 3.5 on both the 2011 and 2012 surveys.

1.3 Assessing the Effectiveness of the OLCF User Survey

Before sending the survey, the OLCF 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 impact the meaning and utility of the data they would generate. As already discussed in Section 1.2.2, questions were added to the survey to solicit better feedback to assist in the OLCF's continuous improvement efforts.

Several targeted notifications were sent to those eligible to participate in the survey. The initial survey invitation from ORISE was sent on September 4, 2012, and subsequent follow-up reminders were sent by James Hack (National Center for Computational Sciences [NCCS] Division Director), Arthur Bland (OLCF Project Director), Ashley Barker (UAO Group Lead), ORISE, and the OLCF User Council. The survey was advertised on the OLCF website and mentioned in the weekly communications e-mail sent to all users. Survey responses were tracked on a daily basis to identify the effectiveness of the various communication methods (see Table 1.5). As also seen in 2011, the message sent from the OLCF User Council was among one of the most productive in terms of soliciting responses. During the 2012 OA onsite meeting, the panel commented that the OLCF should work to increase survey participation in 2012. At the end of survey period, 386 users completed the survey out of 1,029 possible respondents, with an overall response rate of 38%. The response rate was much improved over last year (31%). We attribute the increase to more frequent reminders and the fact that unlike previous years, we provided a direct link to the survey in our weekly e-mails and on our website.

Table 1.5. Timeline of Survey Communications and Responses

Survey Timeline	Date	Number of Respondents	Percent of responses (N = 1,029)	Description of Reminder
Day 1	September 4	50	4.86%	Initial e-mail invitation sent by ORISE evaluator, Dr. Erin Burr.
Day 2	September 5	20	1.94%	—
Day 3	September 6	55	5.34%	Initial e-mail invitation resent by Dr. Burr due to ORNL security filter issues.
Days 4–8	September 7–11	45	4.36%	—
Days 9–10	September 12–13	12	1.17%	E-mail reminders sent by NCCS UAO Leader, Ashley Barker, to INCITE users (included in regular quarterly report update request e-mail).
Day 11	September 14	5	0.49%	Survey link sent to users in Friday Announcements.
Days 14–15	September 17–18	3	0.29%	—
Day 16	September 19	1	0.10%	Link to survey posted on OLCF.ORN.LGOV.
Day 17	September 20	39	3.79%	E-mail reminder sent from OLCF User Council.
Day 18	September 21	10	0.97%	Survey link sent to users in Friday Announcements.
Days 19–24	September 22–27	10	0.98%	—

Table 1.5. Timeline of Survey Communications and Responses (continued)

Survey Timeline	Date	Number of Respondents	Percent of responses (N = 1,029)	Description of Reminder
Day 25	September 28	7	0.68%	Survey link sent to users in Friday Announcements.
Day 28	October 1	30	2.92%	E-mail reminder sent by ORISE evaluator, Dr. Erin Burr.
Days 29–31	October 2–4	10	0.97%	—
Day 32	October 5	2	0.19%	Survey link sent to users in Friday Announcements.
Day 36	October 9	3	0.29%	—
Day 39	October 12	0	0.00%	Survey link sent to users in Friday Announcements.
Day 42	October 15	21	2.04%	E-mail reminder sent by ORISE evaluator, Dr. Erin Burr.
Days 43–45	October 16–18	16	1.56%	—
Day 46	October 19	4	0.39%	Survey link sent to users in Friday Announcements.
Days 48–49	October 21–22	3	0.29%	—
Day 53	October 26	2	0.19%	Survey link sent to users in Friday Announcements.
Day 56	October 29	1	0.10%	E-mail reminder sent by OLCF Project Director, Arthur Bland, and NCCS Director, Dr. James Hack.
Day 57	October 30	2	0.19%	—
Day 60	November 2	0	0.00%	Survey link sent to users in Friday Announcements.
Day 63	November 5	0	0.00%	Reminder sent to OLCF users at ORNL by OLCF Director of Science, Dr. Jack Wells.
Day 64	November 6	18	1.75%	Final e-mail reminder sent by ORISE evaluator, Dr. Erin Burr.
Day 65	November 7	3	0.29%	—
Day 66	November 8	13	1.26%	Final e-mail reminder sent by Ashley Barker.
Day 67	November 9	1	0.10%	Survey closed by ORISE evaluator, Dr. Erin Burr.

The OLCF has a relatively equally balanced distribution of users in terms of their length of time using the systems (Table 1.6).

Table 1.6. User Survey Participation

	2011 Survey	2012 Survey
Total Number of Respondents (Total Percentage Responding to Survey)	252 (31%)	386 (38%)
New Users (OLCF User <1 Year)	31%	33%
OLCF User 1–2 Years	30%	25%
OLCF User >2 Years	39%	42%

Survey respondents were asked to classify the program types with which they were affiliated (Table 1.7).

Table 1.7. User Survey Responders by Program Type

Program	Response Rate*
INCITE	69%
DD	25%
ALCC	14%
Other	6%

* Total is greater than 100% because survey respondents can be associated with more than one type of project.

Statistical Analysis of the Results

Statistical analysis of four key survey areas is shown in Table 1.8. These reflect overall facility satisfaction, services, and computational resources.

Table 1.8. Statistical Analysis of Key Results

	Overall Satisfaction	Overall Experience with User Assistance and Outreach	Effectiveness of Problem Resolution	Overall System Performance of the Cray XT5
Number of Survey Respondents	386	386	386	386
Number of Respondents to This Specific Question	353	267	268	330
Mean	4.2/5.0	4.2/5.0	4.3/5.0	4.2/5.0
Variance*	0.48	0.89	1.00	0.48
Standard Deviation*	0.75	0.95	1.00	0.69

*The OLCF examined the variance and standard deviation for several key questions and found them to be within acceptable parameters. We noted a high variance for the rating of “Effectiveness of Problem Resolution,” which received a mean rating of 4.3 but had a variance of 1.0. The larger variance is due in part to the lower response rate for this particular question; however, we also saw a significant range of expectations for problem resolution. While a large number of the 268 respondents were satisfied with the effectiveness of problem resolution, about 6% were not satisfied and requested, for example, after-hours support and faster response times.

1.4 Problem Resolution Metrics

The OA metrics for problem resolution are the following:

- Average satisfaction ratings for problem resolution related questions on the user survey 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 three business days.

1.4.1 Problem Resolution Metric Summary

In the majority of instances, the OLCF can resolve the reported problem directly, which includes identification and execution of the necessary corrective actions such that the problem is resolved from the users’ perspective. Occasionally User Assistance receives problem reports for which its ability to resolve the root cause of the issue is limited due to factors beyond its control. In this scenario “addressing the problem” requires that User Assistance has identified and carried 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 for User Assistance might be to recreate the issue; open a bug ticket with the product vendor; provide the

vendor necessary information about the issue; and then provide a workaround to the user, if possible.

The OLCF uses request tracker (RT) 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. In addition, the software collates statistics on tickets issued, turnaround times, etc., to produce weekly reports, allowing the OLCF staff to track patterns and address anomalous behaviors before they have an impact on additional users. The OLCF issued 2,323 tickets in response to user queries for CY2012 (Figure 1.1). The center exceeded the problem resolution metric and responded to 92.3% of these queries within three business days (Table 1.9).





Figure 1.1. Number of Helpdesk Tickets Issued per Month.

Table 1.9. Problem Resolution Metric Summary

Survey Area	CY2011		CY2012	
	Target	Actual	Target	Actual
Percent of problems addressed in three business days	80%	89.8%	80%	92.3%
Average of problem resolution ratings	3.5/5.0	4.2/5.0	3.5/5.0	4.4/5.0

Each ticket is assigned to one user assistance or account analyst, who establishes customer contact and tracks the query from first report to final resolution, providing not just fast service, but also service tailored to each customer’s needs. While UA0 is dedicated to addressing queries promptly, user assistance and account analysts consistently strive to reach the “right” or best solution rather than merely a quick turnaround. Tickets are categorized by their most common types. The top reported problem in 2012 (as well as 2011) was related to jobs/batch queues (Figure 1.2).



CY 2012
(2,323 New Inquiries)

Figure 1.2. Categorization of Helpdesk Tickets.

1.5 User Support and Outreach

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

- Anecdotal evidence confirms in-depth collaborations between facility staff and the user community.
- Summary of the training events conducted during this period is provided.

The OLCF recognizes there are four pillars of user support and outreach. The first is user support staff made up of account management liaisons, UAO analysts, and Scientific Computing Group (SciComp) liaisons. The second is multiple vehicles to communicate with users, sponsors, and vendors. The third is developing and delivering training to current and future users. And the last is the strong outreach component needed to interface with the next generation of HPC users, the external media, and the public. This section discusses key activities and contributions for all four areas.

1.5.1 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 two complementary OLCF user support vehicles: UAO and SciComp, which includes the scientific and visualization liaisons. Scientific liaisons are a unique OLCF response to high-performance scientific computing problems faced by users.

1.5.2 Scientific Computing Liaisons

The OLCF pioneered a total user support model widely recognized as a best practice for HPC centers: the SciComp liaison program, comprising experts in their scientific discipline, including PhD-level researchers, who are also specialists in developing code and optimizing HPC systems. Support ranges from basic support—access to computing resources—to complex, multifaceted support for algorithm development and performance improvement. Scientific liaisons support the research focus of projects, while visualization liaisons frequently find themselves developing custom software and algorithms to address unique user challenges in data analysis. Selected examples of the in-depth collaborations between OLCF staff and the user community are described here.

CAAR Continues to Deliver Results

The OLCF teams with users to exploit hierarchical parallelism within applications to better map to next-generation architectures. To this end, the OLCF established the Center for Accelerated Application Readiness (CAAR). For a detailed description of CAAR, reference the 2011 OLCF Operational Assessment Report (OAR), Section 4.1.[†]

SciComp staff members worked assiduously with application teams throughout 2012, leading to application performance improvements of 1.4x to 3.3x on GPU-based systems (see Table 1.10).

Table 1.10. Application Performance Benchmarking

	XK6 (w/ GPU) versus XE6	Cray XK6: Fermi GPU plus Interlagos CPU Cray XE6: Dual Interlagos and No GPU
Application	Performance Ratio	Comment
S3D	1.4	Turbulent combustion 6% of Jaguar workload
Denovo	3.3	3D neutron transport for nuclear reactors 2% of Jaguar workload
LAMMPS	3.2	High-performance molecular dynamics 1% of Jaguar workload
WS-LSMS	1.6	Statistical mechanics of magnetic materials 2% of Jaguar workload 2009 Gordon Bell Winner
CAM-SE	1.5	Community atmosphere model with SE dycore 1% of Jaguar workload

To illustrate the type of CAAR activities carried out by the OLCF staff, we provide a representative example: SciComp liaison Mike Brown performed research and development on methods to improve performance for the widely used LAMMPS molecular dynamics (MD) package. Improvements Brown implemented included methods to accelerate parallel particle-particle particle-mesh calculations, improved algorithms for neighbor list calculation, improved algorithms for MPI process mapping for Cartesian topologies, acceleration for new force-fields needed by materials science investigators, and new methods for calculating electrostatics on hybrid supercomputers. Brown adapted methods for topology-aware partitioning of k-

[†] <http://info.ornl.gov/sites/publications/Files/Pub35226.pdf>

space and real-space calculations in MD for use with accelerators. The research leading to these improvements has been published and all of the improvements have been incorporated into the main LAMMPS distribution. In addition to the CAAR efforts, Brown and other SciComp staff work with numerous projects at the OLCF.

How Many Isotopes Can There Be? Simulations Provide an Answer

The INCITE project led by James Vary, Iowa State University, this year produced a *Nature* article on mapping the nuclear landscape[‡] and significant impact to experiment in calculating oxygen-23.[§] This and other significant scientific achievements are a result of close collaboration between the OLCF and domain scientists. The SciComp liaison to this project, Hai Ah Nam, made significant contributions toward code development (debugging difficult race conditions, incompatibility between FORTRAN derived types and OpenMP) and is a co-author on a recent paper from this team.

Accelerating Turbomachinery

SciComp computer scientist Norbert Podhorszki worked with Mathieu Gontier from NUMECA International to implement an efficient I/O module for the FINE/Turbo solver for computational fluid dynamics (CFD) applications. Ramgen Power Systems, LLC, used the new solver on Jaguar, with a tenfold increase in checkpoint/restart performance. The unusual characteristic of the solver is such that, for load balancing reasons, each processor holds multiple and a variable number of pieces of each variable because of a nonuniform balancing of the structured model over the processes. This renders all traditional I/O solutions very inefficient, especially their original host-slaves I/O approach based on CFD General Notation System (CGNS) on a very large number of sub domains, and which becomes the bottleneck at scale. Ramgen significantly advanced its shock-wave-based compression aerodynamic design process and revealed designs that exhibit valuable new aerodynamic characteristics by working with OLCF. Ramgen also ran a 3.7 billion grid cell resolution two-body test case simulation, which would not have been possible without the Adaptable I/O System (ADIOS) implementation (see Section 3.2.8).

Simulations Map Nuclear Pasta in A Core-Collapse Supernova

The behavior of matter at very high densities (i.e. at densities greater than that of an atomic nucleus) is an important ingredient in understanding the birth of neutron stars in core-collapse supernovae. Load imbalance—both in runtime per task and memory footprint per task—was crippling progress with the Hartree-Fock Equation of State (HF-EOS) code in the INCITE project led by Anthony Mezzacappa, ORNL. In fact, many runs would not complete with the code if the sampled phase space led to the node memory being exhausted. OLCF staff implemented a new load-balancing scheme in the HF-EOS code, based on a master-slave algorithm, ensuring all runs would complete and improving overall performance of the code by roughly 2x.

In addition to doubling code performance for the Mezzacappa project, visualization SciComp team member Ross Toedte and OLCF postdoc Chaoli Wang did much of the programming on a custom visualization tool that rapidly rendered high-quality volume visualizations of low resolution computational data. Because of the complexity of the equation of state, executing a large number of simulations was the only way to

[‡] <http://www.olcf.ornl.gov/2012/06/27/ornlutk-team-maps-the-nuclear-landscape/>

[§] J. Erler, N. Birge, M. Kortelainen, W. Nazarewicz, E. Olsen, A.M. Perhac, and M. Stoitsov, “The limits of the nuclear landscape,” *Nature* **486**, 509 (2012).

understand the morphologies among the different pasta phases in very small proton and neutron volumes. Typically, volume visualizations are computationally intensive yet difficult to generate from sparse or low-resolution data. Chaoli's tool not only produced high-quality results that facilitated visual understanding of the transitions between different pasta phases, but it also used a novel data structure that sped production of the visualizations by at least an order of magnitude. This work resulted in a 2012 publication in *Physical Review Letters* (see Section 3.2.6).

Scaling the Bridge between Quantum and Classical Turbulence

SciComp computer scientist Norbert Podhorszki and ORNL researcher Qing Gary Liu worked with Min Soe from Rogers State University to improve the I/O of the QLG2Q code through the application of ADIOS. ADIOS checkpoint/restart achieved a data bandwidth of 43GB/s on Jaguar, overcoming the I/O bottleneck and enabling scaling beyond 30k cores. QLG2Q is a mesoscopic unitary algorithm code to study quantum turbulence. The DD project led by George Vahala, The College of William and Mary, examines the bridge between quantum and classical turbulence. Related to lattice Boltzmann, unitary collision and streaming operators are used and propagate the wave-function information on lattice grids.

Other Sample Activities

It is impossible to enumerate all of the SciComp support activities. In addition to those highlighted above, other sample activities include the following.

- Collaborating with investigators at the University of Connecticut to use GPU acceleration for their studies of polyelectrolyte brushes.
- Performing the first simulations of liquid copper on graphite at the same time and size scales as the Center for Nanophase Materials Sciences experiments. The simulations identified a manufacturing issue with the experiments (copper lines smaller than they had intended).
- Software development and performance tuning of the RAPTOR-LES large eddy simulation code for the INCITE project led by Joseph Oefelein, Sandia National Laboratories.
- Building, testing and supporting SIERRA continuum mechanics suite on TitanDev for one of the allocation project teams.

1.5.3 UAO Analysts

As already discussed in the problem resolution section, UAO analysts are responsible for addressing user queries. Some of the most common UAO activities include the following:

- Enabling access to OLCF resources
- Helping users compile and debug large science and engineering applications
- Identifying and resolving system-level bugs in conjunction with other technical staff and vendors
- Installing third-party applications and providing documentation for usage
- Engaging other OLCF staff to ensure users have up-to-date information about OLCF resources and to solicit feedback
- Researching, developing, and maintaining reference and training materials for users

- Communicating with users
- Developing and delivering training
- Acting as user advocates

Preparing for Titan

UAO focused its efforts in 2012 on preparing for Titan. In response to some of the survey feedback received on the 2011 survey and making way for new Titan documentation, the OLCF made comprehensive changes to the design of the OLCF support website.** The changes allowed the team to better accommodate new content for the Titan upgrade and present existing support content in a format that is simpler to navigate. The changes were completed in November 2012, and initial feedback on the new support site has been very positive, as evidenced by the following comments:

“The documentation was an incredible resource and my thanks go out to whoever wrote it.”

“In general, I have found most of the information that I have needed in the user’s guides; they are very well put together.”

New System User Guides Section

Prior to 2012 support content was primarily contained in individual knowledge-base articles. Articles in the knowledge base are searchable and simple to maintain, but end-users expressed a desire for a single, definitive source of information for any given OLCF system. In March 2012 the OLCF implemented new code within its website content management system to present a curated subset of articles as a single system-user guide for an individual OLCF system. This technique leverages existing content (thereby reducing duplication of effort) and retains the searchability of the individual articles for those users who prefer that approach while enabling the creation of a single definitive guide with a logical flow for those users who prefer a narrative structure. Currently the OLCF has a system-user guide for its flagship supercomputer, Titan, and its analysis and visualization cluster, Lens.

New Programming Tutorials Section

With new challenges associated with GPU programming on the horizon for Titan end-users, there was a need for more in-depth support information deemed too detailed for inclusion within a system-user guide. In October 2012 the OLCF created a new section of the support website for programming tutorials to house such content. Most of these dozen or so tutorials present users with detailed examples illustrating how to transform non-accelerated code into accelerated code using a variety of state-of-the-art techniques. Tutorials on a range of different subjects are planned for the future.

Updated Software Section

Prior to 2012 the software section of the OLCF support website contained a long system-specific list of software installed at the OLCF. Users commented that the list was a bit overwhelming and difficult to search. Additionally, software installed on multiple systems had multiple knowledge-base articles, which made maintenance of the section difficult. In October 2012 the OLCF debuted a new software section that is organized by software category (e.g., compilers, data management) and provides

** <http://www.olcf.ornl.gov/support/>

more content for each installed software package. Each package has one article that serves as the single point of reference. Usage support information about the software package is combined with version and system availability information to create the final article displayed on the website. The process of publishing these software articles is integrated with the tools the center uses to maintain the software, ensuring changes to the software stack are automatically reflected on the website.

Updated Training Events Section

Before 2012 training events were listed on a calendar with links to individual event pages. The event pages were not prominently featured on the support website, and archived material was not always easy to find and access. In October 2012 an updated training-events section was added to the main support page to simplify the retrieval of archived training presentations and materials. Training events are all listed in one location and separated into upcoming and past events. For upcoming events, information about registration and accommodations is included. For past events, presentation slides and other support material are posted.

New Policies Section

Until 2012 OLCF policies were scattered throughout support documentation. The new policies section of the site consolidates all OLCF policies into a single policy guide for easy reference. Each policy is clearly designated and includes version information and a list of the types of end-users affected.

User Access Improvement

The OLCF participated in an internal Operational Performance Review (OPR) in 2012. Based on user interaction, additional survey feedback, and the center's observations, the OLCF felt that user access was an area that should be assessed as part of the OPR. During the OPR, the OLCF evaluated the entire process for granting a project and/or user access to the facility. The center focused on a few key tasks including export control, identity proofing, and user agreements.

One of the operational improvements that resulted from the OPR was a change in how Oak Ridge National Laboratory (ORNL) conducts export control reviews for OLCF projects. At the time of the review, the OLCF was experiencing long delays in receiving export control reviews for new or renewing projects. OLCF representatives met with staff from the ORNL export control office and identified two changes to the process for improved turnaround time. The first was that renewed projects that do not experience a change in scope would not require a new review, decreasing the number of reviews needed. The second change was to invite the ORNL Export Control Officer to the weekly Resource Utilization Council meeting so that he could become more familiar with the project as it was discussed. This additional familiarity saved the export control reviewer time when writing the review and gave him the background necessary to perform a more informed review. Both changes have led to a much improved turnaround time for export control reviews.

The second item investigated was Level 2 identity proofing. Level 2 identity proofing procedures are implemented to assurance confidence in the asserted identity's validity. For users local to ORNL, a member of the Accounts Management team can do this proofing in person. For the majority of the users who are not local, this process is done through the use of notary publics, who verify the serial number of the RSA token and photo identification of users. The users must then send the original notarized form back to the OLCF (no faxed copies are accepted) before the RSA token is activated. Many of the center's international users find it challenging to

locate a notary public. As a result of the review, the OLCF determined that Skype-like technology can be used to verify a user's identity. This technology is now available to users who cannot easily access a notary public for identity proofing.

The third item investigated the ability to quickly execute a user agreement. There are times when a project needs rapid access to the OLCF for a short period of time, as in the case of the center's "INCITE Prep" projects, which request access to run benchmarks for their INCITE proposals. The OLCF is working with the ORNL Partnerships Directorate to see if it would be possible to implement a rapid-access version of the user agreement that could be executed by the user rather than the user's institution for cases where rapid access is needed for a short period of time.

1.5.4 Training and User Engagement

Workshops and seminars are other important components of the customer support model. They provide an additional opportunity to communicate with users and serve as vehicles to reach out to the next generation of HPC users. In 2012 the OLCF hosted or participated in several activities in the form of educational outreach, training, and user engagement. A summary of these events is shown in Table 1.11.

The OLCF continued to focus on preparing users for Titan and held three separate workshops during the year. A total of 270 people participated in these events. At the request of the users and the OLCF User Council, the center continues to incorporate more hands-on exercises into the training classes so users get a chance to put into practice what they learn during the day.

In addition to training events, the OLCF gathered top experts in science, engineering, and computing from around the world to discuss research advances that are now possible with extreme-scale hybrid supercomputers at the Accelerating Computational Science Symposium 2012 (ACSS 2012) held March 28–30, 2012, in Washington, DC. The conference was cohosted by the OLCF, National Center for Supercomputing Applications and the Swiss National Supercomputing Center. Attendees explored how hybrid supercomputers speed discoveries, such as deeper understanding of phenomena from earthquakes to supernovas, and innovations, such as next-generation catalysts, materials, engines, and reactors. Presenters, including several OLCF users, shared recent advances enabled by hybrid supercomputers in chemistry, combustion, biology, nuclear fusion and fission, seismology, and other fields. They also discussed the new breadth and scope of research that will be possible as petascale systems continue to increase in computational performance.

The OLCF, along with the University of Tennessee's Joint Institute for Computational Sciences cohosted a workshop, "Electronic Structure Calculation Methods on Accelerators," at ORNL February 5–8, 2012 to bring together researchers, computational scientists, and industry developers. The participants attended presentations and training sessions on the advances and opportunities that accelerators—dedicated, massively parallel hardware capable of performing certain limited functions faster than central processing units (CPUs)—bring to HPC. The workshop's purpose was to respond to the challenge of using innovative accelerator hardware and multicore chips in electronic-structure-theory research, which investigates the atomic structure and electronic properties of materials.

Table 1.11. Training Event Summary

Event Type	Event Description	Event Date	Participants
Educational Outreach	Southeast Conference for Women in Physics*	January 12–13, 2013	114
Workshop	Titan Workshop	January 23–27, 2013	93
User Conference Call	User Council Conference Call	February 15, 2013	N/A
Workshop	Electronic Structure Calculation Methods on Accelerators	February 6–8, 2012	64
Conference	Industry–National Laboratory Workshop on Modeling and Simulation*	March 7, 2012	167
Workshop	Performance Analysis Tools Workshop	March 20–21, 2012	30
Workshop	INCITE Webinar	March 26, 2012	26
Conference	Accelerating Computational Science Symposium 2012 (ACSS 2012)	March 28–30, 2012	103
Educational Outreach	OLCF New Users Training and Titan Development and OLCF Users Meeting	April 16–20, 2012	73
Workshop	INCITE Webinar	April 24, 2012	60
Educational Outreach	Crash Course in Supercomputing	June 12–13, 2012	85
Workshop	A Preview of MPI 3.0: The Shape of Things to Come	June 24–28, 2012	60
Workshop	NCCS HPC Fundamentals Series	June 27, 2012	24
Workshop	R Workshop	June 28–29, 2012	78
Workshop	NCCS HPC Fundamentals Series	July 2, 2012	18
Workshop	NCCS HPC Fundamentals Series	July 9, 2012	23
Educational Outreach	ARC Students	July 9–20, 2012	10
Workshop	NCCS HPC Fundamentals Series	July 16, 2012	26
User Conference Call	User Council Conference Call	July 18, 2012	N/A
Workshop	NCCS HPC Fundamentals Series	July 23, 2012	18
Workshop	NCCS HPC Fundamentals Series	July 30, 2012	10
Workshop	NCCS HPC Fundamentals Series	August 10, 2012	10
User Conference Call	User Council Conference Call	September 13, 2012	N/A
Workshop	OLCF New User Training	October 8, 2012	44
Workshop	Cray Technical Workshop on XK6 Programming*	October 9–10, 2012	104

* The OLCF co-hosted this event.

ARC = Appalachian Regional Commission.

N/A = Not applicable.

Training the Next Generation

The OLCF maintains a broad program of collaborations, internships, and fellowships for young researchers. Twenty-nine faculty, student interns, and postdoctoral researchers were supported from January 1, 2012, through December 31, 2012. Examples of user engagement and outreach include the following:

- Ten students and teachers from around Appalachia gathered at ORNL this past summer for interactive training with volunteers from the OLCF. In its fourth year, the program is a partnership between ORNL and the Appalachian Regional Commission Institute for Science and Mathematics.

Students built a Beowulf cluster using Mac minis and were then challenged to write a parallel program and compile and execute the program on the cluster.

- During SC12 OLCF nuclear physicist Hai Ah Nam and HPC user support specialist Fernanda Foertter participated in a “birds-of-a-feather” session speaking on women in computing and stressing the need for better recruiting and retention policies for women working in science, technology, engineering, and mathematics.
- For the second year in a row, the OLCF offered an 8-week introductory course in supercomputing for interested interns and employees at ORNL. The course is designed to give more insight into HPC, provide a glimpse into what the OLCF’s mission is, and introduce participants to the basic concepts of HPC.
- OLCF Director of Science, Jack Wells, and Bronson Messer spoke to computational science graduate students at the DOE Computational Science Graduate Fellowship’s annual conference about HPC.
- For the past 3 years Dustin Leverman has participated in the Student Cluster Competition as a Supercomputing (SC) committee member, the annual supercomputing conference. The goal of the competition is to expose students to the field of HPC. By participating in this event, the OLCF has the opportunity to engage with students interested in HPC and as a result has hired two former participants in the competition into the facility, including Dustin Leverman himself.
- Over the summer Science Undergraduates Laboratory Internships interns Ryan Laurenza of the University of North Carolina–Chapel Hill and Mark Keele of Middle Tennessee State University worked under the mentorship of OLCF computer scientists Mitchell Griffith and Adam Simpson to write code for the Resource Allocation and Tracking System, which handles all the project and user-allocation information for the OLCF. At the conclusion of their internships, Laurenza and Keele presented their work at the ORNL Summer 2012 Poster Session and Graduate Student Recruitment Fair. The two made valuable contributions that will aid OLCF users while gaining experience that will help them both in their future careers.
- PhD students Thomas Papatheodore and Austin Harris (both University of Tennessee physics) and Masters student Rohan Garg (Northeastern University computer science) participated in the Higher Education Research Experiences program this summer, mentored by Bronson Messer (Papatheodore and Harris) and Judy Hill (Garg). Harris refactored portions of the nuclear kinetics code XNET for acceleration on GPUs and tested several different programming models and libraries for GPU-accelerated linear algebra. Papatheodore worked on extending the capabilities of the FLASH code for astrophysical reactive flow simulations. These extensions included the incorporation of Harris’ work in XNET into the larger FLASH codebase. Garg investigated the use of HMPP compiler directives for accelerators in the software framework, MADNESS, and compared the performance of the directives-based code to that of hand-tuned CUDA kernels. The students expressed a strong interest in the type of work undertaken by SciComp at the OLCF and will likely return for the summer of 2013.

1.5.5 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, and center publications. The OLCF provides tours to groups throughout the year for visitors that range from middle-school students through senior-level government officials. The center gave tours for 533 distinct groups in CY2012 and highlighted the research from many different projects. These highlights can be found on the OLCF website. In 2012 the OLCF produced more than 55 new highlights and the 2011/2012 annual report.

The center's outreach team also helped users and nonusers alike prepare for the unveiling of the world's most powerful computer. To make the general public aware of the achievement and the science possible on Titan, the OLCF began a media launch well in advance of Titan's unveiling. This effort was undertaken in close conjunction with the center's industry partners, NVIDIA and Cray. The media launch was a combination of targeted press releases, scheduled on-site visits by journalists, and videos and interviews with researchers and computing experts. Besides HPC standards such as *HPCWire* and *InsideHPC*, Titan was featured in a variety of media outlets including Yahoo, *The Washington Post*, *USA Today*, *Forbes*, NPR, *Popular Science*, *Wired*, Fox News, the BBC, *Time*, and *National Geographic*. The media launch is profiled on the Titan media webpage,^{††} which lists many of the articles, videos, and images associated with the OLCF's Titan media effort. The OLCF also produced a poster for the National User Facility Organization User Science Exhibition on Capitol Hill. The event was organized to highlight the significant and important role that scientific user facilities play in science education, economic competitiveness, fundamental knowledge, and scientific achievements. The center contributed a poster that highlighted both the science undertaken at the center and its resources and provided video images of the OLCF.

^{††} <http://www.olcf.ornl.gov/titan/>

Business Results

HIGH PERFORMANCE COMPUTING FACILITY
2012 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2013

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 highly capable and reliable systems to the user community. The effective delivery of these resources is further demonstrated by the business result metrics, which were met or exceeded. These leadership-class computational resources support scientific research through production simulation across many scientific domains, providing the key computing resources that are critical to their success.

2.1 Business Results Summary

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

2.2 Cray XT Compute Partition Summary

In accordance with the Critical Decision (CD)-3 signed by the Acquisition Executive, the OLCF upgraded the Cray compute system from a model XT5 to a model XK in the fourth quarter of 2011. The initial portion of this upgrade, *Phase 1*, included the installation of 4,672 new XK compute blades, each configured with four AMD Opteron™ 6274 processors and the upgrade of the system interconnect fabric from SeaStar to Gemini. The Phase 1 upgrade provided a new baseline of 18,688 compute nodes for 2012. In February 2012, near the conclusion of the Phase 1 upgrade, 960 of the XK nodes were also upgraded with the installation of an NVIDIA X2090 (Fermi) accelerator. The NVIDIA accelerator is a hosted accelerator, i.e. the Cray XK compute node can be configured as either a standalone CPU or as a CPU-GPU pair. These 960 Fermi-equipped nodes were managed separately, as a logical partition, for the entire term of their installation with specific access and use requirements for a restricted set of users. Utilization information for this logical partition is provided.

Phase 2 of the upgrade, conducted during the fourth quarter of 2012, included the removal of the 960 NVIDIA Fermi accelerators and the installation of 18,688 NVIDIA K20X (Kepler) accelerators. The revised cabinet configuration increased the peak electrical demand load for each cabinet from approximately 36kW/cabinet to about 45kW/cabinet. Modifications were made to the electrical power distribution and mechanical (chilled water) distribution systems to accommodate this load. A new 3.0MW transformer and switchboard were installed, and 53 cabinets were relocated

from existing switchboards to the new switchboard. In addition, the chilled water flow control valves for the 48 Liebert XDPs were upgraded to allow up to 120 gallons per minute, commensurate with the additional anticipated load being shed to the chilled water system. The facility modifications were accomplished as a rolling upgrade in summer 2012.

The Phase 1 and Phase 2 activities impacted full system availability in 2012. The timeline that reflects these impacts to system availability in 2012 is shown in Figure 2.1. The area in green indicates system availability; red areas reflect partial, short term, or full system interruptions. Despite these significant activities, the OLCF exceeded its commitments to the INCITE, ALCC, and DD programs for AMD Opteron-based compute hours, delivering more than 1.45 billion compute hours in 2012.




Figure 2.1. Commissioning Activities Affecting System Availability in 2012.

The OLCF provided the following computational resources in 2012 for scientific research (reference Table 2.1). The Cray XK with the 960 Fermi accelerators was managed as a separate logical partition, and is described separately. That particular configuration of Gemini+Fermi is a Cray XK6 and was managed as *TitanDev*. The large partition of CPU-only XK blades was managed as *JaguarPF*. The particular configuration of Gemini+Kepler is a Cray XK7. That system, to be accepted in 2013 will be managed as *Titan*. All are members of the XK series. The Cray XK6 system

described here is the smaller partition made available in December 2011 and decommissioned in October 2012.

Table 2.1. OLCF Production Computer Systems, 2012

System	Access	Type	CPU	GPU	Computational Description			Interconnect
					Nodes	Node Configuration	Memory Configuration	
Titan	Not generally available. Staff, Vendor, and CAAR teams.	Cray XK7	2.2 GHz AMD Opteron 6274 (16-core)	732 MHz NVIDIA K20X (Kepler)	18,688	16-core SMP + 14 streaming multiprocessor (SM) GPU (hosted)	32 GB DDR3-1600 and 6 GB GDDR5 per node; 598,016 GB DDR3 and 112,128 GB GDDR5 aggregate	Gemini (160 GB/sec)
TitanDev	Selected and Qualified Users	Cray XK6	2.2 GHz AMD Opteron 6274 (16-core)	1000 MHz NVIDIA X2090 (Fermi)	960	16-core SMP + 14 SM GPU (hosted)	32 GB DDR3-1600 and 6 GB GDDR5 per node; 30,720 GB DDR3 and 5,760 GB GDDR5 aggregate	Gemini (160 GB/sec)
JaguarPF	Full Production	Cray XK6	2.2 GHz AMD Opteron 6274 (16-core)	None (empty SXM slot)	17,728	16-core SMP	32 GB DDR3-1600 per node; 567,296 GB aggregate	Gemini (160 GB/sec)

The OA Business Results reported here are based on the length of time the computational resource has been in production. The OLCF production computational systems entered into production according to the following schedule (reference Table 2.2). This includes historical and forward-looking 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. Entries with the lighter shaded font were not generally available during this OA period.

Table 2.2. OLCF HPC System Production Dates, 2008–Present

System	Type	Production Date	Performance End Date	Notes
Titan	Cray XK7	NULL	NULL	Release to production at conclusion of Acceptance Test in 2013.
JaguarPF	Cray XK6	September 18, 2012	October 7, 2012	Production at 240,000 cores until September 18, when partition size was reduced to 120,000 AMD Opteron cores. Additional Kepler installation. TitanDev access terminated.
JaguarPF	Cray XK6	February 13, 2012	September 12, 2012	Full production until September 12, when partition size was reduced to 240,000 AMD Opteron cores. Beginning of Kepler installation
JaguarPF	Cray XK6	February 2, 2012	February 13, 2012	Stability Test. Restricted user access. 299,008 AMD Opteron 6274 cores. Includes 960-node Fermi-equipped partition.
JaguarPF	Cray XK6	January 5, 2012	February 1, 2012	Acceptance. No general access. 299,008 AMD Opteron cores

Table 2.2. OLCF HPC System Production Dates, 2008–Present (continued)

System	Type	Production Date	Performance End Date	Notes
JaguarPF	Cray XK6	December 12, 2011	January 4, 2012	142,848 AMD Opteron cores
JaguarPF	Cray XT5	October 17, 2011	December 11, 2011	117,120 AMD Opteron cores
JaguarPF	Cray XT5	October 10, 2011	October 16, 2011	162,240 AMD Opteron cores
JaguarPF	Cray XT5	September 25, 2009	October 9, 2011	224,256 AMD Opteron cores
JaguarPF	Cray XT5	August 19, 2008	July 28, 2009	151,000 AMD Opteron cores

The production date used for computing statistics is either the initial production date or the production date of the last upgrade to the computational resource. The performance end date is the last calendar day that user jobs were allowed to execute on that partition. For a period of one year following either system acceptance or a major system upgrade, the targeted scheduled availability for that HPC computational or storage system is at least 85% and the targeted overall availability is at least 80%.

Business Results are provided for the OLCF computational resources, the High-Performance Storage System (HPSS) Archive System, and the external Lustre File Systems (reference Tables 2.3 to 2.5).

Table 2.3. OLCF Business Results Summary for HPC Systems

	Measurement	2011 Target	2011 Actual	2012 Target	2012 Actual
Cray XE6/XK6 (JaguarPF)	Scheduled Availability		NIP	85.0%	98.11%
	Overall Availability		NIP	80.0%	91.45%
	MTTI (hours)		NIP	NAM	132.89
	MTTF (hours)		NIP	NAM	225.59
	Total Usage		NIP	NAM	84.39%
	Core Hours Used		NIP	NAM	1,452,936,146
	Core Hours Available		NIP	NAM	1,721,620,377
	Capability Usage				
	INCITE Projects		NIP	NAM	48.36%
	All Projects		NIP	30.0%	50.67%
Cray XT5 (JaguarPF)	Scheduled Availability	95.0%	96.37%	NIP	NIP
	Overall Availability	90.0%	92.88%	NIP	NIP
	MTTI (hours)	NAM	60.38	NIP	NIP
	MTTF (hours)	NAM	79.66	NIP	NIP
	Total Usage	NAM	87.11%	NIP	NIP
	Core Hours Used	NAM	1,428,874,052	NIP	NIP
	Core Hours Available	NAM	1,640,290,505	NIP	NIP
	Capability Usage				
	INCITE Projects	NAM	47.8%	NIP	NIP
	All Projects	40.0%	54.0%	NIP	NIP

Table 2.3. OLCF Business Results Summary for HPC Systems (continued)

	Measurement	2011 Target	2011 Actual	2012 Target	2012 Actual
Cray XT4 (Jaguar)	Scheduled Availability	95.0%	97.58%	NIP	NIP
	Overall Availability	90.0%	97.09%	NIP	NIP
	MTTI (hours)	NAM	78.67 hours	NIP	NIP
	MTTF (hours)	NAM	87.80 hours	NIP	NIP
	Total Usage	NAM	90.73%	NIP	NIP
	Core Hours Used	NAM	39,079,672	NIP	NIP
	Core Hours Available	NAM	43,070,274	NIP	NIP
	Capability Usage				
	INCITE Projects	NAM	39.1%	NIP	NIP
	All Projects	NAM	57.1%	NIP	NIP

MTTF = Mean time to failure.

MTTI = Mean time to interrupt.

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

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

Table 2.4. OLCF Business Results Summary for HPSS

	Measurement	2011 Target	2011 Actual	2012 Target	2012 Actual
HPSS	Scheduled Availability	95.0%	99.81%	95.0%	99.57%
	Overall Availability	90.0%	98.65%	90.0%	98.46%
	MTTI (hours)	NAM	224.73	NAM	228.6
	MTTF (hours)	NAM	628.03	NAM	588.85

MTTF = Mean time to failure.

MTTI = Mean time to interrupt.

Table 2.5. OLCF Business Results Summary for the External Lustre File Systems

	Measurement	2011 Target	2011 Actual	2012 Target	2012 Actual
Widow 1	Scheduled Availability	95.0%	99.26%	95.0%	99.88%
	Overall Availability	90.0%	97.95%	90.0%	98.25%
	MTTI (hours)	NAM	536.27	NAM	719.15
	MTTF (hours)	NAM	785.84	NAM	2,924.48
Widow 2	Scheduled Availability	95.0%	99.93%	95.0%	99.81%
	Overall Availability	90.0%	99.34%	90.0%	98.69%
	MTTI (hours)	NAM	966.92	NAM	722.42
	MTTF (hours)	NAM	1750.78	NAM	2,191.89
Widow 3	Scheduled Availability	95.0%	99.95%	95.0%	99.89%
	Overall Availability	90.0%	99.36%	90.0%	98.95%
	MTTI (hours)	NAM	967.10	NAM	869.14
	MTTF (hours)	NAM	1751.09	NAM	1,754.82

MTTF = Mean time to failure.

MTTI = Mean time to interrupt.

2.3 Resource Availability

2.3.1 Scheduled Availability

2012 Operational Assessment Guidance

For HPC Facilities, scheduled availability (reference formula #1) is the percentage of time a designated level of resource is available to users, excluding scheduled downtime for maintenance and upgrades. To be considered a scheduled outage, the

user community must be notified of the need for a maintenance event window no less than 24 hours in advance of the outage (emergency fixes). Users will be notified of regularly scheduled maintenance in advance, on a schedule that provides sufficient notification, and no less than 72 hours prior to the event, and preferably as much as seven calendar days prior. If that regularly scheduled maintenance is not needed, users will be informed of the cancellation of that maintenance event in a timely manner. Any interruption of service that does not meet the minimum notification window is categorized as an unscheduled outage.

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

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

As shown in Table 2.6, the OLCF has exceeded the scheduled availability targets for the facility's computational resources for 2011 and 2012.

Table 2.6. OLCF Business Results Summary: Scheduled Availability

	System	2011 Target	2011 Actual	2012 Target	2012 Actual
Scheduled Availability	Cray XK6	NIP	NIP	85.0%	98.11%
	Cray XT5	95.0%	96.37%	NIP	NIP
	Cray XT4	95.0%	97.58%	NIP	NIP
	HPSS	95.0%	99.81%	95.0%	99.57%
	Widow 1	95.0%	99.26%	95.0%	99.88%
	Widow 2	95.0%	99.93%	95.0%	99.81%
	Widow 3	95.0%	99.95%	95.0%	99.89%

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

Assessing Impacts to Scheduled Availability

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

2.3.2 Overall Availability

2012 Operational Assessment Guidance

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

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

As shown in Table 2.7, the OLCF has exceeded the overall availability targets for the facility's computational resources for 2011 and 2012.

Table 2.7. OLCF Business Results Summary: Overall Availability

	System	2011 Target	2011 Actual	2012 Target	2012 Actual
Overall Availability	Cray XK6	NIP	NIP	80.0%	91.45%
	Cray XT5	90.0%	92.88%	NIP	NIP
	Cray XT4	90.0%	97.09%	NIP	NIP
	HPSS	90.0%	98.65%	90.0%	98.46%
	Widow 1	90.0%	97.95%	90.0%	98.25%
	Widow 2	90.0%	99.34%	90.0%	98.69%
	Widow 3	90.0%	99.36%	90.0%	98.95%

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

2.3.3 Mean Time to Interrupt

2012 Operational Assessment Guidance

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

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

where time in period is start time – end time

start time = end of last outage prior to reporting period

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

The Mean Time to Interrupt summary is shown in Table 2.8.

Table 2.8. OLCF Business Results Summary: Mean Time to Interrupt

	System	2011 Target	2011 Actual	2012 Target	2012 Actual
MTTI (hours)	Cray XK6	NIP	NIP	NAM	132.89
	Cray XT5	NAM	60.38	NIP	NIP
	Cray XT4	NAM	78.67	NIP	NIP
	HPSS	NAM	224.73	NAM	228.6
	Widow 1*	NAM	536.27	NAM	719.15
	Widow 2	NAM	966.92	NAM	722.42
	Widow 3	NAM	967.10	NAM	869.14

*Due to the extremely long uptime of the Widow files systems, the formula for MTTI can produce artificially skewed results using the calendar year period defined in the formula. Values presented here as “Actual” for Widow 1, Widow 2, and Widow 3 were calculated based on a calendar year period without regard for potential skew.

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

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

2.3.4 Mean Time to Failure

2012 Operational Assessment Guidance

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

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

where time in period is start time – end time

start time = end of last outage prior to reporting period

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

The Mean Time to Failure summary is shown in Table 2.9.

Table 2.9. OLCF Business Results Summary: Mean Time to Failure

	System	2011 Target	2011 Actual	2012 Target	2012 Actual
MTTF (hours)	Cray XK6	NIP	NIP	NAM	225.59
	Cray XT5	NAM	79.66	NIP	NIP
	Cray XT4	NAM	87.80	NIP	NIP
	HPSS	NAM	628.03	NAM	588.85
	Widow 1	NAM	785.84	NAM	2,924.48
	Widow 2	NAM	1750.78	NAM	2,191.89
	Widow 3	NAM	1751.09	NAM	1,754.82

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

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

2.4 Resource Utilization

2012 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 concept of core hours is applicable to current sites. Subsequent versions of this calculation may need to be revised to better reflect the specific systems at a particular Facility.

For the Cray XK for the OA period January 1–December 31, 2012, 1,452,936,146 traditional CPU core hours were delivered from a scheduled maximum of 1,721,620,377 core hours. This resulted in total system utilization for the Cray XK of 84.39%. Note: due to the service reductions associated with the Phase 1 and Phase 2 upgrades this number reflects a production start date of February 2, 2012, and a production end date of October 7, 2012.

Understanding Resource Utilization Measurements

In 2012, the Cray XK systems, delivering resources as both the 17,728-node CPU-only partition JaguarPF and as the smaller 960-node partition TitanDev, exacerbated an existing difference in the manner in which the computer resources are managed, and how that resource utilization is reported. The current job scheduler for the OLCF compute resources is Adaptive Computing's Moab, coupled to the Cray resource manager, Torque. Moab/Torque allocates resources at the granularity of a single node, regardless of the composition of that node.

For CPU-only partitions, the calculation of resource utilization for 2012 is quite straightforward. For each node-hour consumed, there is a direct translation (x16) to the number of CPU core hours consumed. An application that consumes one node-hour on JaguarPF is charged at an effective rate of 16 core hours for that use. This allows programs such as INCITE, which allocated cycles for CY2012 to their users in core hours to have an easy translation to allocated node-hours for system administration and management purposes. The utilization information that is provided to the users retains the distinction of core hours in 2012, with the translation from node-hours to core hours managed by the OLCF.

For CPU+GPU partitions such as TitanDev, and with Titan in 2013, there is now the concept of the hybrid node that combines 16 traditional CPU cores and 14 streaming multiprocessors (SMs). From a system accounting perspective, a job that runs on a hybrid node accumulates 16 traditional core hours and 14 SM-hours. For 2013, INCITE program allocations on Titan were made in core hours with the explicit

stipulation that a node-hour was the equivalent of 30 Titan core hours. It is important that we do not equate a traditional CPU core hour with an SM-hour, as the architectural concepts are considerably different.

Hours accumulated on the TitanDev partition during its availability in 2012 were managed from a system accounting perspective in terms of node-hours. These hours were reported to the users as x16 equivalents, not x30. This was intended to preclude any perception that using the hybrid nodes was punitive. In 2013, with a uniformly configured system, hours will be reported to users in terms of both node-hours consumed, and then with the equivalent x30 multiplier.

It is the opinion of the OLCF that there should be a structured transition to the allocation programs (INCITE, ALCC, DD, and others), where the principle allocable unit is a node-hour. For consistency of reporting across multiple years, the concept of core hours can be retained. The OLCF will work with the program manager to consider revision of the definition for resource utilization measurement units beginning in 2013.

2.4.1 Total System Utilization

2012 Operational Assessment Guidance

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

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

The system utilization, by program, and by system are shown in Table 2.10. This table reflects combined system utilization for the XT4 and XT5 across programs, since there is no separate allocation by both program and system, and the assessment of those total hours by system.

Table 2.10. 2012 OLCF System Utilization

Program	Measurement Period	CPU Hours Allocated	CPU Hours Consumed	CPU Hours Available	Percent of Allocation Consumed	Program Consumption
INCITE	CY2012	940,000,000	1,012,273,087*		107.69%	79.04%
ALCC	CY2012	N/A. ALCC allocates on split-year period	201,778,586			15.76%
DD	CY2012	N/A. DD allocates on a continuous basis	92,264,636			7.20%
Subtotal			1,280,643,283			100.00%
Staff and Vendor	CY2012	0.5% of all hours	172,292,863			
Total			1,452,936,146	1,721,620,377	84.39%	100%

N/A = Not applicable.

* CPU Hours consumed by the INCITE Program includes 31,615,890 core hours that were consumed by non-renewing 2012 INCITE projects in the period ending 02/05/2013.

The OLCF tracks the consumption of core hours by job. This can then be extended to track the consumption of core hours by program, project, user, and system with high fidelity. Figure 2.2 describes the utilization by week and by Program for all of 2012. No adjustment is made to exclude any user group, including staff and vendors.




Figure 2.2. 2012 XT5 Resource Utilization – Opteron Core Hours by Program.

Assessing Total System Utilization

Allocation programs INCITE, ALCC, and DD are aggressively monitored to ensure that projects within these allocation groups maintain adequate consumption rates. This is reflected in both the successful delivery of more than 100% of the INCITE allocations, and by the steady consumption by these programs week to week.

Note that non-renewed INCITE projects from 2011 were allowed by OLCF policy to continue running at low priority early in the first quarter of 2012 so that those projects could complete while new 2012 INCITE projects ramped up. Not only is this a user-friendly policy for non-renewed projects that have not quite exhausted their allocation, it serves to increase utilization while new projects establish a more predictable consumption routine.

2.5 Capability Utilization

2012 Operational Assessment Guidance – Capability Utilization

The Facility shall describe the agreed definition of capability, the agreed metric, and the operational measures that are taken to support the metric.

Leadership Class (capability) is defined by the minimum number of nodes allocated to a particular job on the OLCF computing resources. Leadership-class jobs must use at least 20% of the available nodes of the largest system to qualify.

The capability metric is defined by the number of node-hours that are delivered by leadership-class jobs. In 2011, the metric stipulated that no less than 35% of the delivered node hours on the Cray XT5 would reflect leadership-class jobs. For the first year of Cray XK production (2012), the metric stipulates that no less than 30% of the delivered node hours reflect leadership-class jobs. This is proposed to increase to 35% in subsequent production years.

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

The OLCF Capability Utilization Definition is summarized in Table 2.11.

Table 2.11. OLCF Capability Utilization Definition

System	Year 1		Subsequent Years	
	Definition for Leadership Class (Capability)	Capability Metric	Definition for Leadership Class (Capability)	Capability Metric
Cray XK6 (JaguarPF and TitanDev)	20%	30% of delivered hours	20%	35% of delivered hours

The OLCF continues to exceed expectations for capability usage of its HPC resources (Table 2.12). Keys to the growth of leadership usage include the liaison role provided by the SciComp Group members, who work hand-in-hand with users to port, tune, and scale code, and ORNL 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.

Table 2.12. OLCF Leadership Usage on the Cray XT and XK Systems

	Leadership Usage	CY2011 Target	CY2011 Actual	CY2012 Target	CY2012 Actual
Cray XK6	INCITE	NIP	NIP	NAM	48.36%
	Total	NIP	NIP	30%	50.67%
Cray XT5	INCITE	NAM	47.8%	NIP	NIP
	Total	40.0%	54.0%	NIP	NIP
Cray XT4	INCITE	NAM	39.1%	NIP	NIP
	Total	40.0%	57.1%	NIP	NIP

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

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

The average consumption of hours by leadership-class jobs was well above the CY2012 target of 30% at 50.67%. This consumption varies during the year, affected by factors including system availability and the progress by the various projects within their research. The distribution of the consumption of hours by month is shown in Figure 2.3.




Figure 2.3. Effective Scheduling Policy Enables Leadership-Class Usage.

Leadership-class jobs are not restricted to the INCITE program. There are leadership-class jobs across the ALCC and DD programs as well. The contribution to capability utilization by program is shown in Figure 2.4.




Figure 2.4. Capability Usage by Project Type.

2.6 Management of INCITE Selection Process (LCFs only)

2.6.1 Process

The Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program promotes transformational advances in science and technology through large allocations of computer time, supporting resources, and data storage at the Argonne and Oak Ridge Leadership Computing Facilities (LCFs) for computationally intensive, large-scale research projects. A detailed description of the selection process is available online.[‡]

The INCITE awards committee is comprised of the LCF directors, INCITE manager, LCF directors of science, and senior management. The committee identifies the top-ranked proposals by (a) peer-review panel ratings, rankings, and reports and (b) additional considerations, such as the desire to promote use of HPC resources by underrepresented communities.

Figure 2.5 schematically outlines the decision-making process of the INCITE awards committee. Input from the peer-review panels and computational-readiness reviewers is combined to yield an initial list of projects sorted by panel-ranked order. The INCITE awards committee identifies the top-ranked proposals by (a) peer-review panel rating and reports and (b) additional considerations, such as the desire to promote use of HPC resources by underrepresented communities. A balance is struck to ensure that each awarded project has sufficient allocation to enable all or part of the proposed scientific or technical achievements and to maximize the scientific support provided to each INCITE project.

[‡] www.doeleadershipcomputing.org/policies/incite_overview_policies.pdf

Decision-making process



Figure 2.5. INCITE Award Decisions Workflow.

The INCITE award period is January through December of the calendar year following the date of the call for proposals. (For example, “2011 INCITE” refers to the call for proposals, reviews, etc., that took place in calendar year 2010 in support of awards for the January through December 2011 time frame). The timeline of activities for the reporting period of the LCF operational assessment are summarized in Table 2.13.

Table 2.13. 2013 Call for Proposals and Awards Activities

April 11, 2012	2013 INCITE Call for Proposals Opened
June 27, 2012	2013 INCITE Call for Proposals Closed
September 25, 26, 27, 2012	2013 INCITE Peer-Review Panel (Rockville, MD)
October 22, 2012	Announcement of Awards to Principal Investigators for CY2013

2.6.2 Peer Reviewers

The INCITE manager will convene independent peer-review panels to evaluate each proposal’s potential for impact. Proposals will be evaluated on scientific quality, proposed impact, appropriateness of the proposed method or approach, competence of the principal investigator (PI) and proposed research team, computational plan, and reasonableness and appropriateness of the proposed request for computational resources. Scientific review panels are composed of application domain experts from national laboratories, universities, and industry who have a working knowledge of the current computational challenges and opportunities in their fields.

83 science experts participated in the 2013 INCITE Peer-Review Panel. More than half of the reviewers are Society fellows (AAAS, APS, SIAM, IEEE, etc.), agency

awardees (e.g., NSF Early Career), Laboratory fellows, National Academy members, and National Society presidents. Forty-one percent of these reviewers also participated in the 2012 INCITE review. Figure 2.6 illustrates the organizational affiliation of the reviewers.




Figure 2.6. 2013 INCITE Reviewer Affiliation.

The INCITE manager conducts a survey of the reviewer participants at the INCITE Peer-Review Panel. The questions and responses are summarized in Table 2.14. The scores indicate that, in the opinion of the diverse set of science, engineering, and computer science experts who participated in the one-day review of proposals, the INCITE program represents cutting-edge computational work, the proposals are of high quality, the panel itself was sufficiently diverse, and the reviewers are satisfied with the panel review process.

Table 2.14. INCITE Reviewer Survey Results (>80% Response Rate)

Survey Question	2012 INCITE	2013 INCITE
The INCITE proposals discussed in the panel represent some of the most cutting-edge computational work in the field.	4.5	4.5
The proposals were comprehensive and of appropriate length given the award amount requested.	4.2	4.1
The science panel was sufficiently diverse to assess the range of research topics being considered.	4.4	4.3
Please rate your overall satisfaction with the [2012, 2013] INCITE Science Panel review process.	4.8	4.6

Scores of “1” indicate “strongly disagree” and “5” indicate “strongly agree.” (Or, for the final question, “1” indicates “very dissatisfied” and “5” means “very satisfied.”)

2.6.3 Proposal Allocation

Statistics on the applications submitted and awarded are reported in Table 2.15. INCITE grants one-, two-, and three-year awards: Multiyear projects are eligible to submit a renewal application. Projects that have completed the term of their award may submit a new proposal for the next allocation period.

Table 2.15. INCITE Submittal and Award Statistics

Data Description	2012 INCITE	2013 INCITE
Total number of proposals submitted (percentage accepted)	84 (33%)	123 (33%)
Total number of renewals submitted (percentage accepted)	35 (91%)	20 (100%)
Total number of awards	60	61
Total number of hours requested	5,163M	14,679M
Total number of hours awarded	1,672M	4,678M
Oversubscription	309%	314%

2.7 Assessing GPU Utilization

The upgrade of the existing Cray XT5 to the XK series system provided a key new capability to users, allowing them to exploit a new hybrid compute node that contains both a CPU and a NVIDIA accelerator. Beginning with the introduction of the 960 Fermi-equipped nodes, and the associated updates to the software stack, users were provided the ability to offload specific portions of their application to the accelerator. On any hybrid node, the GPU is an option for the user. There is no explicit requirement to use it, and a hybrid node can be used in the exact same manner as an Opteron-only node.

As identified in the 2011 OA Recommendations, the “OLCF should quickly develop one or more methods to measure GPU utilization, deploy on a trial basis to evaluate the efficacy of the measurement approaches, and implement the most useful measures for reporting in the future...” From a practical standpoint, the OLCF has approached this as the need to track, on a per-job basis, whether (a) applications are actively exploiting the GPU; and then (b) to what degree those applications are actively exploiting the GPU. At the time of the Fermi installation, neither capability existed. However, during the first half of 2012, the OLCF defined a path-forward that addresses both needs.

2.7.1 Identifying the Use of GPU-Specific Libraries with ALTD

In 2009, ORNL developed and deployed the Automatic Library Tracking Database (ALTD). ALTD actively monitors the compilation phase of individual applications, and at link time, creates a unique record for that application that contains a list of each of the libraries that were linked against that particular binary. When this application is executed via aprun, a new ALTD record is written to the database that contains the name of the executable, the batch job id, and other supporting information. To determine whether a specific executable takes advantage of the GPU, we examine whether an executed job, for which we have all of the per-job scheduling information, was linked against an accelerator-specific library. For 2012, this includes any library that matches the following identifiers:

libacc*, libOpenCL*, libmagma*, libhmp*, libcuda*, libcupti*, libcula*,
libcublas*

Jobs whose executables are linked against one of the above are deemed to have used the accelerator. From this information, per-job utilization can be derived and aggregated in to reports that describe system utilization across GPU-enabled and CPU-only qualifiers.

While this method does provide a mechanism for examining CPU-only and GPU-accelerated contributions to System Utilization, there are some limitations.

- While ALTD is enabled by default, it must be disabled in certain instances. Then, job compilations executed without inclusion in ALTD will not produce corresponding records, and subsequent execution of that binary will return a NULL lookup result.
- Debugging sessions cannot generally tolerate the wrapped aprun, so will contribute to an unknown result.
- Jobs that are executed outside of the job scheduler, such as in a dedicated mode, will not generate job records and cannot be correlated. These will contribute to an unknown result.

The results of the implementation of this method are shown in Figure 2.7, where the information is aggregated on a weekly schedule, and the results are expressed in Titan core hours. In the Figure, node hours that were accumulated on compute nodes without using the GPU are aggregated as “No.” Node hours that were consumed by applications that met the known triggers for needing and using the GPU are aggregated as “Yes.” Because much of the measurement period coincided with Acceptance Test preparation, there are large periods of time where job execution was managed outside of Moab, contributing to the significant number of jobs that are identified as “Unknown.” In addition, the terms of the Acceptance Test include significant numbers of CPU-only applications that were being prepared for Acceptance. These impact the jobs in the “No” category. The actual data in the Figure is not tied to any target, metric, or goal, but does provide a representative description of the capability to measure and identify the job distribution. The period in October 2012 with no production is the Cray upgrade/diagnostics period and significant dedicated testing. Since December 12, 2012, the OLCF has been tracking system utilization with additional detail that tracks the contribution of both accelerated and non-accelerated leadership-class jobs to the total system utilization.




Figure 2.7. Tracking GPU Usage on Titan.

2.7.2 Accessing Limited NVML Data through Cray’s Resource Utilization Reporting

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

To meet this requirement, the OLCF actively engaged with both NVIDIA and Cray in 2012 to define what information was needed from a system accounting

perspective. The results of those discussions are driving revisions to the NVIDIA device driver, the development of an accompanying API and library, and changes to the Cray Resource Utilization software so that additional information about the GPU utilization, on a per-job basis, will be available at the conclusion of each job. Anticipated information available includes the number of GPU contexts opened; the number of GPU kernels run; the amount of time (GPU-seconds) accrued while running the kernels; and the high water mark of memory used on the GPU. This development effort is progressing, with changes to the NVIDIA driver on schedule for release in the second quarter of 2013, and changes to the Cray software stack on schedule for a beta release in summer 2013, and a production release in the third quarter of 2013.

2.8 Safety

The provisioning of a safe working environment and the demonstrated safety-conscious attitude of all subcontractors and employees remains an important consideration. In the face of the very high volume of work required by the Phase 1 and Phase 2 Cray XK upgrade, the ability to foster and promote a safe work environment remains paramount. The following quote, received from a Staff Member in the ORNL Safety Services Division described his perception of the success of the effort that we have expended to ensure a safe working environment.

“This year with all of the massive computer upgrades dealing with all of the subcontractors working in the facility, we had no safety problems, and more importantly, no injuries. This proves that the safety culture has arrived. The facility and lab space managers have done an excellent job communicating safety to all. The computer staff have meetings almost every week, and safety is of utmost concern to all. The cooperation from bottom up and top down has been a pleasure. As a safety professional, I have been allowed to be a part of the operations. Working together has grown the safety culture and that is innovative.” – David Edds, ORNL Safety Services.

2.9 Internal ORNL Operational Performance Review

In 2012, the NCCS, which manages the OLCF Project, participated in an internal OPR. The OPR is a collaborative, yet independent peer review of a division or directorate’s operational performance and risk management inclusive of contractor assurance, work control, and integrated performance management. This independent review evaluates the effectiveness and maturity of ORNL’s division or directorate’s implementation of Contractor Assurance and Integrated Safety Management requirements in terms of functionality, effectiveness and efficiency. The review is tailored to the performance and risks of each organization and includes assist opportunities on known problem areas.

The OPR team focused on the identified risks and operational challenges identified by NCCS including

- Managing risks in the User’s Program
- Facility infrastructure and maintenance/upgrades
- Responding to abnormal situations in the computer center
- Physical and cyber security, information protection

The OPR team identified no Findings, 17 Noteworthy Practices, and 29 Opportunities for Improvement (OFIs) during the assessment. The NCCS assessed the OFIs and, where relevant, updated or implemented changes to documentation and policies.

Strategic Results

HIGH PERFORMANCE COMPUTING FACILITY
2012 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2013

3. STRATEGIC RESULTS

CHARGE QUESTION 3: Is the facility enabling scientific achievements consistent with the Department of Energy 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.

To be sure, the projects and user programs operating within the OLCF advance DOE's mission to ensure America's security and prosperity by addressing its energy, environmental, and nuclear challenges through transformative science and technology solutions. In this section on strategic results, we describe and select a modest number of accomplishments that serve to communicate how OLCF is advancing two of DOE's four strategic goals, and associated targeted outcomes, of the DOE Strategic Plan:

- Goal 1: Catalyze the timely, material, and efficient transformation of the nation's energy system and secure US leadership in clean energy technologies.
- Goal 2: Maintain a vibrant US effort in science and engineering as a cornerstone of our economic prosperity with clear leadership in strategic areas.

3.1 Science Output

The facility tracks and reports the number of refereed publications written annually based on using (at least in part) the facility's resources. This number may include publications in press or accepted, but not submitted or in preparation. In 2012 the number of referred publications reportable within OAR guidance was 321. For comparison, in 2011 the number of reportable publications was 300.

The OLCF currently follows the recommendation in the 2007 report of the ASCR Advisory Committee Petascale Metrics Panel to report and track user products including, for example, publications, project milestones (requested quarterly; also examined in the INCITE renewal process), and code improvement. Publications are listed in Table 3.1. At the end of the year, a library search was carried out to identify additional publications based on work using OLCF resources. This library search was augmented by an automated search of the Web of Science database using user names and facility keywords as the basis for the search (see Section 4.2).

Table 3.1. OLCF Publications

	2011	2012
Number of refereed publications reportable within OAR guidance	300	321

3.2 Scientific Accomplishments

The OLCF advances DOE's science and engineering enterprise through robust partnerships with its users. The following sections provide brief summaries of and resources for obtaining more information on selected accomplishments that advance the state of the art in science and engineering research and development and are advancing DOE's science programs toward their targeted outcomes and mission goals. These selected highlights cannot capture the scope and scale of achievements enabled at the OLCF in 2012. As an additional indication of OLCF achievements, OLCF users published many breakthrough publications in high-impact journals in 2012, including three in *Nature*, one in *Nature Physics*, one in *Nature Climate Change*, and two in *Proceedings of the National Academy of Sciences*.

3.2.1 Understanding Solar Storm/Magnetosphere Interaction

William Daughton, Los Alamos National Laboratory, INCITE

Objective: Develop a better understanding of magnetic reconnection physics in the Earth's magnetosphere, where it plays a key role in space weather.

Impact: Space weather poses a continual threat to the Earth's technological systems and has already caused more than \$4 billion in satellite losses alone. We currently lack accurate forecasting capabilities, due in part to an inadequate understanding of the major driver of space weather, the so-called magnetic reconnection process. Magnetic reconnection enables explosive release of energy stored in magnetic fields as field lines break and reconfigure. It is the mechanism causing solar flares and is also operative in the Earth's magnetosphere. While considerable progress has been made within two-dimensional (2D) models of reconnection, very little is known regarding the influence of realistic three-dimensional (3D) dynamics. The advent of petascale computing has, for the first time, enabled large-scale 3D kinetic simulations, which treat the problem at the most basic level. These new predictions will be testable with high time resolution measurements from the National Aeronautics and Space Administration's (NASA's) upcoming Magnetospheric Multiscale Mission (MMS),^{§§} being developed for launch in 2014 to permit detailed observations of electron-scale physics embedded within the larger magnetohydrodynamic scales (Moore 2012). The simulation predictions are guiding the final planning of the MMS mission and contributing to the refinement of our 3D multiscale picture of reconnection, yielding improved understanding of the microscopic physics controlling the onset or quenching, variability, and mean rate of reconnection. This in turn will enable improved predictability of the structural features created by transient reconnection and their space weather consequences.

Accomplishments: New results from this project suggest that reconnection can spontaneously generate turbulence, which is dominated by coherent structures in the form of kinetic-scale current sheets. In 3D simulations, these sheets are unstable to the formation of magnetic flux ropes, which interact and generate new current sheets. The resulting turbulence enables a more efficient heating and transport of

^{§§} <http://mms.gsfc.nasa.gov/>

energetic particles into the magnetosphere, thus exaggerating the severity of space weather effects. Recent petascale particle-in-cell simulations of reconnection encompass three spatial dimensions while resolving electron kinetic scales. The simulations offer a number of striking predictions in which electron-scale physics is important, including the development of 3D flux ropes during component reconnection and the generation of volume-filling current layers in regions with strong flow shear (Roytershteyn 2012; Wan 2012).




Figure 3.1. Isosurface of Electron Vorticity Showing the Development of Vortex Tubes Along the Top (Magnetospheric Side of the Layer). (These vortices wrap up magnetic field lines, providing an alternative way of generating flux ropes.)

OLCF Contributions: On the upgraded Jaguar in 2012, Daughton et al. were able to use 264,000 cores to evolve 17 billion cells and 3.3 trillion particles. These leadership-class calculations are the largest VPIC simulations to date and required ~80 TB of disk space for this run alone. For the campaign, a minimum of 400 TB of scratch disk space was required. Visualization and analysis were performed on the Lens cluster using ParaView.

Total 2012 Usage: 75 million Jaguar core hours

Capability Utilization: <20% = 12%, 20% to 60% = 87%, 60% to 100% = 1%

Related Publications:

- V. Roytershteyn, W. Daughton, H. Karimabadi, F.S. Mozer, “Influence of the lower-hybrid drift instability on magnetic reconnection in asymmetric configurations,” *Phys. Rev. Lett.* **108**, 185001 (2012).
- M. Wan, W.H. Matthaeus, H. Karimabadi, V. Roytershteyn, M. Shay, P. Wu, W. Daughton, B. Loring, S.C. Chapman, “Intermittent dissipation at kinetic scales in collisionless plasma turbulence,” *Phys. Rev. Lett.* **109**, 195001 (2012).
- T.E. Moore, J.L. Burch, W.S. Daughton, S.A. Fuselier, H. Hasegawa, S.M. Petrinenec, Z. Pu, “Multiscale studies of the three-dimensional dayside X-line,” submitted to *J. Atmos. Sol. Terr. Phys.* (2012).
- H. Karimabadi, V. Roytershteyn, M. Wan, W.H. Matthaeus, W. Daughton, P. Wu, M. Shay, B. Loring, J. Borovsky, E. Leonardis, S. Chapman, T.K.M. Nakamura, “Coherent structures, intermittent turbulence and dissipation in high-temperature plasmas,” *Phys. Plasmas* **20**, 012303 (2013).

Online Story: <https://www.olcf.ornl.gov/2012/02/06/when-worlds-collide/>

3.2.2 How Many Isotopes Can There Be? Simulations Provide an Answer

James Vary, Iowa State University, INCITE

Objective: Using density functional theory and ORNL's Jaguar system, calculate the number of isotopes that are theoretically possible under the laws of physics.

Impact: The nuclear drip line identifies the limits of nuclear stability. For each number of protons in a nucleus, there is a limit to how many neutrons can bind to the nucleus, however fleetingly. Likewise, there is a limit to the number of protons that can be added to a nucleus with a given number of neutrons. With Jaguar, researchers were able to calculate these drip lines with confidence for heavy elements for which experimental data is not available. Not only does this work help us understand the nuclear landscape, but it also opens the way for designing “designer” nuclei with unique and useful properties. The nuclear configurations predicted near the neutron drip line will be studied for years at DOE's Facility for Radioactive Ion Beams currently under construction at Michigan State University.

Accomplishments: Answering one of the fundamental questions of nuclear structure physics, this project predicted the limits of nuclear stability by determining there are approximately 7,000 possible combinations of protons and neutrons allowed in bound nuclei with up to 120 protons (Erler 2012). Statistical and systematic uncertainties in the position of the drip lines were quantified. Extrapolations for drip-line positions and selected nuclear properties, including neutron separation energies relevant to astrophysical processes, were very consistent between the models employed. This project closely related to Scientific Discovery through Advanced Computing's (SciDAC's) Universal Nuclear Energy Density Functional project.




Figure 3.2. Map of Bound Even-Even Nuclei as a Function of Z and N. (Shown are stable nuclei [black], radioactive nuclei [green], mean drip lines with uncertainties [red], and two-neutron separation line [blue]. The inset shows the irregular behavior of the two-neutron drip line around Z=100.)

ORNL Contributions: These leadership-class simulations are some of the largest nuclear structure studies ever performed. See Section 1.5.2 for a description of OLCF liaison activities.

Total 2012 Usage: 51 million Jaguar core hours

Capability Utilization: <20% = 31%, 20% to 60% = 65%, 60% to 100% = 4%

Related Publication:

J. Erler, N. Birge, M. Kortelainen, W. Nazarewicz, E. Olsen, A.M. Perhac, and M. Stoitsov, “The limits of the nuclear landscape,” *Nature* **486**, 509 (2012).

Online Stories: <http://science.energy.gov/np/highlights/2012/np-2012-06-a/> and <http://www.olcf.ornl.gov/2012/06/27/ornlutk-team-maps-the-nuclear-landscape/>

3.2.3 Simulating Paleoclimate Shows Carbon Dioxide Drove Warming at End of Last Ice Age

Warren Washington, National Center for Atmospheric Research, INCITE

Objective: The correlation between carbon dioxide and temperature in Antarctic ice cores has created a mystery surrounding carbon dioxide’s role in the last deglaciation. These simulations were run to explain the global temperature record during that time to discover whether carbon dioxide drives warming. Other objectives included producing a more comprehensive global temperature dataset and simulating Earth system energy transport mechanisms.

Impact: This research was the first to definitively show the role that carbon dioxide played in helping to end the last ice age. Simulations found that global temperature mirrored and generally lagged behind rising carbon dioxide during the last deglaciation, which points to carbon dioxide as the major driver of global warming.

Accomplishments: This research, which started in 2007 on Phoenix, transitioned to Jaguar in late 2009, and concluded early in 2012, shows for the first time that an Intergovernmental Panel on Climate Change–class Coupled General Circulation Model used to predict climate’s future is capable of reproducing its past (Shakun 2012). Simulations created an innovative reconstruction of the global temperature record and explained the lag in global surface temperature in response to carbon dioxide. Simulations also showed that increased insolation disrupted the Atlantic Meridonal Overturning Current (AMOC), demonstrating how AMOC disruption contributed to the “bipolar seesaw” in global temperature patterns (He 2013). The transient Earth System Model simulation was a 22,000-year simulation reproducing reconstructed global temperature response to carbon dioxide and explaining the misleading lag of carbon dioxide behind Antarctic temperature records and pointing to carbon dioxide’s role in driving global climate change over glacial cycles. This work represents important elements within Feng He’s PhD dissertation (He 2011).

OLCF Contributions: This paleoclimate study accumulated a simulation dataset of nearly 300 TB that was analyzed on Lens to produce these results. “Our project could have only been done using [OLCF] resources given the computational and storage requirements.”—Thomas Bettge in project’s Quarterly Report to the OLCF

Total 2012 Usage: 48 million Jaguar core hours

Capability Utilization: <20% = 99%, 20% to 60% = 1%, 60% to 100% = 0%




Figure 3.3. Global Paleoclimate Reconstructure of Temperature and CO₂ and Jaguar Simulation at OLCF. (Simulation suggests that CO₂ drove global warming at the end of the last ice age.)

Related Publications:

- F. He, J. Shakun, P. Clark, A. Carlson, Z. Liu, B. Otto-Bliesner, J. Kutzbach, "Northern hemisphere forcing of southern hemisphere climate during the last deglaciation," *Nature* **494**, 81 (2013).
- J. Shakun, P. Clark, F. He, S. Marcott, A. Mix, Z. Liu, B. Otto-Bliesner, A. Schmittner, E. Bard, "Global warming preceded by increasing carbon dioxide concentrations during the last deglaciation," *Nature* **484**, 49 (2012).
- P. Clark, J. Shakun, P. Baker, P. Bartlein, S. Brewer, E. Brook, A. Carlson, H. Cheng, D. Kaufman, Z. Liu, T. Marchitto, A. Mix, C. Morrill, B. Otto-Bliesner, K. Pahnke, J. Russell, C. Whitlock, J. Adkins, J. Blois, J. Clark, S. Colman, W. Curry, B. Flower, F. He, T. Johnson, J. Lynch-Stieglitz, V. Markgraf, J. McManus, J. Mitrovica, P. Moreno, and J. Williams, "Global climate evolution during the last deglaciation," *P. Nat. Acad. Sci.*, published online before print (February 13, 2012).
- S. Marcott, P. Clark, L. Padman, G. Klinkhammer, S. Springer, Z. Liu, B. Otto-Bliesner, A. Carlson, A. Ungerer, J. Padman, F. He, J. Cheng, A. Schmittner, "Ice-shelf collapse from subsurface warming as a trigger for Heinrich events," *P. Nat. Acad. Sci.* **108**, 13415 (2011).
- F. He, "Simulating transient climate evolution of the last deglaciation with CCSM3 (TraCE-21K)," PhD dissertation, University of Wisconsin–Madison (2011).

Online Stories: <https://www.olcf.ornl.gov/2012/04/04/carbon-dioxide-caused-global-warming-at-ice-ages-end-pioneering-simulation-shows/> and <https://www.olcf.ornl.gov/2013/02/08/lessons-from-the-past/>

3.2.4 Showing Sea-Level Rise Will Continue Even with Aggressive Emission Mitigation

Warren Washington, National Center for Atmospheric Research, INCITE

Objective: Based on four climate change mitigation scenarios, researchers aimed to determine how much sea level will rise due to thermal expansion, which occurs as the water molecules in the sea warm and expand. The simulations were used to examine the mechanisms involved in sea-level rise, including ice melt and thermal expansion.

Impact: The research helped to draw attention to the under-appreciated variance involved in global temperature differences and sea-level-rise mitigation. The researchers were able to show that even if aggressive mitigation measures are taken, sea level is predicted to rise even after global average temperatures level off, but that sea-level rise will not be as dramatic if less carbon is emitted and more is sequestered.

Accomplishments: This project predicted, with quantified uncertainties, the amount of sea-level rise that will occur due to thermal expansion in response to mitigation strategies (Meehl 2012). This breakthrough has spurred greater research into contributions to sea-level rise from ice melt. The results of the simulations can also be used to push for more aggressive emission mitigation measures in order to avoid increased damage to property and life from rising oceans.




Figure 3.4. Sea Level Rise through 2300. (Sea level will continue to rise due to thermal expansion into the year 2300 under the most aggressive mitigation scenario, [cooling: RCP2.6], but the rise will be slowed enough to allow adaptation measures. Under less aggressive mitigation [stable: RCP4.5] or business as usual [warming: RCP8.5], there will be much less time for adaptation.)

OLCF Contributions: To draw their conclusions, the scientists needed models to go out to year 2300 running five simulations for each of the four mitigation scenarios. All told, they simulated a total of 4,500 years, which would have been impossible without a powerful supercomputer like Jaguar, whose speed increased the realism of climate simulations.

Total 2012 Usage: 48 million Jaguar core hours

Capability Utilization: <20% = 99%, 20% to 60% = 1%, 60% to 100% = 0%

Related Publications:

G.M. Meehl, A.Hu, C. Tebaldi, J.M Arblaster, W.M. Wahsington, H. Teng, B.M. Sanderson, T. Ault, W.G. Strand, J.B. White III, "Relative outcomes of climate change mitigation related to global temperature versus sea-level rise," *Nature Clim. Change* **2**, 576 (2012).

G.A. Meehl, J.M. Arblaster, J.T. Fasullo, A. Hu, K.E. Trenberth, "Model-based evidence of deep-ocean heat update during surface-temperature hiatus periods," *Nature Clim. Change* **1**, 360 (2011).

Online Stories: <https://www.olcf.ornl.gov/2013/01/24/sea-level-rise-will-continue-even-with-aggressive-emission-mitigation/> and <https://www2.ucar.edu/atmosnews/research/7557/expanding-seas>

3.2.5 Simulations Map Nuclear Pasta in a Core-Collapse Supernova

Tony Mezzacappa, Oak Ridge National Laboratory and University of Tennessee, INCITE

Objective: Map the critical densities and temperature for neutron-rich dense matter for both the onset of the inhomogeneous “pasta phase”—consisting of neutron-rich heavy nuclei and a free neutron and electron gas—and its dissolution to a homogeneous neutron, proton, and electron liquid. A core-collapse supernova is the last stage of life of a massive star. In the dying star, matter is compressed to densities exceeding the density of atomic nuclei and exposed to extreme temperatures and pressures. It has been proposed that, at a certain stage of the collapse, matter at these high densities self-organizes into what is known as “nuclear pasta,” a collection of bizarre structures, such as rods, slabs, and cylindrical and spherical holes (bubbles), which may constitute 10–20% of the inner core of the collapsing star.

Impact: The formation of the nuclear pasta phase of nuclear matter has profound consequences for neutrino transport. Because neutrinos are understood to play a crucial role in a supernova explosion, the model predicted in this work will contribute toward a more realistic description of the dynamics of a collapsing star. This research created the equation of state to be used as an input to core-collapse supernova simulations. This information will enable the understanding of the cataclysms responsible for seeding the universe with most of its elements.

Accomplishments: This team developed a fully self-consistent microscopic theory that describes the formation of nuclear pasta as the density increases. Its model predicted all the structures identified in previous studies and found evidence for a never-before-seen formation with a “cross-rod” shape (Pais 2012). The project identified, fully self-consistently, the onset of the pasta phase in inhomogeneous core-collapse supernova matter consisting of neutron-rich heavy nuclei and a free neutron and electron gas and its dissolution to homogeneous neutron, proton, and electron liquid. This work represents important elements of Helena Pais’s PhD dissertation (Pais 2013).




Figure 3.5. Nuclear Matter Density Distributions for Several Different “Pasta Phases.”
(Pais and Stone discovered the center “cross-rod” shape in these breakthrough simulations.)

OLCF Contributions: See Section 1.5.2 for contributions from OLCF’s Liaison Program in topics such as visualization and load balancing. The OLCF Resource Utilization Council prioritized this work to enable a PhD student to meet critical deadlines in the development of her PhD dissertation.

Total 2012 Usage: 28 million Jaguar core hours

Capability Utilization: <20% = 50%, 20% to 60% = 50%, 60% to 100% = 0%

Related Publication:

H. Pais and J.R. Stone, “Exploring the nuclear pasta phase in core-collapse supernova matter,” *Phys. Rev. Lett.* **109**,151101-1 (2012).

H. Pais, “Exploring the nuclear pasta phase in core-collapse supernova matter,” PhD dissertation, University of Tennessee (2013).

Online Story: <http://physicsbuzz.physicscentral.com/2012/10/nuclear-pasta-now-available-at-your.html>

3.2.6 Simulations Lead to Experimental Verification of Bose Glass

Tommaso Roscilde, Ecole Normale Supérieure de Lyon, INCITE

Objective: Using ORNL’s Jaguar supercomputer, researchers simulated a doped magnet with strong electronic correlation near absolute zero temperature, clarifying the conditions necessary for a Bose glass. In this state impurities allow quantum quasiparticles to condense into local regions, rather than across the magnet as a whole. Using information from the simulations, both the conditions under which a Bose glass would manifest and the observables—such as magnetization and specific heat—were determined.

Impact: Strongly correlated materials are a wide class of electronic materials that show unusual (often technologically useful [e.g., superconductivity and superfluidity]) electronic and magnetic properties, such as metal-insulator transitions or half-metallicity. The essential feature that defines these materials is that the behavior of their electrons cannot be effectively described in terms of noninteracting entities. Rather, the problem must be recast into new entities called “quasiparticles.” Theoretical models of the electronic structure of strongly correlated materials must include electronic correlation to be accurate. That is, each electron has a strong influence on its neighbors.

The Bose–Hubbard model gives an appropriate, approximate description of the physics of interacting bosons on a lattice. It is closely related to the Hubbard model, which originated in solid-state physics as an approximate description of superconducting systems and the motion of electrons (which are fermions) between the atoms of a crystalline solid. At very low temperatures, the Bose–Hubbard model (in the absence of disorder) is in either a Mott insulating state at weak particle-particle coupling or in a superfluid state at large coupling. The Mott insulating phases are characterized by integer boson densities, the existence of an energy gap for particle-hole excitations, and zero compressibility. In the presence of disorder, a third, “Bose glass,” phase exists. The Bose glass phase is characterized by a finite compressibility—the absence of an energy gap—and by infinite superfluid susceptibility. It is insulating despite the absence of a gap, as low tunneling probability prevents the generation of excitations, which, although close in energy, are spatially separated.

Accomplishments: This project performed a comprehensive quantum Monte Carlo study of the Bose glass in a doped quantum magnet within a strong magnetic field and calculated the quantitative features associated with such a state. These calculations enabled the experimental verification of Bose glass in real material systems as well as the tuning of the system between the Bose glass phase and Mott insulating phase by varying the dopant level in the magnetic material. Previous

realizations of this strongly correlated phenomenon were achieved in ultracold gas atoms confined in optical lattices, not in real quantum material systems.




Figure 3.6. Phase Diagram in the Magnetic Field (H) vs. Temperature (T) Plane. (Experimental phase diagram of the quantum magnet material compared with quantum Monte Carlo [QMC] simulations on Jaguar: Bose-Einstein condensate [BEC], Bose glass [BG], and Mott insulator [MI]).

Total 2009 INCITE Usage: 1.3 million Jaguar core hours

Capability Utilization: <20% = 31%, 20% to 60% = 69%, 60% to 100% = 0%

Total 2011 Director's Discretionary Usage: 12 million Jaguar core hours

Capability Utilization: <20% = 23%, 20% to 60% = 77%, 60% to 100% = 0%

Related Publications:

R. Yu, L. Yin, N.S. Sullivan, J.S. Xia, C. Huan, A. Paduan-Filho, N.F. Oliveira Jr., S. Haas, A. Steppke, C.F. Miclea, F. Weichert, R. Movshovich, E-D. Mun, B.L. Scott, V.S. Zaph, T. Roscilde, "Bose glass and Mott glass of quasiparticles in a doped quantum magnet," *Nature* **489**, 379 (2012).

R. Yu, C.F. Miclea, F. Weickert, R. Movshovich, A. Paduan-Filho, V.S. Zapf, R. Roscilde, "Quantum critical scaling at a Bose-glass-superfluid transition: Theory and experiment for a model quantum magnet," *Phys. Rev. B* **86**, 134421 (2012).

3.2.7 Accelerating Turbomachinery

Allan Grosvenor, Ramgen Power Systems, ALCC

Objective: Design a turbocompressor and turbogenerator with shock-wave-based technology leading to dramatically lower costs and higher efficiency while at the same time contributing to DOE goals for reducing carbon capture and sequestration costs.

Impact: Shock-wave-based turbomachinery has the potential to reduce the capital cost of carbon dioxide compression for carbon capture and sequestration by 50% to reduce the operating costs by 25%. Similarly, using this new turbomachinery technology in a 400-megawatt clean coal plant could lead to capital costs savings of \$22 million and an estimated \$5 million in cost-of-operation savings.

Accomplishments: Ramgen, NUMECA International, and the OLCF have transformed the workflow of this turbomachinery design project in a way that exploits the strengths of Jaguar/Titan. This process has involved performance improvements in the simulation code and memory reductions per core to fully utilize nodes. All of these improvements have enabled the use of intelligent optimization

techniques in which ensemble simulations of varying design parameters are combined into a single run on Jaguar/Titan capable of utilizing more than 240,000 cores. This extreme level of performance is in stark contrast to when Ramgen first partnered with the OLCF and could run the NUMECA CFD code on just 40 cores. The analysis of these ensembles drives the Ramgen designs toward more optimal configurations, leading to accelerated timelines for product development and deployment. Without the use of Jaguar, Ramgen's aerodynamic design timeline would not have been possible. The collaboration took an accelerated computational design cycle for turbomachinery from months to 8 hours. "The use of Jaguar has cut the projected time from concept to a commercial product by at least 2 years and the cost by over \$2 million," said Ramgen's CEO and Director Doug Jewett.

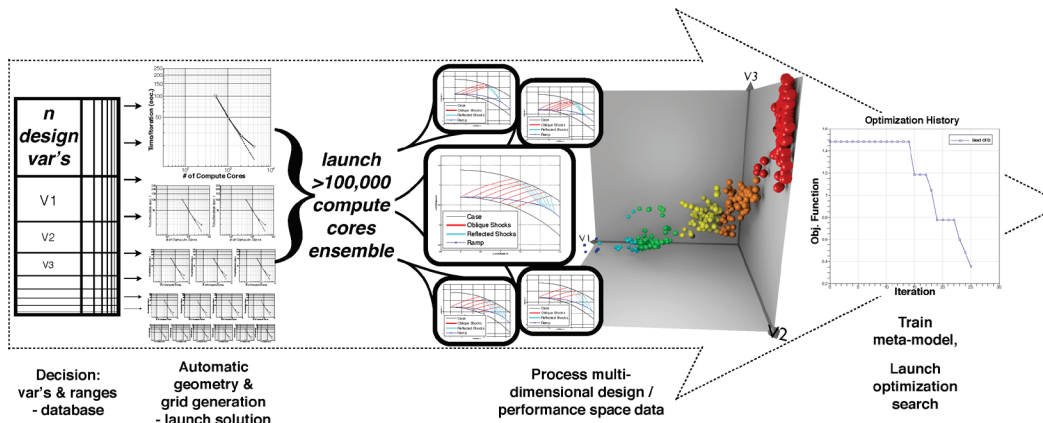


Figure 3.7. Ramgen Applies Aerodynamic Design Optimization Techniques on Jaguar. (The workflow begins with the specification of design variables. A large database is constructed of designs that reflect combinations over specified ranges. CFD solutions are performed and results collected to construct an approximate or "meta-model" of the resultant performance space. This space is then searched for designs predicted to offer high performance.)

OLCF Contributions:

The immense amount of data generated by the Ramgen project has created the need for improved analysis capabilities. The OLCF has provided expertise in visualization in areas such as shock wave volume rendering, development of new boundary layer detection techniques, and statistical analysis. (See Section 1.5.2 for details.) Mike Matheson performed much of the work described here. Ramgen sponsored Matheson through a Work for Others project for effort above and beyond that which is customary for an ALCC project.

Total 2012 Usage: 35 million Jaguar core hours

Capability Utilization: <20% = 27%, 20% to 60% = 19%, 60% to 100% = 54%

Related Publication:

A.D. Grosvenor, A.A. Zheltovodov, M.A. Matheson, L.M. Sailer, M.Krzysztopic, D.P. Gutzwiller, "Verification for a series of calculation 3D shock wave/turbulent boundary layer interactions flows," in Proceedings of 4th European Conference for Aerospace Sciences (EUCASS).

Online Story: <https://www.olcf.ornl.gov/2012/08/14/ramgen-simulates-shock-waves-makes-shock-waves-across-energy-spectrum/>

3.2.8 Cutting Costs and Risks in Drug Discovery and Development

*Jerome Baudry, University of Tennessee and Oak Ridge National Laboratory,
Director's Discretionary*

Project Name: High-Performance Computing for Rational Drug Discovery and Design

Objective: Find the best match between a molecular compound and its targeted protein receptor in hopes of creating drugs with a higher degree of specificity and less cross-reactivity, as well as exploring alternative uses for existing drugs. A team led by Jerome Baudry of the University of Tennessee–ORNL Center for Molecular Biophysics adapted widely used existing software to allow supercomputers such as ORNL's Jaguar to sift through immense molecular databases and pinpoint chemical compounds as potential drug candidates.

Impact: Less than 1% of drugs starting in the lab make it to market. Bringing a new drug to market can cost a few hundred million dollars to more than a billion and take from 10 to 15 years. Virtual high-throughput simulations may drastically decrease the time and money spent bringing new, more effective drugs to market. This work is supported by the National Institutes of Health (NIH) to look for chemicals that could treat prostate cancer. The research is funded by an NIH Clinical Translational Science Award, which was awarded to Georgetown and Howard Universities and includes ORNL, Med/Star Health, and the Washington, DC, Veterans Affairs Medical Center as key partners.

Accomplishments: The Jaguar supercomputer successfully screened 2 million compounds against a targeted receptor in a matter of days, as opposed to the months that would have been required for computing clusters and even longer for conventional test tube methods. The calculations allowed scientists to account for specific binding in protein receptors as well as the structural variations that occur within the receptor. The project created a vast library of molecular compounds that can be used for future screenings of potential drug candidates. This project is the dissertation thesis topic for PhD candidate Sally Ellingson, University of Tennessee.

ORNL scientists hope to be running the codes with performance of approximately 20 petaflops in 2013, a nearly tenfold increase in computational power. At that rate Titan would be able to screen 20 million compounds against a targeted receptor in just 1 day. Baudry says, "We will be able to simulate going inside patients and inside their cells instead of the test tube, and that's a revolution."

Total 2012 Usage: 8 million Jaguar core hours

Capability Utilization: <20% = 2%, 20% to 60% = 98%, 60% to 100% = 0%

Related Publications:

- S. Ellingson. "Acceleration and accessibility of virtual high-throughput molecular docking," PhD dissertation, University of Tennessee, doctorate expected 2014.
- S. Ellingson. 2012. "Accelerating virtual high-throughput ligand docking: Screening one million compounds using a petascale supercomputer," presented at 2012 Emerging Computational Methods for Life Sciences Workshop at HPDC12 in Delft, Netherlands.
- B. Collignon, R. Schulz, J. Smith, J. Baudry, "Task-parallel message passing interface implementation of Autodock4 for docking of very large databases of

compounds using high-performance super-computers,” *J. Comp. Chem.* **32**, 1202 (2010).

Online Story: <https://www.olcf.ornl.gov/2012/10/18/big-computing-cures-big-pharma/>

3.2.9 Providing Modeling and Simulation for CASL: Large-Eddy Simulation Investigation of Grid-to-Rod Fretting

John Turner and Doug Kothe, Oak Ridge National Laboratory, ALCC

Objective: This work has focused on a series of efforts progressing toward a simulation capability using Hydra-TH for the grid-to-rod fretting (GTRF) problem in light water reactors. GTRF is a complicated phenomenon that includes flow excitation force, nonlinear mechanical vibration, tribology, changing irradiated material properties, and fuel assembly geometry. As the coolant flows through the fuel assembly, fluid forces generated by the flow field induce fuel rod vibration. The flow-induced vibration causes small relative motions between grid supports and fuel rods, leading to fretting wear and, ultimately, fuel rod leakage. The primary objective of this work was to perform implicit large-eddy simulations (LESs) of the high Reynolds number turbulent flow in a fuel rod bundle and around the supporting spacer-grid structure. Unstructured meshes ranging from 2 million to 100 million grid points have been used to study the turbulent flow field and compute time-dependent forces on the fuel rod and spacer grid. These forces are used to compute the wear at the contact surface between the fuel rod and support structures in the reactor core.

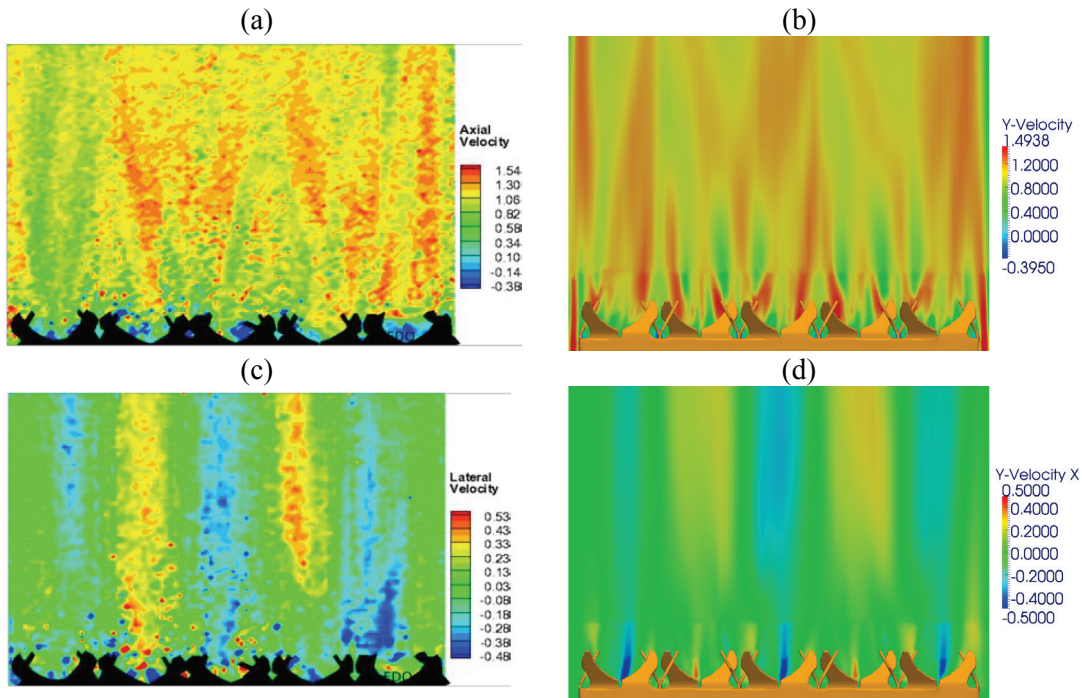


Figure 3.8. Experimental and Computer Axial (y-Direction) Time-Averaged Velocities. [a] experimental axial velocity, [b] Hydra-TH y-velocity, [c] experimental lateral velocity, [d] Hydra-TH x-velocity.)

Impact: The Consortium for Advanced Simulation of Light Water Reactors’ (CASL’s) primary modeling-and-simulation science driver is to enable nuclear reactor power uprates, life extensions, and higher fuel burn-up. Accordingly, a series of ten

challenge problems have been defined that address these areas. To solve the challenge problems, the modeling and simulation requirements for the required physics have been codified in a series of benchmark test problems. By solving these test problems, CASL will be positioned to investigate the challenge problems. This accomplishment is related to one of these benchmark problems.

In addition to providing validation of the Hydra-TH simulation capabilities, this work will also provide a sensitivity analysis of the fluid forces on the computed work wear rates on the fuel rods, which will enable CASL to develop advanced GTRF methods with higher-fidelity structural mechanics simulations coupled to the LES forces provided by Hydra-TH.

Accomplishments: This work demonstrated that fuel-rod acceleration, velocity, and displacements, using the fluid forces computed with LES, are predicted with a high degree of accuracy. The Westinghouse structural dynamics code, VITRAN, was used to test the sensitivity of the work wear rate for the fuel rods to the fluid forces, and it was determined that the largest differences were approximately 0.17%. In addition, use of experimental data for a 5-by-5 rod bundle and spacer grid has provided experimental validation of the LES capabilities provided by Hydra-TH.

OLCF Contributions:

Total 2012 Usage: 21.9 million Jaguar core hours

Capability Utilization: <20% = 76%, 20% to 60% = 24%, 60% to 100% = 0%

Related Publications:

M.A. Christon, J. Bakosi, N. Barnett, M.M. Francois, R.B. Lowrie, *Initial Assessment of Hydra-TH on Grid-to-Rod Fretting Problems*, Los Alamos National Laboratory, LA-UR 11-07034, December 2011.

J. Bakosi, M.A. Christon, M.M. Francois, R.B. Lowrie, R.R. Nourgaliev, *GTRF Calculations using Hydra-TH*, Los Alamos National Laboratory, LA-UR 12-24526, September 2012.

R. Nourgaliev, M.A. Christon, J. Bakosi, R.B. Lowrie, L.A. Pritchett-Sheats, "Hydra-TH: A thermal-hydraulics code for Nuclear Reactor Applications," submitted to NURETH-15 Conference, May 12–17, 2013.

J. Bakosi, M.A. Christon, R.B. Lowrie, L.A. Pritchett-Sheats, R.R. Nourgaliev, "Large-eddy simulations of turbulent flow for grid-to-rod fretting in nuclear reactors," submitted to *Nuclear Engineering and Design*, 2013.

3.3 Allocation of Facility Director's Reserve

2012 Operational Assessment Guidance

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

The OLCF allocates time on leadership resources primarily through the INCITE program and through the facility's DD program. The OLCF seeks to maximize scientific productivity via capability computing through both programs. Accordingly, a set of criteria are considered when 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, up to 30% of the facility's resources are allocated by the ASCR office through the ALCC program.

3.3.1 Director's Discretionary Program

The goals of the DD program are threefold: development of strategic partnerships, preparation for leadership computing competitions (i.e., INCITE and ALCC), and application performance development and measurement. These goals are aligned with particular strategic goals for the OLCF, namely the expansion of the leadership computing science community and enhancement of the pervasive use of leadership computing in a variety of scientific fields.

Strategic partnerships are those aligned with strategic and programmatic ORNL directions. They may be entirely new areas with respect to HPC or ones in need of nurturing. Example candidate projects are those associated with the ORNL Laboratory Directed Research and Development Program, programmatic science areas (bioenergy, materials discovery, nanoscience, climate, nuclear engineering, and energy technology), and key academic partnerships (e.g., that with the ORNL Joint Institute for Computational Sciences). The newly awarded Critical Materials Institute hub led by Ames National Laboratory is an example of an emerging strategic partnership. Included in this broad category is the Industrial HPC Partnerships Program (see below), providing opportunities for researchers in industry to access the leadership-class systems to carry out work that would not otherwise be possible.

The DD program is also accessible by the general HPC community to carry out porting and development exercises for nascent and less-efficient applications. These performance enhancement projects range in scope from immediate INCITE preparation—designed to allow investigators the opportunity to test their codes' scalability on INCITE platforms—to somewhat longer-term projects involving improvement in algorithms and implementations. As examples of DD program outcomes in expanding the leadership computing science community, two of the five new INCITE PIs at the OLCF in 2013 (Jeroen Tromp, Princeton University, and Zan Luthey-Schulten, University of Illinois, Urbana-Champaign) had their initial experience on OLCF leadership computing resources through the DD program

The OLCF DD program also supports a variety of “data-only” projects that require data storage and bandwidth capabilities, but few compute resources. Ongoing data-only projects include the Earth System Grid Federation, Data Sharing Project for the Center for Exascale Simulation of Combustion in Turbulence (ExaCT) codesign project, and the Majorana Demonstrator Secondary Data Archive. In addition, infrastructure software such as frameworks, libraries, and application tools and support research 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 for 1 year or less. The typical size of DD awards is roughly 1 million core hours but can range from tens of thousands of hours to 4 million hours or more.

Since its inception in 2006, the DD program has granted allocations in virtually all areas of science identified by DOE as strategic for the nation (Table 3.2). Additional allocations have been made to promote science education and outreach. Requests and awards have grown steadily each year (Table 3.3).

Table 3.2. Director's Discretionary Program: Domain Allocation Distribution

Time Period	Biology	Chemistry	Computer Science	Earth Science	Engineering	Fusion	Materials Science	Nuclear Energy	Physics
2008	19%	8%	28%	4%	8%	15%	3%	1%	14%
2009	5%	3%	19%	6%	8%	6%	33%	1%	19%
2010	9%	6%	10%	8%	19%	6%	16%	3%	23%
2011	7%	1%	10%	19%	14%	0%	9%	13%	26%
2012	6%	1%	21%	14%	25%	5%	10%	1%	18%

Annual DD allocations are typically less than the available hours; that is, all of the DD time is not allocated at the beginning of the calendar year. With this approach, the OLCF can remain flexible and responsive to new project requests and research opportunities that arise during the year. The leadership computing resources continue to be effectively used under this approach, as INCITE and ALCC users are not “cut off” when they overrun their allocation. Rather, they are allowed to continue running at lower priority to make use of potentially available time.

Table 3.3. Director's Discretionary Program: Awards and User Demographics

Year	Project Awards	Project Requests	Hours Available (million)	Hours Allocated (million)	User Demographics (%)	
2008	36	38	18.33	8.5	42.7	DOE
					3.8	Government
					6.4	Industry
					47.1	Academia
2009	47	51	125	38	55.9	DOE
					0.7	Government
					9.9	Industry
					33.5	Academia
2010	77	85	160	85	46.0	DOE
					2.3	Government
					12.2	Industry
					39.5	Academia
2011	57 had carryover and 43 had new awards for a total of 100 projects	57	160	139	41	DOE
					4	Government
					4	Industry
					50	Academia
					1	Other
2012	31 had carryover and 79 had new awards for a total of 110 projects	82	160	148	35	DOE
					1	Government
					8	Industry
					54	Academia
					2	Other

3.3.2 Industrial HPC Partnerships Program

The Industrial HPC Partnerships Program completed its fourth year in 2012 and continues to attract new users as well as new projects from previous users. Twenty-one projects were under way during the year (nine continued from previous years), with companies applying for time through INCITE, ALCC, and the DD allocation

processes. Companies like GE and United Technologies, who have seen measurable results from multiple projects, now begin planning for INCITE and ALCC prior to calls for proposals, including budgeting internal funds in advance to support this work should proposals be selected.

This year four companies applied for and received ALCC awards at the OLCF: United Technologies Research Center, Ramgen Power Systems, GE Global Research, and Global Foundries. Out of 29 ALCC awards at all three DOE/ASCR HPC centers, Ramgen received the fifth largest allocation (40 million core hours) and GE received the sixth largest (34 million core hours), validating that industry has large, complex problems that require leadership-level computational resources and that companies can successfully compete for large awards. For instance, Procter & Gamble (P&G) is participating in an INCITE project led by Temple University (“Coarse Grained Molecular Dynamics Studies of Vesicle Formation and Fusion”).

New DD awards were made to GE Global Research, United Technologies Research Center, Global Foundries, FM Global, Ford, and GM. The discretionary allocation option continues as a jumping-off point for companies to gain experience with large-scale modeling and simulation. Ford and GM each received their second discretionary awards. Ford is studying engine cycle-to-cycle variability, and GM is using Titan to optimize and design fuel injectors. Both are also collaborating with an OLCF researcher and the laboratory’s National Transportation Research Center with support from DOE/Energy Efficiency and Renewable Energy’s Vehicle Technology Program. Results from these projects will help these firms design more fuel-efficient automobiles.

Ford also completed a project, “Large Scale Engine Bay Package Optimization,” that enabled the firm to develop a new methodology to optimize a vehicle’s cooling package and for the first time to run an optimization of a complete auto cooling system. In 2012 Ford used 1 million core hours on Jaguar to perform some of its largest simulations to date, which are now being analyzed. These analyses will help accelerate product development, bringing concepts to market faster and at lower costs. “We ran a set of conditions that needed 1,500 to 1,600 analyses. These were more than other systems could deliver in a reasonable time,” said Alex Akkerman, senior technical specialist at Ford Motor Company. Incorporating Jaguar helped demonstrate the value of new computationally intensive optimization techniques. For the first time researchers were able to compute a high number of design parameters with multiple vehicle operating conditions in a single run. The team also established a validated computational fluid dynamics (CFD) design method for developing future Ford vehicles. Results have also provided important return-on-investment justification for Ford to upgrade its internal HPC computing capability so that researchers can run larger-scale optimization problems in house.

New industrial user FM Global is a top-rated global property insurance firm. Its project, “CFD Modeling of Industrial Scale Fire Growth and Suppression,” harkens back to the company’s origins in 1835 when founder Zacharia Allen, a prominent textile mill owner, made significant property improvements to minimize the chance of fire loss. When his request for a reduction in insurance rates was denied, he and a group of mill owners launched their own insurance firm, known today as Factory Mutual Insurance company (conducting business as FM Global). The company is using OLCF resources to study combustion to better understand how large-scale industrial fires spread.

Two United Technologies Research Center projects spawned peer-reviewed papers. “Nanostructured Catalyst for WGS and Biomass Reforming Hydrogen

Production” fostered a new article entitled “GGA+U method from first principles: Application to reduction–oxidation properties in ceria-based oxides,” which appeared in 2012 in the *Journal of Materials Science*. (Two other articles previously appeared in *Nano Letters* and *The Journal of Physical Chemistry Letters*.) This project applied for time on Jaguar at least in part to gain access to a version of Vienna Ab initio simulation package (VASP) optimized by ORNL’s Paul Kent. “VASP on Jaguar was probably the best optimized compilation of this program that exists,” PI Amra Peles said. “Using VASP on Jaguar we were able to increase our model size by a factor of 10.” Results from the project “Next Generation Turbulent Reactive Flow Simulation Capability” were captured in “Is LES of Reacting Flows Predictive? Part 1: Impact of Numerics.” This paper was accepted for publication for the 2013 conference proceedings of the 51st AIAA Aerospace Sciences Meeting, sponsored by the American Institute of Aeronautics and Astronautics. This prestigious industry event gathers aerospace scientists and engineers from around the world to share and disseminate the latest scientific knowledge and research.

Access to OLCF systems and experts through our Industrial HPC Partnership Program continues to help companies address cutting edge science and engineering problems that are making their products, their companies, and the country more competitive. In turn, this enhances the return the nation receives from its investment in the OLCF.

3.3.3 2012 Director’s Discretionary Allocations

Table 3.5 provides a list of the DD allocations for 2012.

Table 3.4. Industry Projects at the OLCF

Corporate Partner	Program	Description
P&G (co-PI)*	INCITE Renewal	Coarse Grained Molecular Dynamics Studies of Vesicle Formation and Fusion
Ramgen Power Systems*	ALCC	Supercomputer-Enabled Accelerated Development of Revolutionary Supersonic Shock Wave-Based Turbomachines: Achieving DOE Goals for Compressing Carbon Dioxide and Achieving High Energy Efficiency via High Resolution CFD
Boeing*	ALCC	Reliable Predication of Performance of High Lift Systems of Commercial Air
GE Global Research*	ALCC	High Fidelity Simulations of Gas Turbine Combustors for low emissions
United Technologies Research Center*	ALCC	Large-eddy simulation for turbomachinery - advancing state-of-the-art
Global Foundries	ALCC	Petascale Atomistic Simulations of Ultra Scaled Transistors
GE Global Research*	ALCC	Non-icing Surfaces for Cold Climate Wind Turbines
United Technologies Research Center	ALCC	Massively Parallel High-Fidelity Simulation of Spray Atomization
GE Global Research	ALCC	Impact of the Inlet Boundary Condition on High-Pressure Turbine Temperature Predictions
Caitin†	DD	Parallel Computing Performance Optimization for Complex Multiphase Flows in Strong Thermodynamic Non-equilibrium
BMI Corporation‡	DD	Smart Truck Optimization
Ford Motor Company*	DD	Large Scale Engine Bay Package Optimization
GE Global Research	DD	Large Eddy Simulation for Wind Turbine Interactions
United Technologies Research Center*	DD	Multiphase Injection
United Technologies Research Center	DD	Next Generation Turbulent Reactive Flow Simulation Capability
GM Global R&D	DD	Multi-hole injector optimization for spark-ignited direct-injection gasoline engines
Ford Motor Company	DD	Cycle-to-cycle Combustion Variation Modeling
FM Global	DD	CFD Modeling of Industrial Scale Fire Growth and Suppression
GE Global Research	DD	Tacoma Scalability for INCITE-sized Problems (INCITE preparatory project)
GE Global Research	DD	Fully explicit transient dynamics simulation of gas turbine structures with frictional contact on massively parallel supercomputers
Global Foundries	DD	Density functional studies of Si/SiGe interface structures

*Project launched prior to 2012.

†Project withdrawn early in 2012 before completion due to PI illness.

‡Majority of work completed prior to 2012.

Table 3.5. 2012 Director's Discretionary Allocations

PI	Affiliation	2011 Allocation	Carryover to 2012	New 2012 Allocation	Total 2012 Allocation	2012 Usage	Project Name
Michael Bussman	Helmholtz-Zentrum Dresden-Rossendorf			200,000	200,000	0	Laser-Wakefield Simulations Using PICONGPU
Mike Henderson	BMI Corporation	2,695,917	2,429,299		2,429,299	28,940	Smart Truck Optimization
Thomas Gieda	Caitin Inc.	500,000	430,895		430,895	19,852	Parallel Computing performance Optimization for Complex Multiphase Flows in Strong Thermodynamic Non-equilibrium
Alexander Akkerman	Ford Motor Company	1,000,000	999,770		999,770	953,531	Large Scale Engine Bay Package Optimization
Rainald Lohner	George Mason University	1,000,000	816,126		816,126	535,310	Highly Detailed Simulations of Blasts on Offshore Platforms
Dominic Von Terzi	GE Global Research			2,000,000	2,000,000	696,395	LES for Wind Turbine Interactions
Dana Hammond	NASA, Langley Research Center			600,000	600,000	577,764	Scaling of NASA CFD Application for Aeronautics
Bronson Messer	ORNL			6,000,000	6,000,000	4,601,669	Explosive Nucleosynthesis and Deflagration to Detonation in Type Ia Supernovae
Patrick Fragile	Oak Ridge Associated Universities	1,000,000	0	4,000,000	4,000,000	4,021,434	Radiation Transport in Numerical Simulations of Black-Hole Accretion Disks
Paul Sutter	University of Illinois	5,000,000	3,204,658		3,204,658	0	Exploring the origins of galaxy cluster magnetic fields
Ward Manchester	University of Michigan			500,000	500,000	155,793	Simulating Coronal Mass Ejections from the Convection Zone to the Earth
Michael Warren	Los Alamos National Laboratory			500,000	500,000	862,108	The Dark Sky Simulations
Casey Meakin	Los Alamos National Laboratory			750,000	750,000	948,609	An Investigation of Turbulence and Mixing in Stellar Interiors: Code Performance Study for INCITE 2013
Tiziana Di Matteo				1,000,000	1,000,000	0	Petascale Cosmology with P-Gadget
Benjamin Preston	ORNL			200,000	200,000	52,872	Quantifying Economic Losses Associated with Climate Extremes under Conditions of Climatic and Socioeconomic Change

Table 3.5. 2012 Director's Discretionary Allocations (continued)

PI	Affiliation	2011 Allocation	Carryover to 2012	New 2012 Allocation	Total 2012 Allocation	2012 Usage	Project Name
Moestasim Ashfaq	ORNL			5,000,000	5,000,000	1,611,975	A hierarchical regional modeling framework for decadal-scale hydro-climatic predictions and impact assessments
John Michalakes	National Renewable Energy Laboratory			500,000	500,000	0	Simulator for Offshore Wind Plant Applications (SOWFA)
Jason Hill	University of Minnesota			200,000	200,000	0	Air Quality Impacts of Conventional and Alternative Energy for Transportation
Jerome Baudry	ORNL			7,000,000	7,000,000	7,666,142	High-Performance Computing for Rational Drug Discovery and Design
Klaus Schulten	University of Illinois–Urbana-Champaign			3,000,000	3,000,000	5,213,500	Simulation of very large biomolecular assemblies
Gustavo Seabra	Universidade de Federal de Pernambuco			500,000	500,000	118,050	Elucidation of the Molecular Mechanism of Enzymatic Reactions by Molecular Dynamics and Hybrid Quantum Mechanical and Molecular Mechanics Simulations
Giuseppe Milano	Universita degli Studi di Salerno			50,000	50,000	0	GPU Accelerated Hybrid Particle Field Molecular Dynamics Simulations
Zaida Luthey-Schulten	UIUC			280,000	280,000	69,169	Optimization and benchmarking for SSU Biogenesis and lattice microbes
Stephan Irle	Nagoya University			200,000	200,000	313,308	Thermodynamic vs kinetic control in nanostructure growth and degradation
Erik Deumens	University of Florida	8,000,000	1,290,567		1,290,567	366,934	EOM-CC calculations on diamond nano crystals
Zhengyu Liu	University of Wisconsin Madison	2,000,000	0		0	62,277	Assessing Transient Global Climate Response using the NCAR-CCSM3: Climate Sensitivity and Abrupt Climate Change
Atul Jain	University of Illinois	30,000	28,448		28,448	15,617	Land Cover and Land Use Change and its Effects on Carbon Dynamics in Monsoon Asian Region
James Joseph Hack	ORNL	15,000,000	10,099,444		10,099,444	11,679,547	Ultra High Resolution Global Climate Simulation to Explore and Quantify Predictive Skill for Climate Means, Variability and Extremes
Aytekin Gel	ALPEMI Consulting	600,000	0		0	49,232	Mitigation of CO2 Environmental Impact Using a Multiscale Modeling Approach

Table 3.5. 2012 Director's Discretionary Allocations (continued)

PI	Affiliation	2011 Allocation	Carryover to 2012	New 2012 Allocation	Total 2012 Allocation	2012 Usage	Project Name
Katherine Evans	ORNL			3,500,000	3,500,000	3,274,344	A Scalable, Efficient, and Accurate Community Ice Sheet Model (SEACISM)
Francisco Doblas-Reyes	Institut Catala de Ciencies del Clima			500,000	500,000	547,294	High-resolution global climate simulation and prediction with EC-Earth
Colin Jones	Swedish Meteorological and Hydrological InstituteFolkborgsvagen			500,000	500,000	5,968	HIRES-CORDEX
Thomas Henderson	National Oceanic and Atmospheric Administration, Geophysical Fluid Dynamics Laboratory			10,000	10,000	1,652	GPU Computing for Numerical Weather Prediction
Robert Cook	ORNL			6,000	6,000	0	Modeling and Synthesis Thematic Data Center (MAST-DC)
Marios Soteriou	United Technologies Research Center	2,500,000	1,013,838		1,013,838	930,733	Multiphase Injection
Suresh Menon	Georgia Institute of Technology	1,000,000			0	32,010	Simulations of Detonation to Deflagration Transition (DDT) in Two-Phase Reactive Mixture and Supercritical Combustion in High Pressure Shear Co-axial Injector
Vaidyanathan Sankaran	United Technologies Research Center	1,000,000	967,329		967,329	133,781	Next Generation Turbulent Reactive Flow Simulation
Samuel Paolucci	University of Notre Dame	1,000,000	1,000,000		1,000,000	1,002,558	Reactive flows with detailed chemistry using an adaptive multiscale wavelet method
Ramanan Sankaran	ORNL			500,000	500,000	127	Simulating Combustion in Automotive Engines with Real Fuel Chemistry
Martin Berzins	University of Utah			1,000,000	1,000,000	1,110,143	Explosive Hazard Predictions with the Uintah Framework
Tang-Wei Kuo	General Motors			1,000,000	1,000,000	0	Multi-hole Injector Optimization for Spark-ignited Direct-injection Gasoline Engines
Brad VanDerWege	Ford Motor Company			1,000,000	1,000,000	0	Cycle-to-Cycle Combustion Variation Modeling
Sreekanth Pannala	ORNL			500,000	500,000	0	Computational Infrastructure for Parallel Simulations of Cycle-to-Cycle Variations of In-cylinder Combustion

Table 3.5. 2012 Director's Discretionary Allocations (continued)

PI	Affiliation	2011 Allocation	Carryover to 2012	New 2012 Allocation	Total 2012 Allocation	2012 Usage	Project Name
Sreekanth Pannala	ORNL			500,000	500,000	1,722	Parallel Computational Infrastructure for Optimizing Multi-hole Injector for Spark-ignited Direct-injection Gasoline Engines
Yi Wang	FM Global			1,000,000	1,000,000	0	CFD Modeling of Industrial Scale Fire Growth and Suppression
Thomas Maier	ORNL	10,000,000	0		0	23	Predictive simulations of cuprate superconductors
Gabriel Kotliar	Rutgers University			1,000,000	1,000,000	45,650	Calculation of Strongly Correlated Systems Using DMFT(CTQMC+WIEN2K) Method
Michael Widom	Carnegie Mellon University			10,000	10,000	0	VASP on GPU Systems
Jens Glaser	University of Minnesota			0	0	0	Optimization of a general-purpose molecular dynamics code running on multiple GPUs
Kalyan Perumalla	ORNL			550,000	550,000	871,463	An Evolutionary Approach to Porting Applications to Petascale Platforms
George I-Pan Fann	ORNL		0		0	78,232	Prototype Advanced Algorithms on Petascale Computes for IAA II
Stephen Poole	ORNL		0		0	0	FASTOS Community Allocation
Zizhong Chen	Colorado School of Mines		500,000		500,000	0	Fault Tolerant Linear Algebra Algorithms and Software for Extreme Scale Computing
Sean Ahern	ORNL		1,000,000		1,000,000	162,518	Large-Scale Data Analysis and Visualization
Terry Jones	ORNL			10,500,000	10,500,000	4,085,596	HPC Colony II
Stephen Scott	ORNL		1,000,000		1,000,000	0	Enabling Exascale Hardware and Software Design through Scalable System Virtualization
Vida Blair Sullivan	ORNL	250,000	248,653		248,653	230,246	Scalable Graph Decomposition and Algorithms to Support the Analysis of Petascale Data
Marc Snir	University of Illinois--Urbana-Champaign	100,000	100,000		100,000	92,963	Damaris
Terry Jones	ORNL	3,000,000	0		0	5,672	Extending Vampir IO for OLCF-3 Class Systems

Table 3.5. 2012 Director's Discretionary Allocations (continued)

PI	Affiliation	2011 Allocation	Carryover to 2012	New 2012 Allocation	Total 2012 Allocation	2012 Usage	Project Name
Joshua New	ORNL				500,000	11,582	Autotune E+ Buildings
Kalyan Perumalla	ORNL				2,000,000	1,142,445	ReveR-SES: Reversible Software Execution Systems for Ultra-scale Computing
Barbara Chapman	University of Houston				150,000	715	A similarity-based analysis tool for pattern derivation and large scale program restructuring
Patrick Joseph Burns	Colorado State University				30,000	448	Grad 511
George Biros	University of Texas–Austin				2,000,000	2,449,348	Fast N-body algorithms in high-dimensions
Olaf Schenk	Universita della Svizzera italiana				500,000	0	Large-Scale Seismic Imaging on HPC Architectures: Applications, Algorithms and Software
David Pugmire	ORNL				1,000,000	51,792	SDAV
Richard Mills	ORNL				1,000,000	1,538,158	Hierarchical Krylov Methods for Ultrascale Computers
Martin Burtcher	Texas State University–San Marcos				50,000	8,492	GPU Application performance and data analytics
Yuji Shinano	Zuse Institute Berlin				100,000	0	ParaSCIP
Jim Tallman	GE Global Research				1,000,000	831,263	Tacoma Scalability for INCITE-sized problems
Andreas Schaefer	Friedrich-Alexander-Universitaet Erlangen-Nuernberg (FAU)				500,000	0	LibGeoDecomp
Rajiv SamPATH	GE Global Research				0	0	Fully Explicit Transient Dynamics Simulation of Gas Turbine Structures with Frictional Contact on Massively Parallel Supercomputers
Choong-Seock Chang	Princeton Plasma Physics Laboratory				2,000,000	1,012,284	Implementation of Multiscale Fusion Gyrokinetic Code XGC1 on OLCF Hybrid Architecture
Zhihong Lin	University of California Irvine				4,000,000	4,035,790	Porting and scaling of GTC code on GPU-based architecture
David Green	ORNL				1,300,000	2,032,028	RF-SciDAC

Table 3.5. 2012 Director's Discretionary Allocations (continued)

PI	Affiliation	2011 Allocation	Carryover to 2012	New 2012 Allocation	Total 2012 Allocation	2012 Usage	Project Name
Steven Shannon	North Carolina State University				300,000	0	Particle-In-Cell Simulation of Radio Frequency Field Structure Near Plasma Facing Antenna Components
William Tang	Princeton University				5,000,000	1,032,447	GPU-CPU Global PIC
Thomas Jordan	University of Southern California	2,000,000	0		0	8,804	Deterministic Simulations of Large Regional Earthquakes at Frequencies up to 4Hz
Jeroen Tromp	Princeton University			100,000	100,000	0	Global Seismic Tomography based on Spectral-Element and Adjoint Method
Balint Joo	Jlab			300,000	300,000	12,225	Porting Lattice QCD Codes to Titan
Rong Tian	Institute of Computing Technology, Chinese Academia of Sciences			3,000,000	3,000,000	1,992,593	Petascale simulation of fracture process
Shok Srinivasan	Florida State University	300,000	300,000		300,000	8,137	Accelerating Quantum Monte Carlo on Massively Parallel Computing Platforms
Bruce Harmon	Ames Lab			1,000,000	1,000,000	385,740	Beyond Rare Earth Magnets (BREM)
Predrag Krstic	ORNL			1,200,328	1,200,328	80,733	Science of the Plasma-Material Interface at Extreme Conditions
Jacek Jakowski	University of Tennessee–Knoxville			2,000,000	2,000,000	603,319	Electronic structure calculation methods on accelerators
Benson Muete	University of Michigan			200,000	200,000	156,251	Numerical investigations of semilinear partial differential equations
Xiaoye Sherry Li	Lawrence Berkeley National Laboratory			200,000	200,000	63,029	Next Generation Computing for X-ray Science
Leonid Zhigilei	University of Virginia			1,000,000	1,000,000	1,950,450	Atomistic Simulations of Laser Material Interactions
Robert Patton	ORNL			1,000,000	1,000,000	0	Modeling & Simulation of Medicare & Medicaid Services
Biswas Sengupta	Indian Institute of Science			0	0	0	The role of constraints in the design of the nervous system
Shaikh Ahmed	Southern Illinois University–Carbondale			1,000,000	1,000,000	899,189	Multimillion-Atom Modeling of Harsh-Environment Nanodevices

Table 3.5. 2012 Director's Discretionary Allocations (continued)

PI	Affiliation	2011 Allocation	Carryover to 2012	New 2012 Allocation	Total 2012 Allocation	2012 Usage	Project Name
Bhagawan Sahu	Global Foundries US Inc.			3,500,000	3,500,000	809,470	Density Functional Studies of Si/SiGe interface structures
John Turner	DOE			7,000,000	7,000,000	3,474,599	Fundamental studies of multiphase flows and corrosion mechanisms in nuclear engineering applications
Calvin Johnson	San Diego State University	500,000	500,000		500,000	171,526	Large-scale configuration-interaction nuclear shell-model code with factorization algorithms
Dipankar Dutta	Mississippi State University			1,000,000	1,000,000	5,221	A New Search for the Neutron Electric Dipole Moment
Kenneth Read	ORNL			100,000	100,000	0	Probing Fluctuating Initial Conditions of Heavy-Ion Collisions
Bobby Sumpter	ORNL			6,000,000	6,000,000	821,389	Computational Resources for the Nanomaterials Theory Institute at the Center for Nanophase Materials Sciences and the Computational Chemical and Materials Sciences group in the Computer Science and Mathematics Division
Sreekanth Pannala	DOE			1,000,000	1,000,000	7	Using Solid Particles as Heat Transfer Fluid in CSP Plants
Homayoun Karimabadi	University of California--San Diego			11,000,000	11,000,000	7,879,953	Enabling Breakthrough Kinetic Simulations of the Earths Magnetosphere through Petascale Computing
Nikolai Pogorelov	University of Alabama--Huntsville			1,000,000	1,000,000	63,227	Flows of Partially Ionized Plasma in the Heliosphere and Astrospheres
Ramesh Balakrishnan	Argonne National Laboratory	1,000,000	1,000,000		100,000	174	The Performance of Turbulence Codes on Massively Parallel Computing Platforms with Multicore Processor Architectures
Oleg Zikanov	University of Michigan			400,000	400,000	408,300	Effect of Liquid-Phase Turbulence on Microstructure Growth during Solidification
George Vahala	College of William and Mary			950,000	950,000	998,141	Lattice Algorithms for Quantum and Classical Turbulence
Praveen Ramaprabhu	University of North Carolina	862,160	858,078		858,078	52,919	Simulations of turbulent mixing driven by strong shock waves
Pui-kuen Yeung	Georgia Institute of Technology	3,000,000	0		0	1,970,806	Frontiers of Computational Turbulence

Table 3.5. 2012 Director's Discretionary Allocations (continued)

PI	Affiliation	2011 Allocation	Carryover to 2012	New 2012 Allocation	Total 2012 Allocation	2012 Usage	Project Name
Misun Min	Argonne National Laboratory	900,000	899,998		899,998	1,405,391	Codes for High Order Methods
Shanti Bhushan	Mississippi State University			400,000	400,000	88,745	Hybrid CPU/GPU Parallelization of a Pseudo-Spectral Solver for Direct Numerical Simulations of Transitional Flow
Lance Collins	Cornell University			500,000	500,000	559,688	Direct numerical simulation of high-Reynolds-number, particle-laden turbulence
Peyman Givi	University of Pittsburg			500,000	500,000	15,094	US National Center for Hypersonic Combined Cycle Propulsion
Antonino Ferrante	University of Washington			200,000	200,000	0	Petascale DNS of high Reynolds number multi-phase turbulent flows
					148,323,431	92,264,636	

Innovation

HIGH PERFORMANCE COMPUTING FACILITY 2012 OPERATIONAL ASSESSMENT OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2013

4. INNOVATION

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

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

4.1 Innovation Summary

The OLCF provides leadership and collaboration in a number of areas that foster enhanced facility operations. Innovations in 2012 included developing tools to more effectively capture user publications and increase the efficiency of job scheduling; piloting projects to manage and analyze large data sets from other user facilities; teaming with vendors to enhance the design and use of the file system, contributions to archival storage system software, networks, and debuggers, where the collaborative results were included in the partners' product lines; and the creation of a workshop to share the expertise of the OLCF in project and program management.

4.2 Innovations in Publication Tracking

The OLCF collaborated with the Computational Data Analytics (CDA) Group at ORNL to develop a search capability that was specifically designed to capture the published results of research carried out at the center. The CDA Group has created a high-speed, multi-domain recommendation system known as Raptor. Given a large data repository, this tool helps the user identify relevant documents from that repository according to the user's defined interest. Originally sponsored by the Office of Naval Research, Raptor was extended to meet the acknowledgement-tracking needs of the user facility. As input to Raptor, the OLCF provided a list of facility users, keywords associated with the facility of interest, and the official acknowledgement statement used by authors in publications for the facility. Raptor then performed an automated, sophisticated, and comprehensive search and analysis of scientific publications and provided the OLCF with a list of the most relevant publications. An early application of this search system identified a paper published in *Nature* in September 2012 related to an OLCF user project with a large allocation of time on Jaguar in CY2009. We recognize that publications will continue to be generated years after a project is completed, and the challenge is to collect this valuable output long after the project team has evolved, moved to new organizations, or otherwise changed or dispersed. Therefore, searches of the scientific literature combined with standard user reporting channels provide a much more comprehensive survey of user publications than previously available. The search

capability was launched in 2012; in 2013 the OLCF plans to continue its development and share the technology with other user facilities.

4.3 Innovations in Scheduler Technology Transfer

ORNL has been a leader in the development of features for batch schedulers on Cray system. Adaptive Computing, Inc.'s Moab and TORQUE are used on every sizable Cray X-series machine in the field. The first port of Moab and TORQUE products for the Cray X-series was performed at the OLCF, and OLCF staff have been involved in subsequent design enhancements. In 2012 OLCF staff collaborated with Adaptive Computing to develop a design that would alleviate synchronization issues between the native Cray application-level placement scheduler (or ALPS) and the Moab batch scheduler. The new design also provides the benefit of allowing the Moab and TORQUE servers to easily exist outside the Cray machine on external service nodes, which enables users to submit and manipulate jobs when the Cray compute partition is unavailable. The OLCF hosted the beta testing of this new design at scale on Jaguar and provided numerous patches to Adaptive Computing for the redesigned code. This redesign has now been incorporated in Adaptive Computing's product line.

4.4 Innovations in Streaming Data Management: Technology Export to the Spallation Neutron Source

In addition to being home to the world's fastest supercomputer, ORNL also operates the world's brightest neutron source, the Spallation Neutron Source (SNS). Funded by the DOE Office of Basic Energy Science (BES), this national user facility hosts hundreds of scientists from around the world, providing a platform to enable breakthrough research in materials science, sustainable energy, and basic science. Because of their expertise in large-scale data management, OLCF personnel have been engaged to help manage and analyze the large data sets (over 1 TB per data set) generated by the intense pulses of neutrons.

OLCF collaborated with the SNS staff to successfully complete the Accelerating Data Acquisition, Reduction, and Analysis (ADARA) Lab-Directed Research and Development (LDRD) project. As a result of the ADARA project, a new data infrastructure was created that enhances users' ability to collect, reduce, and analyze data as it is taken; create data files immediately after acquisition, regardless of size; reduce a data set in seconds after acquisition; and provide the resources for any user to do post-acquisition reduction, analysis, visualization, and modeling without requiring users to be on-site at the SNS facility. On August 29, 2012, during a formal BES review of the SNS, reviewers responded very positively to a live demonstration of the ADARA system on HYSPEC (a beam line at SNS).

ADARA provides a streaming data backplane, allowing scientists to go from experiment to data reduction to obtaining an energy spectrum or diffraction pattern nearly instantaneously and while the experiment is still running. Rather than the previous approach of saving data in "buckets" and, once the bucket is full, handing the bucket off to the next process, ADARA uses a streaming approach. As data is being captured, translation is done concurrently. Every single event coming off a detector is translated to a common data format as the experiment progresses. While performing translation, ADARA also does live data reduction, so as neutron events are coming off the detectors, that same data is reduced live into an energy spectrum or diffraction pattern. To accomplish this, the ADARA architecture leverages a combination of many techniques commonly used in HPC with other techniques from

traditional distributed computing such as publish/subscribe. Moving forward, ADARA will be deployed in production on several other beam lines at SNS.

4.5 Innovations in the Lustre Parallel File System

4.5.1 Shaping the Lustre Community through Leadership in OpenSFS

Adoption of the Lustre parallel file system continues to expand in the wider HPC community. Many Lustre code improvements, whether bug fixes, maintenance, or feature requests are being driven through funding by the OpenSFS organization, which the OLCF co-founded. OLCF personnel lead the Open SFS Technical Working Group and contribute to the Release Working Group. These working groups are primarily responsible for determining the future direction for Lustre development and ensuring that contracted work is accomplished. OLCF staff members have overseen the development of features that significantly benefit the OLCF and the broader Lustre community. In 2012 these enhanced features included a metadata server (MDS) survey tool, imperative recovery, and general metadata improvements.

One of the greatest challenges to applications on the HPC systems is MDS performance. An MDS survey tool was developed through a contract with Whamcloud, Inc. to simulate standard workloads, permitting performance testing without requiring any clients to create metadata operations. Several metadata handling defects were identified and addressed during 2012. One such defect was the single lock on a single directory. This lock would serialize access to a directory if several processes simultaneously attempted to access this location in the file system. The solution was to introduce more granular locking on ext4 directory structures, thus allowing more concurrent operations.

Today's servers are normally multicore, but Lustre had not taken advantage of these resources. To address this issue, another 2012 performance improvement was to add a new symmetric multiprocessing (SMP) layer to minimize the penalties that occur when processes migrate from one core to another.

OLCF Lustre Development Efforts and Its Impact on Center Operations

In addition to the work leveraged by OpenSFS, internal work at the OLCF has also helped enhance Lustre. Titan uses the base Linux distribution SLES11 SP1, and the OLCF plans to support Service Pack 2 (SP2) in the future. However, Whamcloud supports only Red Hat-based servers. The OLCF has helped bridge the gap for support in this unique environment. This work will also benefit the wider Lustre community and prepare for newer kernels supported in the SuSE-based system. In the future, Red Hat-based systems will also use these newer kernels. This work has been successfully completed and will be a part of the Lustre 2.4 release.

The Gemini interconnect is a new networking technology used in the Cray XE and XK systems. While the underlying Gemini technology was supported on Lustre 1.8, newer Lustre versions do not support this technology, limiting the OLCF from taking advantage of its new features and performance improvements. To overcome this limitation, OLCF staff have ported and improved the Gemini Lustre Networking (LNET) driver for Lustre. The base work has been completed and will be merged in Lustre 2.4 support. Experimental work is being done on SMP scaling to study what configurations could be used to improve the performance of the Gemini LNET driver. The goal is to have that work also merged into the Lustre 2.4 version.

Monitoring and reporting of disk usage on Lustre file systems is a resource-intensive task and can affect metadata performance if not done in a centralized and scalable way. LustreDU is a nonintrusive tool developed by OLCF personnel to address this issue. It provides an end-user utility that queries a database of file- and directory-size information. This database is updated daily from a separate process that runs on the Lustre servers. This approach is significantly more efficient than going through the Lustre client-side application programming interface (API) that the normal “du” utility would have to use. This tool has had a significant positive impact on OLCF operations as system administrators use it to trim the scratch space usage of the Spider Lustre parallel file system.

4.6 Innovations in the HPSS Archival Storage

The HPSS is the software used by the OLCF for managing more than 30 petabytes of data on disk and robotic tape. HPSS has a large installed base outside of the OLCF. It is developed and maintained through an industry/national laboratory collaboration. It provides highly flexible and scalable hierarchical storage management that keeps recently used data on disk and less recently used data on tape.

For the last several years, OLCF staff members have had the primary responsibility for the development of the following HPSS subsystems:

- Storage System Manager (SSM), the graphical and command-line interface for monitoring, configuring, and controlling the system
- Bitfile Server (BFS), one-third of the Core Server
- Logging subsystem
- Accounting subsystem

During 2012 the OLCF's efforts were devoted to HPSS release 7.4, HPSS release 7.p, and quality improvements, as discussed below.

Release 7.4: The capabilities implemented in release 7.4 had significant contributions from OLCF HPSS developers.

- **Dynamic drive updates.** This update builds upon the dynamic drive's add-and-delete functionality, which was first provided in HPSS 7.3 and allows device configuration updates without system downtime. OLCF HPSS developers made major contributions to this feature in the area of SSM.
- **“hpssadm” enhancements.** “hpssadm” is the command-line interface to SSM. In release 7.4 it was extended to provide complete HPSS configuration capability. Lengthy system configuration changes can now be automated in a batch script, reducing or eliminating downtime. A complete system can now be configured from a script, enabling quick setup of new test systems or of production systems at new sites. The OLCF SSM developers were responsible for this feature in its entirety.
- **Logging enhancements.** Log files were changed from binary to text format, which is a tremendous boon to real-time debugging. Log archiving was improved to be more flexible and avoid potential loss of logging data during times of high activity; previous systems could lose some log data when a log file could not be archived quickly enough. The OLCF logging subsystem developer provided all of these improvements.

- **Redundant array of inexpensive tape (RAIT).** RAIT will improve both bandwidth to tape drives and reliability of the data the OLCF stores and retrieves in production. OLCF developers made small but essential contributions to the support of RAIT in the areas of SSM and logging.

Final testing of the 7.4 release was conducted in 2012, and release general availability is targeted for March 2013.

Release 7.p: The first step in the planning for moving HPSS to exascale is release 7.p, a performance-centered release. The strategy is to partition the metadata database across multiple storage devices and adapt HPSS to exploit the parallelism that this partitioning makes available. Work on release 7.p began in mid-2012 and is ongoing.

Quality: A major focus of HPSS development for 2012 was on improving quality. By redesigning the collaboration Software Development Process (SDP), the HPSS Technical Committee is addressing the challenge of maintaining software quality and agility in the face of declining resources. The OLCF developers have taken a major leadership role in the redesign of the SDP. In addition, in the past year the OLCF developers have significantly expanded, improved, and automated the Logging and BFS test suites—and are in the process of doing the same to the SSM test suite—to ensure greater reliability and higher quality in each release.

4.7 Innovations in Networking: The Common Communication Interface

The OLCF imposes scalability issues for everything from storage to debugging tools. In addition to Titan, the OLCF includes many different types of hardware as well as multiple types of network infrastructures. Each network provides at least two APIs: Berkeley sockets BSD and the network's native interface, which provides better performance through direct access to the network hardware. Titan, for example, uses Cray's General Network Interface (GNI) over Gemini, while the storage system uses Verbs over InfiniBand.

For each new generation of hardware, various groups within the OLCF port (i.e., modify) applications to use each network's native API to obtain the best performance (i.e., lowest latency, highest throughput, and lowest CPU utilization).

OLCF staff members have been working on a new programming interface that will provide a common API for applications, allowing them to take advantage of current networking hardware and next-generation hardware as it is acquired. This new API, known as the Common Communication Interface (CCI), is being jointly developed by ORNL, the University of Tennessee, Myricom, and Cisco. CCI is designed for portability, scalability, and performance. We continue to refine the API and have completed support for Sockets (UDP and TCP), Cray Portals and GNI, and Verbs. The software will be ready for adoption in 2013.

4.8 Innovations in Application Support: Hierarchical Collective Communications Library to Improve Application Performance

Collective operations are used in parallel computing to synchronize processes and for other operations such as broadcasts and reductions in which all processes participate. Collective operations are often among the most sensitive parts of scientific simulations with regard to performance and scalability. Unfortunately, collective operations that have been designed for supercomputers with single- or dual-processor nodes and a particular network fabric can present scalability and performance bottlenecks on today's multicore compute nodes and when used on different network fabrics. The OLCF's Computer Science Research Group (formerly

Application Performance Tools Group) built Cheetah,^{***} a framework for implementing collective operations to address these problems, from the ground up.

The goal of Cheetah is to provide an efficient collective-operations implementation for modern supercomputers and various programming models, including message passing and global address space models. Cheetah achieves this goal by designing the collective operation as a combination of simple collective primitives, which are each optimized for a different homogeneous block in the heterogeneous architecture. In addition, by enabling asynchronous progress, computation progress in parallel with communication and allow communication offload to network hardware so computation can be overlapped with the communication. The reference implementation outperforms the native implementations on Cray XE/XK and InfiniBand systems, and scales to over 100,000 cores on Cray XE/XK systems.

Cheetah is a joint effort of ORNL and Mellanox Technologies, and the Cheetah research shaped the features and capabilities of Mellanox's CORE-Direct technology. Cheetah 1.0.0, the first version to be officially released, will appear in forthcoming releases of Open MPI, a widely used open-source implementation of the MPI-2 standard, and is already available to Open MPI developers.

4.9 Innovations in Application Support: On-Site Support for the DDT Debugger and HMPP Compiler

The OLCF has long-standing partnerships with Allinea Software, Ltd., developers of the distributed debugging tool (DDT), and CAPS Enterprise, developer of the HMPP compiler suite for GPUs. We have implemented subcontracts with Allinea and CAPS to enhance and extend their products to meet the scale and functionality needs of OLCF users. In 2012 we expanded these partnerships to include the full-time placement of partner employees on-site at ORNL to provide in-depth support for these tools to OLCF staff and users.

On-site CAPS and Allinea staff actively work with application teams to address problems that go beyond typical tool support, often assisting in use of the tools in challenging situations. For example, Allinea on-site staff member Dirk Schubert recently helped an application team use DDT to isolate and fix a bug that occurred in its application only when it was run at full scale on Titan. When coupled with a long-term partnership such as OLCF's with Allinea and CAPS, placing support staff on location at the center is a best practice that enhances the partnership, leads to improvements in the partner's products, and facilitates center operations.

4.10 Innovations in Management Knowledge Transfer

The OLCF is able to provide world-class resources in support of science because of its depth of expertise and many years of experience in managing projects to successful, on-time, and on-target delivery. In 2012 the OLCF initiated training workshops to share its expertise related to ensuring the quality of project outcomes with project and program managers at ORNL. Kathlyn Boudwin, deputy project director for the NCCS/OLCF, led the training. The training was divided into two separate class offerings: "Project Management Methodology and Tools" and "Financial Tools for Project Management," each of which was offered multiple times. Eighty-seven people took at least one of the classes. Feedback from the class attendees (Computing and Computational Sciences Directorate project managers,

^{***} <http://www.csm.ornl.gov/cheetah/index.html>

principal investigators, project management assistants, and finance officers) was very positive. Attendees completed the workshops with a better understanding of how they could apply sound Project Management fundamentals to their programs.

Risk Management

HIGH PERFORMANCE COMPUTING FACILITY
2012 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2013

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 and rating, and retiring both project- and operations-based risks. Our risk management approach uses the Project Management Institute's best practices as a model. Risks are tracked and, when appropriate, retired, recharacterized, or mitigated. The major risks currently being tracked are listed and described below. Any mitigations planned for or implemented are included in the descriptions. The OLCF has only two "high" operational risks: (1) that the funding picture for 2013 is uncertain and (2) that the center will continue having difficulty finding sufficient staff. To address the first risk, the OLCF will continue to work with its DOE sponsors to understand their projections and adjust its plans accordingly. To address the second, recruiting efforts have received increased emphasis.

5.1 Risk Management Summary

The OLCF's Risk Management Plan (RMP) describes a regular, rigorous, proactive, and highly successful review process first implemented in October 2006. The RMP is reviewed at least annually and updated when 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-register database application 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 these 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 dates associated with the risk. Trigger conditions are stated in the Risk Notes narrative section of the registry 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 their rating level.

The OLCF reports current high- and medium-level risks to the DOE program office as part of its monthly operations report. At the time of this writing, 25 active entries are in the OLCF operations risk register. They fall into two general categories: risks for the entire facility and risks particular to some aspect of it. Across-the-board risks are concerned with such things as safety, funding/expenses,

and staffing. More focused risks are concerned with reliability, availability, and use of the system or its components (e.g., the computing platforms, power and cooling, storage, networks, software, and user interaction). Sixteen risks are tracked for the two projects that are current at the time of this report. A project risk may be listed below if it could also significantly impact operations.

Costs for handling risks are integrated within the budgeting exercises for the entire facility. Risk mitigation costs are estimated as any other effort cost or expense would be. For projects, a more formal bottom-up cost analysis is performed on the Work Breakdown Structure. However, for operations, costs of accepted and residual risks are estimated by expert opinion and accommodated as much as possible in management reserves. This reserve is continually reevaluated throughout the year.

5.2 Current Risk Status

The scope of operations risks remains relatively stable from year to year. Adequate funding is always a concern, and OLCF's mission of continual innovation requires both scientists and OLCF staff to make frequent adjustments to accommodate new technologies. Recently OLCF has experienced some staffing difficulties because of a highly competitive job market.

5.3 Major Risks Tracked in the Current Year (2012)

Risk	Section	Rating	Notes
979 – Insufficient funding to meet DOE commitments (FY2013)	5.3.1	High	Uncertainty is a concern. Annual budgets are set with guidance from the ASCR office, but actual allocated funds are unknown until Congress passes funding bills. Continuing resolutions are common, and we often go several months before actual funding is resolved.
1006 – Inability to acquire sufficient staff	5.3.2	High	Recently added. OLCF is having difficulty acquiring adequate qualified staff because of a highly competitive job market.
361 – Scientists decline to port to heterogeneous architecture	5.3.3	High then Medium	Porting is difficult. Mitigation includes in-house experience and training development. Remains a concern.
917 – Robust support will not be available to ensure portability of restructured applications	5.3.4	Medium	Remains a concern.
906 – Programming environment tools may be insufficient	5.3.5	High then Medium	Mitigation includes subcontracts with tool vendors. Remains a concern.
948 – Lack of infrastructure for an exascale system	5.3.6	Medium	Long lead time will be required to resolve. Remains a concern.
721 – Lustre metadata performance continues to impact applications	5.3.7	Medium	Mitigation includes participation in OpenSFS. Remains a concern.
412 – Inadequate system availability	5.3.9	Medium	Unforeseen problems may arise with new technology.
992 – Leadership computing is not achieved (CY2012)	5.3.10	High	Retired: Leadership targets achieved.
974 – Insufficient funding to meet DOE commitments (FY2012)	5.3.11	High then Medium	Retired: Funding was sufficient for FY2012.
973/975 – Supply chain issues (i.e., flooding in Thailand) may impact planned hard disk acquisitions	5.3.12	High/ Medium	Retired: Risk occurred but with little impact.

5.3.1 ID# 979 – Insufficient Funding to Meet DOE Commitments (FY2013)

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

Annual budgets are set with guidance from the ASCR office, but allocated funds are unknown until Congress passes funding bills. Continuing resolutions are common, and we often go several months before actual funding is resolved. The risk is that we may have to delay some purchases, activities, hiring, etc., or adjust lease payment schedules, resulting in high costs or schedule delays.

Trigger: Intelligence on congressional or DOE funding capabilities and priorities

The FY2013 budget is now being formulated, and substantial reductions from current projections are being discussed. We will maintain close contact with the Federal Project Director and ASCR Program Office to understand the changing funding projections so alternate plans can be made in sufficient time. Where possible we will structure contracts to accommodate flexible payment terms.

5.3.2 ID# 1006 – Inability to Acquire Sufficient Staff

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

The OLCF is having difficulty 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 OLCF lines

Although the cost, schedule, and technical impact ratings are all low, the risk is rated High because of “Other” impacts, such as those to OLCF’s or ORNL’s reputation as a preferred place to work.

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

5.3.3 ID# 361 – Scientists Decline to Port to Heterogeneous Architecture

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

Common to all programming models is the need to structure and/or restructure codes to express increased hierarchical parallelism on today's hybrid multicore architectures. This is necessary on all high-performance architectures to achieve good performance. Beyond this restructuring one 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 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.

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 Tools team is leveraging LDRD and other investments to develop tools to assist users in porting their codes. Recent evidence supports mitigation of this risk. The 2013 INCITE Call for Proposals realized an increase in proposal submissions requesting OLCF resources compared to the last three years. Of the 32 proposals awarded INCITE projects at OLCF during 2013, 21 had a computational readiness score of greater or equal to 4 out of 5. There appears to be many applications teams who are porting their codes.

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

This risk and risk ID# 906 (still active) and ID# 912 (now retired) are related to the introduction of the new heterogeneous architecture by the OLCF-3 project.

The original risk evaluation rated this risk as High. Mitigation with outreach, training, and the availability of libraries and development tools will ameliorate some user resistance. Discretionary resources are allocated for the purpose of porting, tuning, and scaling applications. Current trends in publication venues imply that many development teams are exploring architectures with accelerators, which is contrary to this risk.

5.3.4 ID# 917 – Robust Support Will Not Be Available to Ensure Portability of Restructured Applications

Risk Owner	Bronson Messer, Acting Group Leader, Scientific Computing		
Probability	Medium		
Impact	Cost: Medium	Schedule: Low	Scope/Tech: Medium
Rating	Medium		
Status	Mitigating the risk		

The programming model that we propose 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. Our experience to date shows that we almost always enhance the performance with this kind of restructuring. The use of directives-based methods will allow the lowest level of concurrency to concomitantly be exposed (e.g., vector- or streaming-level programming). This means that the bottom level of concurrency can be directly generated by a compiler. We expect 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)—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, OpenCL could be lacking on OLCF-3’s platform. Applications run on Titan could be developmental “dead ends,” and the improvements to application codes derived from porting might not survive.

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 approach will be abetted by publishing the case studies and experience with the six project applications, coupled with the appropriate training of our user community. Our work with vendors continues to improve compiler technology and other tools. Additionally, we have worked with compiler vendors to help form and promote OpenACC as a new standard aimed at providing a portable way to program for accelerator-based systems that is transparent to nonaccelerator systems.

The risk was modified in late CY2012 to better describe the risk.

5.3.5 ID# 906 – Programming Environment Tools May Be Insufficient

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

The OLCF-3 system (Titan) relies on GPU accelerators for the bulk of its computational capability. The programming environment for OLCF-3 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 original risk evaluation rated this risk as High. To mitigate the risk we created a Software Tools Group within the NCCS to be responsible for addressing the problem. We surveyed users on their requirements in this area and the adequacy of the tools available or planned. We found that plans were in place to extend the useful functionality of most of the primary tools from the OLCF-2 environment to the OLCF-3 system. Where we found gaps, we initiated contracts with vendors to accelerate their development and add key functionality needed for the OLCF-3 system. These activities were moved into the OLCF-3 project to provide initial risk mitigation. We monitor the progress of the tool developers and check out early versions of the tools on new Fermi processors in Jaguar and on other GPU-enabled systems to ensure their compatibility with existing programming models. We are also developing portable programming models (through our vendor partners), such as the directives-based OpenACC standard and the OpenMP directives for accelerators.

5.3.6 ID# 948 – Lack of Infrastructure for an Exascale System

Risk Owner	James H. Rogers, NCCS Director of Operations		
Probability	Low		
Impact	Cost: High	Schedule: High	Scope/Tech: Medium
Rating	Medium		
Status	Accepting the risk		

DOE's long-term plans include pre-exascale and exascale systems before the end of this decade. ORNL has a plan to provide the space, power, and cooling to support these goals, but there is a risk that the systems will be significantly larger or use more power than projected.

Trigger: Intelligence on the size and power requirements of proposed systems

The much-preferred approach would be to construct a new building designed specifically for exascale. The Office of Management and Budget has rejected third-party financing as a method of building such a facility, so such funding would require a congressional line item, therefore, ORNL has a plan to house the exascale system in Bldg. 5600 by moving other systems out of the building.

5.3.7 ID# 721 – Lustre Metadata Performance Continues to Impact Applications

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

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

Trigger: Direct observations reported by users or staff

The OLCF is working with other major Lustre stakeholders through OpenSFS to develop features to improve single metadata server performance and follow-on support of multiple metadata servers for the Lustre file system. We will be deploying Lustre 2.4 along with our storage system upgrade in September 2013.

The OLCF is also working with application teams to reduce their metadata workloads through code restructuring and the use of middleware I/O libraries. Tools have been developed to monitor and respond to metadata performance slowdowns to minimize the impact to the overall user population. Multiple file systems have been deployed, reducing the load on the metadata server. Even with all of this mitigation effort, however, progress toward a solution has been slower than expected, and the risk remains a significant concern.

5.3.8 ID# 412 – Inadequate System Availability

Risk Owner	Kevin G. Thach, Group Leader, High-Performance Computing Operations		
Probability	Medium		
Impact	Cost: Low	Schedule: Low	Scope/Tech: Low
Rating	Medium		
Status	Accepting the risk		

System design, manufacturing flaws/deficiencies, or other unforeseen issues related to hardware may impact the availability and stability of systems, which are critical to users. There is a risk that the system stability and availability may not be sufficient to meet user needs or our DOE operational metrics. Projected FIT rates of the Kepler chip are worse than for CPUs, so there is a risk that the system may not be stable enough to meet these operational requirements. Missing performance targets will disappoint sponsors and users, which might have a lasting impact in terms of reduced system usage and future support for the OLCF.

Trigger: Intelligence on or direct observation of deficiencies

The risk is rated Medium because the “Other” impact rating is Medium, reflecting concerns about meeting INCITE and other user metrics.

The OLCF will continue existing policies that control availability such as minimizing maintenance downtimes, coordinating upgrades, and maximizing fault-tolerant hardware and software. We will measure availability and stability and use the results to detect trends in time to enable remedial action. We will work closely with NVIDIA and Cray to characterize failures and develop responses should FIT rates impact our operational requirements.

5.3.9 ID# 992 – Leadership Computing Is Not Achieved (CY2012)

Risk Owner	Arthur S. Bland, OLCF Project Director		
Probability	High		
Impact	Cost: High	Schedule: High	Scope/Tech: High
Rating	High		
Status	Retired. The risk did not occur in CY2012		

Many users prefer to run smaller-scale jobs to obtain quicker throughput. The risk is that too many application jobs that did not achieve “leadership” status may have been submitted.

Trigger: Periodic performance reports showing trend toward possible missed goals

OLCF’s Scientific Computing Group continued to improve application readiness. Job-queue policies with a high preference for leadership jobs were established. OLCF’s continued involvement with the INCITE proposal-selection process ensured that leadership-class projects received preference.

This is considered a Low risk for CY2013.

5.3.10 ID# 974 – Insufficient Funding to Meet DOE Commitments (FY2012)

Risk Owner Arthur S. Bland, OLCF Project Director
Probability Low
Impact *Cost:* Low *Schedule:* Low *Scope/Tech:* High
Rating Medium
Status Retired. Risk did not occur. Funding was adequate.

Annual budgets are set with guidance from the ASCR office, but actual allocated funds are unknown until Congress passes funding bills. Continuing resolutions are common, and we often go several months before actual funding is resolved. The risk is that we may have to delay some purchases, activities, hiring, etc., or adjust lease payment schedules, resulting in high costs or schedule delays.

As the year progressed, this risk rating was reduced until the risk was eventually retired. Funding was sufficient for FY2012. The OLCF maintained close contact with the Federal Project Director and ASCR Program Office to understand the changing funding projections so alternate plans could be made in sufficient time. Where possible we structured contracts to accommodate flexible payment terms.

5.3.11 ID# 973/975 – Supply Chain Issues (i.e., Flooding in Thailand) May Impact Disk Drive Acquisitions

Risk Owner Al Geist, NCCS Chief Technology Officer
Probability Medium
Impact *Cost:* Medium *Schedule:* Medium *Scope/Tech:* Medium
Rating Medium
Status Retired. Risk event occurred and delayed disk acquisition but with only minor impact on operations.

Planned advanced storage system deployment might have been delayed beyond the fourth quarter of FY2012 because of the continued impact of flooding in Thailand on the availability of hard disk drives.

The original risk (ID# 973) was rated as High. To mitigate the risk we maintained very close contact with suppliers to monitor the situation and continually evaluate price projections. As a result we were able to develop a purchase plan that we hoped to initiate in May 2012. We retired ID# 973, replacing it with ID# 975, which accepted the risk and had a rating of Medium. Had the situation not improved or indicated improvement prior to release of the storage system request for proposal (RFP) in May, we would have considered reducing the scope (less performance due to increased disk prices), continuing to hold the RFP until the market improved, and/or increasing the level of funding for the acquisition. Alternately we could have continued to use the existing Spider file system.

5.4 Risks That Occurred During the Current Year and the Effectiveness of Their Mitigations

Risk	Section	Rating	Notes
990 – Titan could be destabilized when Kepler upgrade installed	5.4.1	Low	Retired: A problem occurred when the Kepler GPUs were installed.
973/975 – Supply chain issues (i.e., flooding in Thailand) may impact planned hard disk acquisitions	5.3.12	High/Medium	Retired: Maintained close contact with Thailand suppliers and their price projections. As a result of this accepted risk, the OLCF developed a 2012 purchase plan that accommodated the

Risk	Section	Rating	Notes
			schedule impact. Risk ID# 973 was retired and replaced with new risk ID# 975, which was also retired at a later time.
721 – Lustre metadata performance continues to impact applications	5.3.7	Medium	We continue to work with other major Lustre stakeholders and with apps to restructure codes. Progress is slow, so risk remains a concern.
407 – Loss of key personnel	N/A	Low	The leader of the Technology Integration Group accepted another assignment at ORNL. It took several months to find a replacement, but other OLCF staff were able to maintain acceptable work progress.
909 – New DOE travel restrictions could negatively impact the OLCF User Training Program	N/A	Low	Retired: Restrictions prevented some potential attendees from attending planned classes. Web-based classes were developed.

5.4.1 ID# 990 – Titan Could Be Destabilized When Kepler Upgrade Installed

Risk Owner Al Geist, NCCS Chief Technology Officer
Probability Low
Impact *Cost:* Low *Schedule:* Low *Scope/Tech:* Low
Rating Low
Status Retired. Risk event occurred.

When the Kepler upgrade to Jaguar in the fall of 2012 was analyzed for risk, it was understood that problems could be introduced that would destabilize the system, resulting in lower availability. At the time, the risk to the OLCF-3 Project was rated as Low because of prior system experiences, preliminary test results, and confidence that we maintained sufficient contingency reserves. The new boards that were installed in February 2012 worked fine until the Kepler GPUs were added to them several months later. When acceptance tests were run in December, results showed that we passed both the functionality and performance tests but did not pass the stability tests. Mechanical problems with the boards did not surface until the GPUs were installed. All of the boards need to be reworked to fix the problem.

In response to this risk occurrence, we have planned the sequence of repair of the boards to minimize the period of disruption by rolling the downtime through a small number of cabinets at a time. This rework may continue through the third quarter of FY2013. In the meantime users may safely continue to run jobs on Titan if they do not access the GPUs. While this issue remains a significant concern, it is believed that our system performance goals for the year still have a good chance of being achieved.

5.5 Risks Retired during the Current Year

Risk	Section	Rating	Notes
973/975 – Supply chain issues (i.e., flooding in Thailand) may impact planned hard disk acquisitions	5.3.12	High	Risk did not occur. We maintained close contact with Thailand suppliers and their price projections. As a result of this accepted risk, OLCF developed a 2012 purchase plan that accommodated the schedule impact. Risk ID# 973 was retired and replaced with new risk ID# 975.
990 – Titan could be destabilized when Kepler upgrade installed	5.4.1	Low	A problem occurred when the Kepler GPUs were installed.
974 – Insufficient funding to meet DOE commitments (FY2012)	5.3.11	Medium	The risk did not occur. Funding was sufficient for FY2012.
992 – Leadership computing is not achieved (CY2012)	5.3.10	High then retired	Risk did not occur. Leadership targets were achieved.
357 – Lustre Client and NVIDIA driver may be incompatible (<i>OLCF-3 risk</i>)	N/A	Low	The risk did not occur. The technologies were compatible.
409 – Leadership computing is not achieved (CY2012)	N/A	Low	Risk did not occur. Continued improvement in application readiness by the OLCF Scientific Computing Group helped achieve targets, as did establishing job-queue policies with high preference for leadership jobs and OLCF's continued involvement with the INCITE proposal -selection process such that leadership-class projects received preference.
908 – Inadequate staff training	N/A	Low	Risk did not occur. We mitigated this risk through feedback and adaptation of our training programs to meet staff needs. After each training opportunity, staff members were asked to participate in a survey on the quality and effectiveness of the training. Adjustments to the training effort were made accordingly.
912 – Lack of hardware availability impacts porting to heterogeneous architecture	5.5.1	Low	Risk did not occur.
915 – Upgrade from XT5 to XK6 takes too long, causing users to seek other alternatives	5.5.2	Medium	Risk did not occur. The XK6 boards had been released to other system owners before the OLCF received them. Although few problems were noticed with these early releases, the risk was that the much larger scale and complexity of the OLCF system could have created problems that delayed completion of the acceptance tests, thus delaying user access. The risk was retired when the system was returned to operations within the scheduled time. As a fallback, we required Cray to keep the existing Seastar-based boards for a period of time to make sure that Gemini was working.
932 – Jaguar could be destabilized when XK6 upgrade installed	N/A	Low	Risk did not occur.

5.5.1 ID# 912 – Lack of Hardware Availability Impacts Porting to Heterogeneous Architecture

Risk Owner	Bronson Messer, Acting Group Leader, Scientific Computing		
Probability	Low		
Impact	Cost: Low	Schedule: Medium	Scope/Tech: Low
Rating	Low		
Status	Retired. Risk did not occur		

OLCF-3's architecture with both Opteron processors and GPUs gives the users the opportunity to port codes from Jaguar, Intrepid, or other traditional systems to run on just the Opteron, while continuing to work on using the GPUs. As pointed out at the July 2009 Lehman review of the project, we will continue to develop a strategy to allow applications to be ported to OLCF-3 and still have portability to more traditional architectures. The risk is that users will be slow to adopt this programming model, resulting in application performance on the OLCF-3 system that would be lower than what it could be.

This risk was restructured in early CY2012 to describe the risk. Its original title was "New Architecture Requires New Programming Model," and it was initially rated as High. That rating was reduced to Medium with mitigation, then Low, and then Retired.

The mitigation involved integrating 960 Fermi cards into Jaguar to allow staff, developers, and users to have access to a GPU-based system to begin early work on porting applications. While this was an operational risk instead of a project risk, it was important to work with users early to allow them to begin porting to the system so the machine would more quickly be judged as successful by delivering breakthrough science.

5.5.2 ID# 915 – Upgrade from XT5 to XK6 Takes Too Long, Causing Users to Seek Other Alternatives

Risk Owner	Ann E. Baker, Previous Group Leader, High-Performance Computing Operations		
Probability	Medium		
Impact	Cost: Low	Schedule: Low	Scope/Tech: Medium
Rating	Medium		
Status	Retired. Risk did not occur.		

The XK6 boards had been released to other system owners before the OLCF received them. Although few problems were noticed with these early releases, the risk was that the much larger scale and complexity of the OLCF system could have created or surfaced problems that delayed completion of the acceptance tests, thus delaying user access. The risk was retired when the system was returned to operations within the scheduled time.

As a fallback, we required Cray to keep the existing Seastar-based boards for a period of time to make sure that Gemini was working.

5.6 Major New or Recharacterized Risks Since Last Review

Risk	Section	Rating	Notes
979 – Insufficient funding to meet DOE commitments (FY2013)	5.3.1	High	Uncertainty is a concern. Annual budgets are set with guidance from the ASCR office, but actual allocated funds are unknown until Congress passes funding bills. Continuing resolutions are common, and we often go several months before actual funding is resolved.
1006 – Inability to acquire sufficient staff	5.3.2	High	Recently added. OLCF is having difficulty acquiring adequate qualified staff because of a highly competitive job market.
917 – Robust support will not be available to ensure portability of restructured applications	5.3.4	Medium	Remains a concern.
412 – Inadequate system availability	5.3.9	Medium	Unforeseen problems might occur with the new technology.

5.7 Major Risks for Next Year

Risk	Section	Rating	Notes
979 – Insufficient funding to meet DOE commitments (FY2013)	5.3.1	High	Uncertainty is a concern.
1006 – Inability to acquire sufficient staff	5.3.2	High	Recently added. OLCF is having difficulty acquiring adequate qualified staff because of a highly competitive job market.
948 – Lack of infrastructure for an exascale system	5.3.6	Medium	Long lead time will be required to resolve. Remains a concern.
721 – Lustre metadata performance continues to impact applications	5.3.7	Medium	Remains a concern.

Summary of the Proposed Metric Values

HIGH PERFORMANCE COMPUTING FACILITY
2012 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2013

6. SUMMARY OF THE PROPOSED METRIC VALUES

CHARGE QUESTION 6: 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 target values that reflect the expectations of the stakeholders in delivering a leadership-class HPC resource.

The OLCF provides (below) a summary table of the metrics and actuals for 2012, and proposed metrics and targets for 2013 and 2014.

2012 Metric and Target	2012 Actual	2013 Metric	2013 Target	2014 Target	Reporting Period
<i>Are the processes for supporting the customers, resolving problems, and Outreach effective?</i>					
<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.	The OLCF exceeded the metric target: 4.2/5.0.	Overall score on the OLCF user survey.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	Annual
Improvement on results that scored below satisfactory in the previous period. Target: Results will show improvement in at least ½ of questions that scored below satisfactory (3.5) in the previous period.	The OLCF exceeded the metric target: No question scored below satisfactory (3.5/5.0) on the 2012 survey.	Improvement on results that scored below satisfactory in the previous period.	Results will show improvement in at least one-half of the questions that scored below satisfactory (3.5) in the previous period.	Results will show improvement in at least one-half of the questions that scored below satisfactory (3.5) in the previous period.	Annual
<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.	The OLCF exceeded the metric target: 4.5/5.0.	OLCF survey results related to problem resolution.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	Annual
OLCF user problem resolution time period. Target: 80% of OLCF user problems will be addressed within three business days, by either resolving the problem or informing the user how the problem will be resolved.	The OLCF exceeded the metric target: 92.3%.	OLCF user problem resolution time period.	Eighty percent of OLCF user problems will be addressed within three business days, by either resolving the problem or informing the user how the problem will be resolved.	Eighty percent of OLCF user problems will be addressed within three business days, by either resolving the problem or informing the user how the problem will be resolved.	Monthly
<i>Customer Metric 3: User Support</i>					
OLCF survey results related to Overall User Assistance and Outreach. Target: Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	The OLCF exceeded the metric target: 4.4/5.0.	OLCF survey results related to Overall User Assistance and Outreach.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	Annual

2012 Metric and Target	2012 Actual	2013 Metric	2013 Target	2014 Target	Reporting Period
<i>Is the facility maximizing the use of its HPC systems and other resources consistent with its mission?</i>					
<i>Business Metric 1: System Availability (for a period of one year following a major system upgrade, the targeted scheduled availability is 85% and overall availability is 80%)</i>					
Scheduled Availability. Target: Jaguar: 85% (lower in FY12 due to the compute system upgrades); HPSS: 95%; External File Systems: 95%.	The OLCF exceeded the metric target. Jaguar: 98.11%; HPSS: 98.46%; Widow1: 99.88%; Widow2: 99.81%; Widow3: 98.95%.	Scheduled availability.	85% (lower in FY12 due to the compute system upgrades).	90%	Monthly
Overall Availability. Target: Jaguar: 80%; HPSS 90%; External File Systems 90%.	The OLCF exceeded the metric target: 91.45% (Jaguar); 98.46% (HPSS); Widow1: 98.25%; Widow2: 98.69%; Widow3: 98.95%.	Overall availability.	Titan: 80%; HPSS 90%; External File Systems: existing, 90%, new, 80%	Titan: 85%; HPSS 90%; External File Systems: existing, 90%	Monthly
<i>Business Metric 2: Capability Usage</i>					
OLCF will report on capability usage. Target: At least 30% of the consumed node hours will be from jobs requesting 20% or more of the available Opteron nodes.	The capability usage was 50.67%. The OLCF exceeded the metric target.	OLCF will report on capability usage.	At least 30% of the consumed node hours will be from jobs requesting 20% or more of the available compute nodes.	At least 35% of the consumed node hours will be from jobs requesting 20% or more of the available compute nodes.	Monthly
N/A	N/A	OLCF will report GPU usage (metric only, no target).	N/A	N/A	Monthly

N/A = Not applicable.