U.S. Department of Energy, Office of Science

High Performance Computing Facility Operational Assessment, CY 2011 Oak Ridge Leadership Computing Facility

February 2012

Prepared by

**Arthur S. Bland
James J. Hack
Ann E. Baker
Ashley D. Barker
Kathlyn J. Boudwin
Doug Hudson
Ricky A. Kendall
Bronson Messer
James H. Rogers
Galen M. Shipman
Jack C. Wells
Julia C. White**

U.S. Department of Energy, Office of Science

**HIGH PERFORMANCE COMPUTING FACILITY
OPERATIONAL ASSESSMENT, FY11 OAK RIDGE
LEADERSHIP COMPUTING FACILITY**

Arthur S. Bland	Ricky A. Kendall
James J. Hack	Bronson Messer
Ann E. Baker	James H. Rogers
Ashley D. Barker	Galen M. Shipman
Kathlyn J. Boudwin	Jack C. Wells
Doug Hudson	Julia C. White

February 2012

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, Tennessee 37831-6283
managed by
UT-BATTELLE, LLC
for the
U.S. 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	v
LIST OF TABLES	vi
EXECUTIVE SUMMARY	1
Communications with Key Stakeholders	2
Summary of 2011 Metrics	3
Responses to Recommendations from the Previous 2011 Operational Assessment Review	6
1. USER RESULTS	12
1.1 User Results Summary	12
1.2 User Support Metrics	13
1.2.1 Overall Satisfaction Rating for the Facility	13
1.2.2 Average Rating across All User Support Questions	14
1.2.3 Improvement on Past Year Unsatisfactory Ratings	15
1.3 Assessing the Effectiveness of the OLCF User Survey	15
1.4 Problem Resolution Metrics	17
1.4.1 Problem Resolution Metric Summary	17
1.5 User Support and Outreach	18
1.5.1 User Support	18
1.5.2 UAO Analysts	19
1.5.3 Scientific Computing Liaisons	20
1.5.4 Communications	23
1.5.5 Training	23
1.5.6 Outreach	25
1.6 User Support Conclusion	26
2. BUSINESS RESULTS	27
2.1 Business Results Summary	27
2.2 Cray XT Compute Partition Summary	27
2.3 Resource Availability	30
2.3.1 Scheduled Availability	30
2.3.2 Overall Availability	31
2.3.3 Mean time to Interrupt	32
2.3.4 Mean Time to Failure	32
2.4 Resource Utilization	33
2.4.1 Total System Utilization	33
2.5 Capability Utilization	35
3. STRATEGIC RESULTS	37
3.1 Science Output	38
3.2 Scientific Accomplishments	38
3.2.1 Subatomic Physics	38
3.2.2 Chemical and Materials Research	39
3.2.3 Bioenergy	42

3.3	Accomplishments in Energy Systems Research	43
3.3.1	Energy Systems Simulation	43
3.4	Allocation of Facility Director's Reserve	47
3.4.1	Director's Discretionary Program	47
3.4.2	Industrial HPC Partnerships Program	49
3.5	Management of INCITE Projects	50
4.	INNOVATION	60
4.1	Application Readiness	61
4.2	Application Support	63
4.3	Outreach	64
4.4	Systems	65
4.5	Leadership	67
4.6	Energy Management	68
5.	RISK MANAGEMENT	70
5.1	Risk Management	70
5.2	Major Risks Tracked in the Current Year	71
5.2.1	ID# 979 – Insufficient Funding to Meet DOE Commitments (FY2013)	71
5.2.2	ID# 974 – Insufficient Funding to Meet DOE Commitments (FY2012)	72
5.2.3	ID# 404 – Insufficient Funding to Meet DOE Commitments (FY2011)	72
5.2.4	ID# 948 – Lack of Infrastructure for an Exascale System	72
5.2.5	ID# 721 – Lustre Metadata Performance Continues to Impact Applications	73
5.2.6	ID# 361 – Scientists Decline to Port to Heterogeneous Architecture	73
5.2.7	ID# 906 – Programming Environment Tools May Be Insufficient	74
5.2.8	ID# 912 – New Architecture Requires New Programming Model	74
5.2.9	ID# 975 – Supply chain issues i.e. flooding in Thailand may impact disk drive acquisitions	75
5.2.10	ID# 915 – System Upgrade from XT5 to XK6 Takes Too Long Causing Users to Seek Alternative Computing Resources	75
5.3	Risk That Occurred during the current year and the effectiveness of their mitigations	75
5.4	Risks retired during the current year	76
5.5	Major New or re-characterized risks since last review	76
5.6	Major Risks for next year	77
6.	SUMMARY OF THE PROPOSED METRIC VALUES	78

LIST OF FIGURES

Figure	Page
Figure 1.1. Number of Helpdesk Tickets Issued per Month.....	17
Figure 1.2. Categorization of Helpdesk Tickets	18
Figure 1.3. Rendering of the Fukushima Reactor Building Spent Fuel Rod Pool.	22
Figure 2.1. Eliminating VRM failures increases system stability.....	32
Figure 2.2. 2011 XT5 Resource Utilization – Core Hours by Program.	34
Figure 2.3. Effective Scheduling Policy Enables Leadership-class Usage.	36
Figure 2.4. Capability Usage by Project Type.	36
Figure 3.1. Atomic-detailed model of plant components lignin and cellulose. The leadership-class molecular dynamics simulation investigated lignin precipitation on the cellulose fibrils, a process that poses a significant obstacle to economically-viable bioethanol production.	42
Figure 4.1. The Effect of Top Hats on CRU Efficiency.....	69

LIST OF TABLES

Table	Page
Table 1.1. 2011 User Result Metrics Summary	13
Table 1.2. OLCF User Support Summary: Metric Targets and Calendar Year Results	13
Table 1.3. Satisfaction Rates by Program Type for Key Indicators.....	14
Table 1.4. Overall Satisfaction Results	15
Table 1.5. User Survey Participation	16
Table 1.6. User Survey Responders by Program Type	16
Table 1.7. Statistical Analysis of Key Results.....	16
Table 1.8. Problem Resolution Metric Summary	17
Table 1.9. Training Event Summary.....	23
Table 2.1. OLCF Production Computer Systems, 2011.....	28
Table 2.2. OLCF HPC System Production Dates, 2007 – Present	28
Table 2.3. OLCF Business Results Summary for HPC Systems	29
Table 2.4. OLCF Business Results Summary for HPSS	29
Table 2.5. OLCF Business Results Summary for the External Lustre File Systems	29
Table 2.6. OLCF Business Results Summary: Scheduled Availability	30
Table 2.7. OLCF Business Results Summary: Overall Availability.....	31
Table 2.8. OLCF Business Results Summary: Mean Time to Interrupt	32
Table 2.9. OLCF Business Results Summary: Mean Time to Failure	33
Table 2.10. 2011 OLCF System Utilization	34
Table 2.11. OLCF Capability Utilization Definition	35
Table 2.12. OLCF Leadership Usage on Jaguar XT5 and XT4.....	35
Table 3.1. List of OLCF Publications	38
Table 3.2. Director’s Discretionary Program: Domain Allocation Distribution.....	48
Table 3.3. Director’s Discretionary Program: Awards and User Demographics.....	48

Table 3.4. Industry Projects at the OLCF.....	49
Table 4.1. The Positive Impact on CRU Return-air Temperatures with Top Hats	69

Executive Summary

HIGH PERFORMANCE COMPUTING FACILITY 2011 OPERATIONAL ASSESSMENT OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2012

EXECUTIVE SUMMARY

Oak Ridge National Laboratory's Leadership Computing Facility (OLCF) continues to deliver the most powerful resources in the U.S. 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) 2011 to researchers around the world for 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. Users reported more than 670 publications this year arising from their use of OLCF resources. Of these we report the 300 in this review that are consistent with guidance provided.

Scientific achievements by OLCF users cut across all range scales from atomic to molecular to large-scale structures. At the atomic scale, researchers discovered that the anomalously long half-life of Carbon-14 can be explained by calculating, for the first time, the very complex three-body interactions between all the neutrons and protons in the nucleus. At the molecular scale, researchers combined experimental results from LBL's light source and simulations on Jaguar to discover how DNA replication continues past a damaged site so a mutation can be repaired later. Other researchers combined experimental results from ORNL's Spallation Neutron Source and simulations on Jaguar to reveal the molecular structure of ligno-cellulosic material used in bioethanol production. This year, Jaguar has been used to do billion-cell CFD calculations to develop shock wave compression turbo machinery as a means to meet DOE goals for reducing carbon sequestration costs. General Electric used Jaguar to calculate the unsteady flow through turbo machinery to learn what efficiencies the traditional steady flow assumption is hiding from designers. Even a 1% improvement in turbine design can save the nation billions of gallons of fuel. Details of these and other scientific discoveries—not possible without access to leadership-class computing resources—are detailed in Section 3 of this report and in the INCITE in Review, available at http://science.energy.gov/~media/ascr/pdf/program-documents/docs/INCITE_IR.pdf.

Effective operations of the OLCF play a key role in the scientific missions and accomplishments of its users. This Operational Assessment Report (OAR) will delineate the policies, procedures, and innovations implemented by the OLCF to continue delivering a petaflop-scale resource for cutting-edge research. This report covers CY 2011 that unless otherwise specified, denotes January 1, 2011 through December 31, 2011. The new XK6/Gemini partition was released to users for the second half of December 2011.

2011 highlights of OLCF operational activities include the following:

- User Support remains a key element of OLCF operations, with the philosophy “whatever it takes” to enable successful research. Impact of this center-wide activity is reflected by the annual user survey results that show users are “very satisfied.” The OLCF continues to aggressively pursue outreach and training activities to promote awareness—and effective use—of U.S. leadership-class resources (Section 1).
- Despite the disruptions of upgrading the Cray XT5 to an XK6, the OLCF continues to exceed DOE metrics for capability usage (40% target in CY 2011, 54% delivered). The Schedule Availability (SA) and Overall Availability (OA) target metrics of 95% and 90%, respectively, for Jaguar were exceeded in CY 2011 (96.37% and 92.88%, respectively). (Section 2).
- The numerous science accomplishments are more fully described in Section 3 and reflect OLCF leadership in enabling high-impact science solutions and vision in creating an exascale-ready center. Center innovations to help deliver science are outlined in Section 4.
- Risk Management (Section 5) is carried out using best practices approved of by DOE. The Center continues to work closely with the DOE Program Manager to refine metrics by which its operational performance is assessed. The proposed metrics for future years are provided in Section 6.

Communications with Key Stakeholders

Communication with the Program Office

The OLCF communicates regularly with the Program Office through a series of established events. These include weekly IPT 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 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.

Communication with the User Community

The role of communications in everything the OLCF does cannot be overstated, whether it is communicating science results to the larger community or communicating tips to users on using OLCF systems more efficiently and effectively. The OLCF uses various avenues, both formal and informal, for communicating with users. Formal mechanisms include the following:

- OLCF support services;
- weekly messages to all users on events;
- monthly OLCF User Council calls;
- quarterly user conference calls;
- annual users meeting;
- workshops and training events; and
- web resources such as system status and update pages, project account summaries, online tutorials and workshop notes, and other documentation such as “frequently asked questions” (FAQs).

Communication with the Vendors

OLCF conducts formal quarterly reviews of projects and operations with Cray, Inc. This 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 key to driving design considerations that positively affect emerging products. Supplementing these formal events, OLCF meets weekly with their 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.

OLCF maintains a robust vendor briefing schedule with other product manufacturers as well, making sure that emerging products that are 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 User Council consists of researchers who have active accounts on the leadership computing facility compute resources. The council meets via a teleconference on a monthly basis. The current User Council is chaired by Balint Joo. The council has been very engaged and provided valuable input to OLCF management this past year in activities such as reviewing queuing policy changes and other operational policies affecting users, assisting in gathering survey results, and participating in outreach activities.

An integral part of the operations of the center is the allocation of time through the INCITE program, which the OLCF jointly manages with the Argonne Leadership Computing Facility. On June 1-2, 2011, the DOE conducted a panel review of the INCITE program. Conversations with the panel participants provided validation of the implementation of the INCITE program: the LCFs and INCITE manager were commended for their oversight of the program.

Although project-focused by nature, the Lehman Reviews also provide an opportunity for communication with external reviewers who advise the center on many matters that cross-cut operations.

Summary of 2011 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 calendar year 2011. The metrics are associated with a series of questions posed to reviewers of the center. The 2011 metrics, target values, projections (through August 2011) and actual results as of December 31, 2011 are summarized below.

Summary of the 2011 Metrics

2011 Metric	2011 Target	2011 Projection	2011 Actual
<i>Are the processes for supporting the customers, resolving problems, and communicating with key stakeholders and Outreach effective?</i>			
Customer Metric 1: Customer Satisfaction	Overall OLCF score on the user survey will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	OLCF exceeded the target in 2010 with a survey mean score for overall customer satisfaction of 4.2, "very satisfied."	OLCF exceeded the target in 2011 with a survey mean score for overall customer satisfaction of 4.2, "very satisfied."

Summary of the 2011 Metrics

2011 Metric	2011 Target	2011 Projection	2011 Actual
	Annual user survey results will show improvement in at least ½ of questions that scored below satisfactory (3.5) in previous period.	OLCF received one question in 2010 that scored below 3.5. OLCF projects to meet or exceed the target in 2011.	No user responses in the 2011 period were below the 3.5 satisfaction level. OLCF met this target.
Customer Metric 2: Problem Resolution	80% of OLCF user problems will be addressed within three working days, either resolving the problem or informing the user how the problem will be resolved.	Through June 30, 2011, 89.5% of queries were addressed within three working days.	Through December 31, 2011, 89.8% of queries were addressed within three working days. OLCF exceeded this target.
Customer Metric 3: User Support	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.	The 2011 survey solicits an overall user satisfaction rating and comments about support, services, and resources. The 2012 user survey incorporates additional specific questions about user support.
<i>Is the facility maximizing the use of its HPC systems and other resources consistent with its mission?</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%)	Scheduled Availability: 95%	Through June 30, 2011: Cray XT5 (93.9%); Cray XT4 (97.6%); HPSS (99.9%); Spider (98.5%); Spider2 (99.9%); Spider3 (99.9%). OLCF projects to meet or exceed each target in 2011.	Through December 31, 2011: Cray XT5 (96.4%); Cray XT4 (97.6%); HPSS (99.1%); Widow 1 (99.3%); Widow 2 (99.9%); Widow 3 (99.95%). OLCF exceeded each target in 2011.
	Overall Availability: 90%	Through June 30, 2011: Cray XT5 (88.7%); Cray XT4 (97.1%); HPSS (98.9%); Spider (96.5%); Spider2 (99.1%); Spider3 (99.2%). OLCF projects to meet or exceed each target in 2011.	Through December 31, 2011: Cray XT5 (92.9%); Cray XT4 (97.1%); HPSS (98.7%); Widow 1 (98.0%); Widow 2 (99.3%); Widow 3 (99.7%). OLCF exceeded each target in 2011.
Business Metric 2: Resource Utilization	OLCF will report on INCITE allocations and usage.	CY 2011 INCITE allocations of 930 million hours. Through June 30, 2011, INCITE usage was 375 million core-hours, or 40.3% of the total allocation.	CY 2011 INCITE allocations of 930 million hours. Through December 31, 2011, INCITE usage was 995,214,895 core-hours, or 107.0% of the total allocation.
Business Metric 3: Capability Usage	At least 40% of the consumed core hours will be from jobs requesting 20% or more of the available cores.	Through June 30, 2011, the capability usage was 54%. OLCF projects to meet or exceed each target in 2011.	Through December 31, 2011, the capability usage was 54%. OLCF exceeded the target in 2011.

Summary of the 2011 Metrics

2011 Metric	2011 Target	2011 Projection	2011 Actual
<i>Is the facility enabling scientific achievements consistent with the Department of Energy strategic goals 3.1 and/or 3.2?</i>			
Strategic Metric 1: Scientific Output	The OLCF will report numbers of publications resulting from work done in whole or part on the OLCF systems.	Through June 30, 2011, 181 publications were reported to the OLCF by users or identified by the OLCF.	Through December 31, 2011, 670 publications were reported to the OLCF by users or identified by the OLCF. 300 are reportable within OAR guidance.
Strategic Metric 2: Scientific Accomplishments	The OLCF will provide a written description of major accomplishments from the users over the previous year.	Reference Section 3	Reference Section 3
Strategic Metric 3: Allocation of Director's Discretionary time	The OLCF will report on how the Facility Director's Discretionary time was allocated.	Reference Section 3	Reference Section 3
<i>Are the costs for the upcoming year reasonable to achieve the needed performance?</i>			
Financial Performance	The OLCF will report on budget performance against the previous year's Budget Deep Dive projections.	This information is included in the OLCF Budget Deep Dive review documentation.	This information is included in the OLCF Budget Deep Dive review documentation.
<i>What innovations have been implemented that have improved the facility's operations?</i>			
Innovation Metric 1: Infusing Best Practices	The OLCF will report on new technologies that we have developed and best practices we have implemented and shared.	Reference Section 4	Reference Section 4
Innovation Metric 2: Technology Transfer	The OLCF will report on technologies we have developed that have been adopted by other centers or industry.	Reference Section 4	Reference Section 4
<i>Is the Facility effectively managing risk?</i>			
Risk Management	The OLCF will provide a description of major operational risks.	Reference Section 5	Reference Section 5
<i>Does the facility have a valid cyber security plan and authority to operate?</i>			
Cyber Security Plan	The OLCF will provide the date of approval and expiration of our authority to operate.	The OLCF maintains an approved cyber security plan. The current authority to operate was granted on June 21, 2011 and expires on June 20, 2012.	The OLCF maintains an approved cyber security plan. The current authority to operate was granted on June 21, 2011 and expires on June 20, 2012.

Responses to Recommendations from the Previous 2011 Operational Assessment Review

In August, 2011, the OLCF delivered a Facility Operational Report to the DOE program sponsor that documented the 2011 calendar year-to-date operational activities of the center. Recommendations provided by reviewers of that report and ORNL actions and DOE ASCR comments are given in the tables below.

1. Are the Facility responses to the recommendations from the previous year's OAR reasonable?

Recommendation	ORNL Action/Comments	HQ Action/Comments
We recommend that OLCF utilize existing external analyses (such as Lehman reviews) to normalize their own risk assessments, and ensure they are realistic. If they are unable to determine any high technology or schedule risks with a major upgrade, they should solicit additional external input or review.	The OLCF has a robust and effective risk management program that we have used since 2006 to assess both project and operations risks. The documentation and management of OLCF Project risks continue to be essential components of our overall strategy for delivering each system on or ahead of schedule and within the allocated budget. External reviews and assessments are taken very seriously and, where applicable, this input is used to adjust existing risk management strategies.	October 20, 2011: Explanation of risk assessment project explained. Response addresses recommendation

2. Are the processes for supporting the customers, resolving problems, and communicating with key stakeholders effective?

Recommendation	ORNL Action/Comments	HQ Action/Comments
OLCF should consider categorizing survey respondents automatically as INCITE, ALCC and/or Discretionary, rather than relying on self-reporting, which produces a large number of "Other" responses and inaccuracies.	The 2011 survey has been changed to require users to include their Project ID. The Project ID will then be mapped to the appropriate allocation category. 2/8/2012: The 2011 user survey was changed to require users to include their Project ID(s). The Project ID was used to map the user to the appropriate project. Some users did not supply accurate Project IDs so we could not account for all users. However, the amount of "other" responses decreased to 6% which was a great improvement over the previous year.	October 20, 2011: Survey modified. Recommendation addressed 10/4/2011: Consider adding additional categories to cover the "Other" responses or inform users of which category they fit into.
The use of the Delphi Technique for validation is interesting. It should be explained.	The OLCF will include a brief summary of the survey validation techniques in next year's OA report. 2/8/2012: Completed.	October 20, 2011: Acceptable solution. Recommendation addressed
OLCF should provide a comprehensive list of the publications it produces (for example the annual report, outreach brochures, media packets)	The OLCF will provide a comprehensive list of Outreach products in the next OA report. 2/8/2012: A list of Outreach products will be included in the next (e.g. 2012) OA report.	October 20, 2011: Acceptable solution. You may consider also putting on your website Recommendation addressed.

3. Is the OLCF maximizing resources consistent with its mission?

Recommendation	ORNL Action/Comments	HQ Action/Comments
Please document the conditions under which a system is considered "down." In particular, describe how partial system outages (i.e. m out	A system is considered "down" if ANY of the following criteria are met: - All batch nodes are down, - All login nodes are down, - The boot node is down, - The batch scheduling node is down, or	October 20, 2011: Explanation provided. Recommendation addressed

Recommendation	ORNL Action/Comments	HQ Action/Comments
of n nodes down) are accounted for in calculating availability.	- 5% or more of the compute nodes are down.	
Overall availability is ~5% less than scheduled availability implying that the Cray XT5 was in scheduled downtime for approximately 2.6 weeks out of the year (1.5 days per month). Additional efforts should be made to minimize scheduled downtime (not counting the planned downtime for the "Titan" upgrade.)	The total scheduled downtime in CY2010 reflected not only the anticipated downtime to perform preventative maintenance on the XT5, but also a series of reservations associated with Lustre performance and scaling testing as the file systems moved to production. OLCF made concerted efforts to consolidate these test reservations as much as possible to minimize downtime. At this point this testing is effectively complete and has contributed to a more stable production file system for the center.	October 20, 2011: Explanation provided. Recommendation addressed
The facility should closely monitor the VRM issue on the Cray XT5 to ensure that the issue does not reoccur.	The OLCF and Cray are closely monitoring this and all causes of down time. To date, there have not been any reoccurrences since implementation of the engineering solution. We are currently meeting both SA and OA metrics for the year.	October 20, 2011: Description of current monitoring actions provided. Recommendation addressed
Per the guidance, ORNL should report the Total System Utilization for the previous calendar year and a projection for the CY 2011 Total System Utilization. The report only included the first six months of the report calendar year. Also, please provide utilization results for the XT5	For CY2010, 1,435,391,110 core-hours were delivered from a scheduled maximum of 1,758,150,000 core-hours. This resulted in total system utilization for the Cray XT5 of 81.64%. For consistency, all metrics were reported for the OA report using the same time period that closed on June 30, 2011. Cray XT5 system utilization through June 30, 2011 was 744,861,807 of 866,290,875 core-hours (85.98%). System utilization through September 30, 2011 was 1,153,398,960 of 1,347,303,188 core-hours (85.61%). System utilization (projected) using a linear method for the period ending December 31, 2011, and not considering any upgrade to the Jaguar system produces a projected system utilization of 1,537,865,280 of 1,796,404,250 core-hours (85.61%). Given that Jaguar is scheduled for an upgrade in the last quarter of 2011, with inflection points due to rolling upgrades on October 10, November 4, and December 15, the total available hours can be projected (with a similar reduction in consumed hours) as 1,364,072,723 of 1,593,394,473 (85.61%). This projection does account for the increased core-size of the upgraded partition(s), but continues to use a linear assumption of the utilization for projecting core-hours consumed. That assumption of calculated utilization may be impacted negatively by a larger number of system interruptions, and impacted positively by greater demand for fewer core-hours.	October 20, 2011: Explanation provided. Recommendation addressed

Recommendation	ORNL Action/Comments	HQ Action/Comments
Please provide projections for the CY 2011 capability results.	For the period ending June 30, 2011, capability usage on the Cray XT5 was 53.6%, well ahead of the prescribed metric of 40%. This is primarily due to aggressive scheduling policies for larger jobs. For the period ending September 30, 2011, capability usage continued to track well ahead of the metric at 56.37%. A linear projection of the last quarter of CY2011 would produce 216,707,489 additional hours of capability-class computing (for the quarter), at this same 56.37% rate. Given that Jaguar will undergo rolling upgrades, there will be inflection points where what constitutes a capability-class job will be reduced. Given that demand will remain high, and that more jobs will qualify as capability-class jobs, the number of capability-class jobs will increase. However, because the largest jobs must be reduced in size, the total delivered core-hours will fall. Assuming that the combination of factors produce an equivalent capability rate (56.37%), we can project that 118,747,937 capability-class hours could be produced in the fourth quarter during the rolling upgrade.	October 20, 2011: Explanation provided. Recommendation addressed
Please describe how the various projections were determined.	System utilization is calculated linearly using the average through the first nine months of the calendar year for projecting the twelve-month total core-hours utilized and system utilization percentage. The capability result is calculated linearly using the average through the first nine months of the calendar year for projecting the twelve-month total core-hours and capability percentage.	October 20, 2011: Explanation provided . Recommendation addressed

4. Is the OLCF meeting the Department of Energy strategic goals 3.1 and 3.2?

Recommendation	ORNL Action/Comments	HQ Action/Comments
Please put the number of refereed publications in context, for example by saying how many projects or how many users produced those publications.	OLCF will explicitly include the number of project in future OA reports. 2/8/2012: Completed.	October 20, 2011: Future action proposed acceptable. Recommendation addressed
OLCF should consider tracking its publications for five years and provide the citations for these papers.	The OLCF will consider the recommendation for tracking publications from INCITE projects for a longer period of time and investigate the feasibility of providing the citations for these papers for the next OA report. 2/8/2012: The OLCF will track publications over an extended time period (e.g. 5 years) and will report initial results in the next (e.g. 2012) OA report.	October 20, 2011: Possible action identified and acceptable . Recommendation addressed

5. How well is the program executing to the cost baseline pre-established during the previous year's Budget Deep Dive? Explain major discrepancies.

Recommendation	ORNL Action/Comments	HQ Action/Comments
OLCF should address whether having a \$7.1M - \$9.4M carry-forward going into FY13 is adequate or advisable. BY comparison the FY11/12 carry forward is \$18.6M. If it needed to be increased, that may impact budget decisions in FY12.	For purposes of calculating the carry forward required for continuing operations under a CR, OLCF generally budgets for one month's routine expenses plus any amount needed to cover known commitments that will occur early in the fiscal year. The \$18.6M budgeted carry forward includes one month's routine operating expenses, plus funds to retire the Jaguar lease and to begin the DME site preparation procurements.	October 20, 2011: Explanation provided. Recommendation addressed
Three tables are provided (Tables 5.1, 5.2, and 5.3) to show planned/actual costs for FY 11 and planned cost FY12. OLCF should incorporate all of the information into a single table.	OLCF will provide the information in a single table in the future.	October 20, 2011: Proposed action acceptable. Recommendation addressed
The FY12 Target budget totals \$104.2M cost (\$8.7M avg./mo), and carryout to FY13 is \$9.4M. The carryout amount of \$9.4M seems to be a little light based on the average monthly cost and a possible CR in FY13. OLCF should explain how this is an adequate amount.	Using a simple average monthly cost does not reliably predict OLCF's ability to cover early in the FY expenses. OLCF uses a more detailed analysis of routine expenses which are fairly consistent month to month, but also includes non routine expenses which may include OCLF-3 DME project expenses or major procurements which are scheduled for a particular time period.	October 20, 2011: Explanation provided. Recommendation addressed
The FY12 Baseline budget totals \$99.5M costs (\$8.3M avg./mo), and carryout to FY13 is \$7.1M. The carryout amount of \$7.1M appears to be insufficient based on the average monthly cost and possible CR in FY13. OLCF should explain the adequacy of the carryover.	See explanations for carry forward amounts above. An explanation of carry forward will be included with the next OA report submission.	October 20, 2011: Explanation provided. Above and proposed action is acceptable. Recommendation addressed
It is noted that OLCF lost eight staff members during FY11. OLCF should provide a hiring plan for replacing these staff members. Also, a table showing current FTE's by group and planned additions would be helpful.	Several of the vacancies mentioned in the 2011 OA report have been filled and others have been posted. Any time a position is vacated, the OLCF group leader and the OLCF management assess the best option for replacing this human resource. If it is determined that the position needs to be filled, the position is posted using the ORNL HR department guidelines. The OLCF will include a table showing current FTEs by group and where additional staff members are planned, in the next OA report submission.	October 20, 2011: Explanation provided and proposed action is acceptable. Recommendation addressed

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

Recommendation	ORNL Action/Comments	HQ Action/Comments
Please provide best practices and innovations in the areas of facilities	In the last several years OLCF has included facilities modifications that we have made, and we will do so in the future when we make additional facility modifications. OLCF didn't make any modifications this year, therefore didn't have anything to report.	October 20, 2011: Explanation provided. Recommendation addressed

7. Is the OLCF effectively managing risk?

Recommendation	ORNL Action/Comments	HQ Action/Comments
OLCF did not talk about the risk of obtaining financing for OLCF-3 and future acquisition projects. Recommend addressing how OLCF is mitigating the possibility of not being able to obtain third party financing. Should be followed in the risk register.	The OLCF risk register includes the risk of not being able to obtain financing for the OLCF-3 system. This risk was mitigated by including an option for financing in the contract with Cray, and therefore is not considered a high risk. OLCF will require manufacturers to provide a lease option in future contract, or contract for third party leases in advance of the contract to acquire the computer.	October 20, 2011: Explanation provided and future actions identified. Recommendation addressed
OLCF should provide a separate list of major risks encountered in the past year, including how those risks were managed.	OLCF will document risks that have occurred and our response in future reports. 2/8/2012: Completed.	October 20, 2011: Proposed action is acceptable. Recommendation addressed
OLCF should provide a separate list of major risks of particular concern for the upcoming year, including how those risks would be mitigated and/or managed.	All OLCF risks in the risk register are either forward looking, or retired. The risks we document are for the upcoming period	October 20, 2011: Explanation provided. Recommendation addressed
OLCF should work with the program office and the other centers to standardize a method of presenting risks in the written report.	OLCF will participate in such discussions	October 20, 2011: OLCF participation acceptable. Recommendation will be addressed with discussion between sites and HQ. 10/4/2011: Headquarters action; will bring up at Facilities Staff meeting

8. Does the OLCF have a valid authority to 'operate'?

Recommendation	ORNL Action/Comments	HQ Action/Comments
OLCF should provide a copy of the ATO memo, similar to NERSC and ALCF OA reports.	The ORNL/OLCF ATO letter was intended to be included as we have done in past years. It was an oversight that it was omitted. It will be included in future reports. 2/8/2012: The ATO letter is not required in the February 2012 report.	October 20, 2011: Explanation provided and proposed action acceptable. Recommendation addressed

9. Are the performance metrics for the next year proposed by the OLCF reasonable?

Recommendation	ORNL Action/Comments	HQ Action/Comments
Customer metric 2 as currently described should be strengthened. It should not contain the phrase “if any”. The CY 2012 Target could be: “OLCF survey results related to Problem Resolution will be satisfactory (3.5/5) based on a statistically meaningful sample.	The metric should be changed as suggested. The metric should read, “OLCF survey results related to Problem Resolution will be satisfactory (3.5/5) based on a statistically meaningful sample.” Question #9 on the 2011 OLCF survey addresses this metric.	October 20, 2011: Identified solution acceptable. Recommendation addressed 10/4/2011: Review the current survey and if there are no questions relating to problem resolution then I suggest deleting metric. If there are questions relating to problem resolution, please enumerate those question numbers to clarify metric.
Customer metric 3 should be strengthened. It should not contain the phrase “if any.” The CY 2012 Target could be: “OLCF survey results related to User Assistance and Outreach and Scientific Liaison Support will be satisfactory (3.5/5) based on statistically meaningful sample.:	The metric should be changed. The metric should read, “OLCF survey results related to Overall User Assistance will be satisfactory (3.5/5) based on statistically meaningful sample.” Question #11 on the 2011 OLCF survey addresses this metric. 2/8/2012: Per the HQ Action/Comments to enumerate, the questions from the 2011 survey related to this metric are listed in Section 6 under Customer Metric 1.	October 20, 2011: Identified solution acceptable. Recommendation addressed 10/4/2011: Review the current survey and if there are no questions relating to User Assistance and Outreach, etc. then I suggest deleting metric. If there are questions relating these areas, please enumerate those question numbers to clarify metric.
Strategic metric 1 should refer to “refereed publications” rather than simply “publications”	As the review committee suggests, Strategic Metric #1 should refer to “refereed publications” rather than simply “publications”. 2/8/2012: Completed.	October 20, 2011: Identified solution acceptable. Recommendation addressed
With the introduction of GPUs in the coming year, ORNL needs to consider how to construct metrics that meaningfully represent GPU usage (core-hours will not be enough) and how to accurately measure that usage (e.g. what fraction of a GPU is being used). It's capability metric should be renegotiated with DOE to account for GPU usage.	OLCF agrees and is working on this. We have contacted each of the other ASCR facilities to discuss future allocation systems (core hours vs. node hours) with heterogeneous nodes. It is also important to understand how to judge capability usage in future architectures in a way that makes sense for each of the centers. 2/8/2012: The response to the recommendation is documented in Section 2.	October 20, 2011: Identified solution acceptable. This discussion needs to include headquarters as well. Recommendation addressed

User Results

HIGH PERFORMANCE COMPUTING FACILITY 2011 OPERATIONAL ASSESSMENT OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2012

1. USER RESULTS

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

OLCF RESPONSE: Yes. The OLCF has a dynamic user support model that is based on continuous improvement and a strong customer focus. We assess the effectiveness of our user support organization through a series of activities. One key element of that internal assessment is the annual user survey that is developed with input from qualified survey specialists and the DOE Program Manager. Again in 2011, OLCF users stated that they are very satisfied with the facility and its services. Details of the survey assessment are described below. The OLCF also 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 their support process for any item that does not meet a minimum satisfactory score. In 2011, the OLCF met or exceeded all metric targets for user satisfaction. Overall Customer Satisfaction rated a 4.2/5.0. 89.8% of reported problems by the users were resolved within three business days. The OLCF continues to implement and maintain operational activities designed to provide technical support, training, outreach and communication to current users and the next-generation of researchers as described below.

1.1 User Results Summary

The OLCF has developed and implemented a dynamic, integrated customer support model. It comprises various 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; comprehensive training programs and user workshops; and tools to reach and train the next generation of computer scientists.

Through a team of communications specialists and writers, the OLCF produces a steady flow of reports and highlights for potential users, the public, and sponsoring agencies. The OLCF communication infrastructure has been identified by ORNL as a best practice and other ORNL facilities are currently exploring ways to implement similar groups.

The OLCF offers many training and educational opportunities throughout the year for both current facility users and the next generation of HPC users (Section 1.5.5). This year, the OLCF's contributions in this area were recognized with several awards.

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

Table 1.1. 2011 User Result Metrics Summary

Metric Description	2010 Actual	2011 Target	2011 Actual
Customer Metric 1: Customer Satisfaction	4.3/5.0	Overall OLCF score on the user survey will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	4.2/5.0
	One question scored below 3.5. The OLCF received a 3.46 in response to “Frequency of unscheduled (unanticipated) outages”	Annual user survey results will show improvement in at least ½ of questions that scored below satisfactory (3.5/5.0) in previous period.	3.5/5.0
Customer Metric 2: User Support	4.3	The average of all user support ratings will be satisfactory (3.5/5.0)	4.1/5.0
Customer Metric 3: Problem Resolution	94.9%	80% of OLCF user problems will be addressed within three working days, either resolving the problem or informing the user how the problem will be resolved.	89.8%
Customer Metric 4: Problem Resolution	4.3/5.0	The average of all problem resolution ratings will be satisfactory (3.5/5.0).	4.2/5.0

1.2 User Support Metrics

The OA metrics for High Performance Computing (HPC) Facilities user support as assessed by the annual user survey are:

- Overall satisfaction rating for the Facility is satisfactory;
- Average of all user support questions on user surveys is satisfactory; and
- Improvement on past year unsatisfactory ratings as agreed upon with the Facility’s DOE Program Manager

The OLCF metric targets and calendar year 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	CY2010		CY2011	
	Target	Actual	Target	Actual
Overall satisfaction rating	3.5/5.0	4.3/5.0	3.5/5.0	4.2/5.0
Avg. of user support ratings	3.5/5.0	4.3/5.0	3.5/5.0	4.1/5.0

1.2.1 Overall Satisfaction Rating for the Facility

Users were asked to rate 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.

Users were asked to rate their overall satisfaction with the OLCF. Mean responses were between 4.0 and 4.3 showing a high degree of satisfaction with OLCF across project classifications (Table 1.3). The calculated mean was 4.2/5.0, well above the stated metric of 3.5. Of the optional questions, this question had one of the highest numbers of responses, with 92% of respondents providing their opinions. Of these, a total of 89% (207 respondents) reported being “Satisfied” or “Very Satisfied” with OLCF overall.

Key indicators from that 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 ^a		
		INCITE	ALCC	Director's Discretionary
Overall Satisfaction with the OLCF	4.2/5.0	4.3/5.0	4.0/5.0	4.3/5.0
Effectiveness of Account Mgmt Staff	4.2/5.0	4.2/5.0	4.2/5.0	4.4/5.0
Effectiveness of problem resolution	4.2/5.0	4.2/5.0	4.1/5.0	4.2/5.0
Overall System Performance of the Cray XT5	4.0/5.0	4.0/5.0	4.1/5.0	4.1/5.0

1.2.2 Average Rating across All User Support Questions

The calculated mean of all answers to all user support questions on the 2011 survey was 4.1/5.0, indicating that OLCF exceeded the 2011 user support metric. In response to open-ended question about the best qualities of the OLCF, User Assistance was listed as the top choice.

After reviewing the results of the 2010 survey in detail, the OLCF determined that additional questions were needed to solicit better feedback regarding user support services, especially in the areas of training and communications. The 2010 survey focused mostly on problem resolution and did not thoroughly represent all of the areas of user support. With feedback from the ORISE survey specialist, the OLCF added additional questions to solicit better feedback in the area of user support including training, communication, and the OLCF websites. A brief summary of some of the findings is listed below. See Table 1.4 for overall satisfaction results from each of these areas.

User Assistance Evaluation

- For support services used, 99% of the 221 respondents reported using the User Assistance Center (UAC), followed by 27% using the Scientific Computing/Liaison service, 10% using visualization, and 4% using End-to-End.
- When asked to rate their overall satisfaction with the user support services provided by the OLCF, the average response was 4.1/5.0. 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 majority of the users (86% and 80%, respectively) were “Satisfied” or “Very Satisfied.”

Training and Education

- The majority of OLCF users said “Yes” (49%) or “Maybe” (44%) to the prospect of attending future OLCF training, based on their previous experience.
- The number one reason users gave for not participating in any live training events was that they do not have the time to attend (56%).
- Documentation was listed as the top choice (69%) for training preference, followed by online training (55%).

OLCF Communications

- Eighty-three percent of respondents (236) rated their overall satisfaction with communications from the OLCF as satisfied or very satisfied.
- Respondents indicated the email message of the week was most useful.

OLCF Website

- Ninety-seven percent of respondents indicated that they had visited the <http://olcf.ornl.gov> web site. Of these users (237), 37% indicated that they visit the site once a week or more, 3% of whom indicated that they visit the site every day. Only seven respondents indicated they had never visited the site.

Table 1.4. Overall Satisfaction Results

	Overall Satisfaction with User Assistance	Overall Satisfaction with Training Activities	Overall Satisfaction with OLCF Communications	Overall Satisfaction with the OLCF Website
Number Surveyed	252	252	252	252
Number of Respondents	232	48	236	224
Rating	4.1/5.0	4.2/5.0	4.1 /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 in both 2008 and 2009, and only one item scored below 3.5 in 2010. This item was related to the frequency of unscheduled outages on the Cray XT5. The rating for this question on the 2011 survey met the metric at 3.5. There were no ratings below 3.5 on the 2011 survey.

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 used appropriately. 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 the user support section, additional questions were added to the survey to solicit better feedback to assist in the OLCF's continuous improvement efforts.

The 2011 survey was launched on October 21, 2011 and remained open for participation through December 16, 2011. The surveys were sent electronically to all individuals with active accounts on an ALCC, DD, or INCITE project. A total of five targeted notifications were sent to each recipient including the initial survey invitation from ORISE and then four follow-up reminders from Buddy Bland (OLCF Project Manager), Ashley Barker (UA Group Lead), ORISE, and the OLCF User Council. The survey was advertised on the OLCF website and mentioned in the weekly communications email sent to all users from the UAO team. Survey responses were tracked on a daily basis to identify the effectiveness of the various communication methods. Reminder notifications from Buddy Bland and the OLCF User Council proved most effective at soliciting responses. At the end of the two-month survey period, 252 users completed the survey out of 813 possible respondents, an overall response rate of 31%.

Information was collected about the various users, the user experience with OLCF, and the OLCF support capabilities. Attitudes and opinions on the performance, availability, and possible improvements for OLCF and its staff were also solicited. Data collected from the users' survey was analyzed by ORISE using

both quantitative and qualitative methods. The two fundamental goals that drove the collection and subsequent analysis were to understand the types of users and to understand their needs and preferences with the systems. Analysis included basic descriptive statistics and qualitative coding of responses to open-ended questions (using grounded theory). 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 2011 survey can be found on the OLCF website at <http://www.olcf.ornl.gov/media-center/center-reports/2011-outreach-survey/>.

OLCF has a relatively equally balanced distribution of users in terms of their length of time using the systems. Reference Table 1.5.

Table 1.5. User Survey Participation

	2010 Survey	2011 Survey
Total Number of Respondents (Total percentage responding to survey)	402 (36%)	252 (31%)
New Users (OLCF User < 1 Year)	31%	31%
OLCF User for 1–2 Years	29%	30%
OLCF User > 2 Years	40%	39%
Used User Assistance Center at least 1 time	80%	77%

Survey respondents were asked to classify the program types with which they were affiliated. Reference Table 1.6.

Table 1.6. User Survey Responders by Program Type

Program	Response Rate
INCITE ¹	69%
Director's Discretionary	25%
ALCC ²	14%
Other	6%

Statistical Analysis of the Results

Statistical analysis of four key survey areas is shown in Table 1.7. These reflect overall Facility satisfaction, services, and computational resources.

Table 1.7. Statistical Analysis of Key Results

	Overall Satisfaction	Effectiveness of Account Mgmt Staff	Effectiveness of Problem Resolution	Overall System Performance of the Cray XT5
Number Surveyed	252	252	252	252
Number of Respondents	232	206	206	210
Mean	4.2/5.0	4.2/5.0	4.2/5.0	4.0/5.0
Variance	0.70	0.77	0.88	0.53
Standard Deviation	0.84	0.88	0.94	0.73

¹ Innovative and Novel Computational Impact on Theory and Experiment

² Advanced Scientific Computing Research Leadership Computing Challenge

1.4 Problem Resolution Metrics

The OA Metrics for Problem Resolution are:

- Average satisfaction ratings for Problem Resolution related questions on the user survey are satisfactory or better; and
- 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 working days.

1.4.1 Problem Resolution Metric Summary

In the majority of instances, the OLCF can resolve the reported problem directly. This 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 in which our ability to resolve the root cause of the issue is limited due to factors beyond our control. In this scenario, "addressing the problem" requires that User Assistance has identified and carried out all corrective actions at their disposal for the given situation. For example, if a user reports a suspected bug in a commercial product, prudent measures 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.

Table 1.8. Problem Resolution Metric Summary

Survey area	CY2010		CY2011	
	Target	Actual	Target	Actual
% of problems addressed in 3 working days	80%	94.9%	80%	89.8%
Avg. of problem resolution ratings	3.5	4.2	3.5	4.2

The OLCF uses Request Tracker software (RT) to track queries and ensure that response goals are met or exceeded. Users may submit tickets via email, the online request form, or by phone. 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 2,460 tickets in response to user queries for CY 2011 (Figure 1.1). The OLCF exceeded the problem resolution metric and responded to 89.8% of queries within three business days (Table 1.8).

Figure 1.1. Number of Helpdesk Tickets Issued per Month.

Each query 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 service tailored to each customer's needs. While UAO 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 reported in 2011 was related to jobs/batch queues (Figure 1.2).

Figure 1.2. Categorization of Helpdesk Tickets

1.5 User Support and Outreach

The OA data for User Support and Outreach include:

- Anecdotal evidence of in-depth collaborations between Facility staff and the User community
- Summary of the training events conducted during this period.

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 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. Last, a strong outreach component is needed to interface with the next generation of HPC users, the external media, and the public. In this section, we will discuss 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 OLCF provides two complementary OLCF user support vehicles: the User Assistance and Outreach Group (UAO) and the Scientific Computing Group (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 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:

- Enable access to OLCF resources
- Help users compile and debug large science and engineering applications
- Identify and resolve system-level bugs in conjunction with other technical staff and vendors
- Install third-party applications and provide documentation for usage
- Engage center staff to ensure users have up-to-date information about OLCF resources and to solicit feedback
- Research, develop, and maintain reference and training materials for users
- Communicate with users
- Develop and deliver training
- User advocates

UAO's regular ticket report meetings, discussed in last year's report, are another OLCF innovation that has paid huge dividends in efficient customer service. The UAO team meets weekly to review problem tickets and discuss enhancements or improvements needed based on their interactions with the users. There were several outcomes from the 2011 UAO meeting including the creation of new mobile phone applications for users that show the real-time status of the OLCF resources. According to the 2011 survey, 88% of respondents indicated they would be likely to download a mobile phone application, and the top two requested features that users would like to see in a phone application include "view a snapshot of the queue" and "view real-time system status notifications." UAO analysts developed applications for both the Android and iPhone platforms. After a period of internal testing, the application was released to users at the end of CY 2011.

In addition, UAO developed "opt-in" notice lists that provide automated notices about the status of OLCF systems, as well as more detailed updates from the OLCF staff as needed. Users can subscribe to receive notifications about particular systems short- or long-term (e.g., for as little as 1 week or for an entire calendar year). Thus, users now have numerous ways to check the status of the machines including checking the website, via email or Tweets, and/or on their mobile phone devices.

Some additional notable UAO contributions in 2011 include the following.

- Developed new materials for Titan website including a FAQ, tutorials, and a timeline (5,346 hits since Nov debut)
- Organized and delivered new Titan training events (248 people collectively attended the events)
- Developed a "Welcome Packet" for new users based on feedback received from the 2010 survey
- Developed a Jaguar user guide and started work on a Titan user guide

- Responded to several suggestions from the previous survey for improvements to the website, documentation, and dashboard, including a “Getting Started” area that consolidated information for new users

1.5.3 Scientific Computing Liaisons

The OLCF pioneered a total user support model widely recognized as a best practice for HPCs: 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. The liaison program is one of the reasons for the success of the OLCF. Several examples of the impact to the user community are described here.

Adapting codes to run on hybrid CPU-GPU architectures

SciComp’s Mike Brown, a molecular dynamics (MD) specialist with a background in both the biomedical and the computer sciences, is working on adapting LAMMPS (Large-Scale Atomic/Molecular Massively Parallel Simulator) and other codes to run on hybrid CPU-GPU machines like the OLCF’s next generation Titan. LAMMPS is a classical MD code that can be used to model atoms or, more generically, as a parallel particle simulator at the atomic, meso, or continuum scale.

This past year, working with Axel Kohlmeyer, ICMS associate director and an expert on MD codes like LAMMPS; NVIDIA’s Peng Wang; SNL’s Steve Plimpton, lead developer of LAMMPS; and Arnold Tharrington, lead on the OLCF LAMMPS CAAR effort, Brown researched algorithms that would allow the LAMMPS MD simulator to run with GPU acceleration on the OLCF’s CPU-GPU test cluster.

The LAMMPS Accelerator Library (<http://users.nccs.gov/~wb8/gpu/lammps.htm>), now distributed as part of the main LAMMPS software package and thus freely available to all MD researchers, is one of the main outcomes of this research to date. The library, which allows simulations to be run between 2 and 14 times faster on Infiniband GPU clusters, is already being applied by LAMMPS users for science applications and will facilitate an early capability for INCITE users to utilize the impressive floating-point capabilities on the Titan machine with full compatibility with all of LAMMPS traditional CPU features.

Improving Performance and Scalability through Profiling Analysis

SciComp staff members like Rebecca Hartman-Baker are aggressively using advanced profiling tools like the Vampir (Visualization and Analysis of MPI Resources) suite of tools. VampirTrace instruments codes and produces trace files when run. The trace files are then loaded into Vampir, which is used to visualize the trace; the output is a timeline trace of the workings of an application with the timeline of the code along the x-axis and processor numbers along the y-axis. Events are represented by colored blocks, dots, and lines, and subroutines or functions of particular interest can be color-coded to stand out.

An example success story in the application of profiling tools is the BIGSTICK configuration-interaction shell-model code, which is used to solve the general many-fermion problem (important in nuclear physics). While the code is supposed to work well on both serial and parallel machines, when Hartman-Baker profiled it using VampirTrace and Vampir, she found that the code had a number of inefficiencies in

its implementation of the Lanczos method for eigenvalues and eigenvectors. She compiled the results and supporting visualizations into a report in which she outlined suggestions for improving the algorithm, based on both the Vampir analysis and her own expertise in numerical algorithms. Hartman-Baker’s analysis and suggestions were discussed at the 2011 UNEDF (Universal Nuclear Energy Density Functional) meeting, and the project team is now planning to test the reformulated code in preparation for an INCITE application in 2013. Because of Hartman-Baker’s initiative, a potential future INCITE awardee has been helped to “get up to speed,” which Hartman-Baker finds particularly gratifying as the OLCF is always looking for new projects. It’s also a great example of how the OLCF and its staff members provide continuous support to the larger HPC community.

In a similar case, Hartman-Baker was asked by the code developers to profile the j-coupled version of NUCCOR. This is a nuclear physics code that takes advantage of symmetries in certain nuclear configurations to perform energy calculations on larger nuclei than can currently be studied with this code in the nonsymmetrical case. Profiling showed that on the small test problem, the code was spending more than half its time in a subroutine called sort. Hartman-Baker’s suggested solution was to implement a heap sort, which would reduce sorting to about 3% of the total run time; however, in consultation with the authors of the code, it was determined that sorting was unnecessary, so the sorting subroutine was removed altogether, resulting in a 30% speedup on the full problem.

Supporting Software Applications

VASP, the Vienna Ab-initio Simulation Package, is a workhorse in the materials science world and one of the premier electronic structure codes used by a number of INCITE and DD projects. VASP isn’t open source software, so OLCF staff can’t really develop it, yet must find a way to optimize it on OLCF platforms. SciComp member Markus Eisenbach provides precompiled versions of both of the commonly used VASP releases (4.6 and 5.2, released in 2010), optimized for OLCF, to licensed users on OLCF systems. The most recent version, 5.2, provides significant new physics capabilities such as exact exchange and hybrid functionals, and while it ports reasonably well, Eisenbach’s background in condensed matter physics, combined with his HPC expertise, enables him to better understand the needs of users and help them get the most from the VASP code on OLCF machines.

Denovo is the ORNL radiation transport code developed to take advantage of the computational power of high-performance computers such as Jaguar. Because of Denovo’s broad applicability to radiation transport modeling, new applications continue to be found, including assistance with the Fukushima reactor (see separate visualization story below). Thanks to SciComp liaison Wayne Joubert and the Denovo team’s work, the latest version of Denovo runs $2 \times$ faster than the previous code on conventional processors, runs an astounding $40 \times$ faster on the NVIDIA Fermi GPU compared to a Jaguar processor core, and is significantly more scalable than its predecessor (scaling up to 200,000 cores).

The Denovo team found that the energy-set reduction operation was the least scalable part of the enhanced code. Team members asked Joubert to help them with a solution to the problem. The code originally used `MPI_Allreduce`, a generic function, for the energy-set reduction operation. Using his knowledge of MPI, Joubert was able to recommend a fairly obscure offshoot, `MPI_Reduce_scatter` that could be used for this case as an alternative method to perform the reduction operation. By simplifying the information that the various processors get, `MPI_Reduce_scatter` improves communication performance and memory usage, making the reduction step

run $3 \times$ faster. This is a classic example of the type of work that liaisons do regularly for their projects.

Responding to Time-Critical Needs

What we do is critically important, not only to national, but also to world security. This was never more evident than in the OLCF's rapid response to the Fukushima nuclear accident. In the days following the March 11, 2011, massive earthquake and subsequent tsunami, DOE staff and experts from ORNL and other national laboratories sprang into action to help collect, analyze, and interpret data to provide the Japanese government and others with critical information. One of these groups consisted of OLCF visualization experts Jamison Daniel, Mike Matheson, and Dave Pugmire. Following the earthquake and tsunami, there was concern that the spent fuel pool had been compromised and that water had leaked out as a result. A loss of water could result in fuel rod heating and damage. Further, because the spent fuel pool consisted of rods that had been removed from the reactor at different times, the response to the level of the water would be different for each set of rods.

Figure 1.3. Rendering of the Fukushima Reactor Building Spent Fuel Rod Pool.

Working with ORNL Reactor & Nuclear Systems Division staff members, the visualization liaisons took blue prints and CAD models of the reactor building, spent fuel pool, and fuel bundle layouts to create a three-dimensional (3-D) model of the Fukushima plant. This 3-D model was then read into high-end rendering packages where camera animation could be added to explore the condition of the reactor (Figure 1.3). Two simulations were incorporated into the visualizations, which showed the temperature of fuel rods, the temperature of the water, and the dose levels as a function of the level of the water. These time-critical simulations were executed in a very short period of time.

Pulling Information from Raw Data

With an INCITE grant of 25 million hours, Jeremy Smith is performing highly parallelized multi-length-scale computer simulations to help understand the physical causes of the resistance of plant cell walls to hydrolysis—the major technological challenge to developing cellulosic bioethanol. This was problematic in two respects: (1) the time dependent nature of the simulations was difficult to understand with simple graphics and (2) some of the large amount of data to be processed obscures other data key to gaining insights. Because advanced visualization techniques, including animation, can aid in the analysis of such data, Mike Matheson, a visualization liaison with a background in engineering, was assigned to the team. Mike's experience with HPC, and especially visualization, enabled him to select the software most suitable to this application. Using Tachyon and the Blender 3D software, Matheson developed a method to deal with the obscuring data in an intelligent manner so Smith and his team could “see” what was important. The high

quality renderings combined with this technique enhanced the team's ability to explain the simulations, especially to others, and enabled them to gather detailed knowledge of the fundamental molecular organization, interactions, mechanisms, and associations of bulk lignocellulosic biomass (Figure 4.11), as well as other insights, from the data. As with other SciComp success stories, the success of this work was based on the close collaboration between Matheson and the project team.

1.5.4 Communications

Eighty-three percent of respondents rated their overall satisfaction with communications from the OLCF as satisfied or very satisfied and submitted an overall rating for communications as 4.1/5.0. The OLCF uses various methods to communicate with users including:

- Weekly email message
- Welcome packet
- General email announcements
- Opt-In email notification lists
- Message of the Day (MOTD)
- OLCF websites
- Phone applications
- Conference calls
- OLCF User Council
- One on one interactions through liaisons and analysts
- Social networking vehicles

Survey respondents indicated that the weekly email message was the most useful form of communication. Due to low attendance in 2011, the OLCF has plans to revamp the user conference calls with the goal of making the calls more useful to the users.

1.5.5 Training

Workshops and seminars are another important component of the customer support model. They provide an additional opportunity to communicate and act as a vehicle to reach out to the next generation of computer scientists. In 2011, the OLCF conducted several activities to communicate and/or train users. Approximately 703 people participated in the communications and training activities.

The OLCF held the following specific training- and outreach-related workshops and seminars in CY2011. A summary of these events is shown in Table 1.9.

Table 1.9. Training Event Summary

Event Type	Event Description	Event Date	Participants
BOF	Heterogeneous Supercomputing on Titan	11/17/11	45
Educational Outreach	HPC Fundamentals	Summer 2011	44
Educational Outreach	Crash Course in Supercomputing	6/16/11	112
Educational Outreach	ARC Mentorship	July 11-22, 2011	10
Seminar	INCITE Proposal Writing Seminar	3/21/12	38

Event Type	Event Description	Event Date	Participants
Seminar Series	LCF Seminar Series: Temporal Debugging via Flexible Checkpointing: Changing the Cost Model, Gary Cooperman, Northeastern University	1/25/11	40
User Meeting	OLCF User Meeting	3/11/11	43
User Teleconferences	OLCF User Council Call	1/18/2011	Not tracked
User Teleconferences	OLCF User Council Call	2/16/2011	Not tracked
User Teleconferences	User Conference Call	3/8/2011	2
User Teleconferences	OLCF User Council Call	4/27/11	Not tracked
User Teleconferences	OLCF User Council Call	5/19/11	Not tracked
User Teleconferences	OLCF User Council Call	6/29/11	Not tracked
User Teleconferences	User Conference Call	6/7/11	0
User Teleconferences	OLCF User Council Call	10/3/11	Not tracked
User Teleconferences	OLCF User Council Call	10/3/11	Not tracked
User Teleconferences	User Conference Call	10/4/11	0
Workshop	Introduction to CUDA	1/20/11	15
Workshop	OLCF Spring Training	Mar 7–10, 2011	80
Workshop	Lustre User Group Meeting	Apr 12–14, 2011	163
Workshop	Vampir Training Class	5/17/11	25
Workshop	Visualization with VisIt 2011	6/14/11	44
Workshop	Tau Workshop	6/29/11	20
Workshop	Introduction to OLCF-3 Webinar	7/26/11	74
Workshop	PGI Workshop	7/25/11	28
Workshop	Introduction to Titan Webinar	7/26/11	66
Workshop	R Training	July 29, Aug 3, & Aug 5, 2011	85
Workshop	Titan Summit	Aug 15-17, 2011	63
Workshop	CAPS Training	Sept 19-22, 2011	15
Workshop	OLCF Fall Training	Oct 18-19, 2011	25

The OLCF began live webcasting of workshops and seminars in 2011 to broaden participation. Initial feedback from the User Council is that webcasting is very helpful for those that do not have the budget or time to attend training events. In further efforts to improve OLCF training events, the OLCF began conducting surveys at the end of each training day to get targeted feedback. The results of the surveys in early 2011 showed that users wanted more hands on training. In response, the OLCF incorporated several hands on activities in the most recent workshops. The surveys from the most recent workshop along with feedback from members of the User Council indicate that users liked the hands on activities that we incorporated. We received numerous comments on the surveys like, “Great job. Like the hands-on training.”

The OLCF offered several new events this year. One of the most notable was the 2011 Lustre User Group meeting (LUG). As a leader in parallel file systems, the OLCF led the organization of the 2011 LUG meeting. This was the first user-led LUG meeting, previously hosted by Oracle, and marked the transition of leadership to the broader user community. With more than 160 attendees from more than 60 organizations, LUG 2011 was a tremendous success. “LUG 2011 is the first LUG that is completely community driven. It opens a promising new area in the Lustre community,” said Jacques-Charles Lafoucrière, Chef de Service at CEA. The LUG offers participants opportunities to share knowledge, ideas, and achievements with a diverse audience,” said Stephen Simms, Data Capacitor project lead at Indiana University.

In addition, the OLCF offered several events to help users prepare for using Titan including a BOF at Supercomputing, a Titan Summit, and two webinars to discuss

the implications of the upgrade for users. Approximately 250 people attended these four Titan events.

Training the Next Generation

The OLCF maintains a broad program of collaborations, internships, and fellowships for young researchers. From January 1, 2010, through December 31, 2010, the OLCF supported 32 faculty, student interns, and postdoctoral researchers. Twenty-eight faculty, student interns, and postdoctoral researchers were supported from January 1, 2011, through December 31, 2011. Several examples of education outreach are listed below.

- For the third straight year, students and teachers from around Appalachia gathered at ORNL this past summer for interactive training from some of the world's leading computing experts. The summer camp, a partnership between ORNL and the ARC Institute for Science and Mathematics, took place July 11–22, 2011.
- OLCF staff helped National Geographic's award-winning middle school science education program "The JASON Project" capture a prestigious CODiE award in early 2011 for the geology curriculum "Operation Tectonic Fury," described in the 2010 OLCF Operational Assessment report. James J. Hack, director of the National Center for Computational Sciences also hosted JASON students and helped them gain a better understanding of the role of climate on our earth's ecosystem.
- OLCF's Doug Fuller was the technical chair of the SC11 Student Cluster Competition (SCC). In addition, Dustin Leverman provided technical support for all competitors. In the SCC, small teams of students each built a computer cluster capable of running open-source software and meeting HPC Center benchmarks.

DOE Recognizes OLCF Outstanding Mentors

The Department of Energy (DOE) recently awarded the Oak Ridge Leadership Computing Facility (OLCF) staff members Jim Rogers and Bobby Whitten with Outstanding Mentor Awards. Coordinated by the SC Workforce Development for Teachers and Scientists, the award recognizes mentors for their personal dedication to preparing students for careers in science and science education through well-developed research projects. Winners are nominated by their mentees.

1.5.6 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 OLCF provided tours for 782 distinct groups in CY 2011. The OLCF highlighted the research from many different projects. These highlights can be found on the OLCF website. In 2011, the OLCF produced more than 60 new highlights. In addition to the highlights, the team also produced the 2010/2011 Annual Report and collaborated on the INCITE in Review document. The team also directed a Titan video that debuted at Supercomputing and developed the Titan website to help inform the users, sponsor, and the general public about the new system. The OLCF also produced a poster for the NUFO 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 and the center resources and provided video images of the facility. Last, the OLCF was highlighted in several different media outlets including CNN and Popular Science.

1.6 User Support Conclusion

In conclusion, the OLCF possesses effective processes for supporting the customers and resolving problems, and conducts a strong training and outreach program as supported by:

- The results of the 2011 OLCF survey indicate that OLCF met and exceeded the metrics for overall user satisfaction (4.2/5.0), user services (4.1/5.0), and problem resolution (4.2/5.0). According to the survey results, approximately 90% of users reported being satisfied or very satisfied overall with the OLCF.
- The OLCF exceeded the problem resolution metric to address user problems within 3 business days.
- In response to open-ended question about the best qualities of the OLCF, User Assistance was listed as the top choice.
- The OLCF effectively communicates with users using numerous channels. 83% of respondents rated their overall satisfaction with communications as “satisfied” or “very satisfied.”
- The OLCF conducts numerous outreach activities throughout the year in the form of workshops, seminars, tours, reports, highlights, websites, conference attendance, conference calls, user meetings, and more. Over 700 people participated in OLCF led training and outreach events in 2011.
- Users indicated they were satisfied with the OLCF training effort and rated it a 4.2/5.0. The majority of OLCF users said “Yes” (49%) or “Maybe” (44%) to the prospect of attending future OLCF training, based on their previous experience.

Business Results

HIGH PERFORMANCE COMPUTING FACILITY 2011 OPERATIONAL ASSESSMENT OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2012

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 their 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

The 2010 OA report described the upgrade of the existing Cray XT5 from AMD Opteron™ 1354 quad-core processors to AMD Opteron 2435 six-core processors, providing a 50% increase in the resources available for OLCF users. Through the period ending October 09, 2011, the underlying Cray XT5 hardware configuration remained unchanged, with steady-state operation delivering well over 1.4 billion compute hours in 2011.

In accordance with the Authorization to Proceed provided by the DOE Office of Science (SC) Office of Project Assessment (OPA), the OLCF began an upgrade of the Cray XT5 to an XK6. This had operational impacts to a portion of the existing Cray XT5 beginning in October 2011 as a partition of the existing system was upgraded with the XK6 compute blades and Gemini interconnect. This new XK6/Gemini partition was released to users for the second half of December 2011 to allow the OLCF to upgrade the remaining XT5 partition. This strategy for partitioning the system and validating the upgrade process and architecture for a portion of the existing system was approved by the SC OPA as part of the formal Critical Decision (CD) process. The full Cray XK6 system will enter production in the period covered by the 2012 OA Review.

The OLCF provided the following computational resources in 2011 for scientific research (reference Table 2.1). The Cray XK6 system described here is the smaller partition made available in December 2011.

Table 2.1. OLCF Production Computer Systems, 2011

System	Type	CPU	Computational Description			Interconnect
			Nodes	Node Configuration	Aggregate Memory	
JaguarPF	Cray XK6	2.2 GHz AMD Opteron 6274 (16-core)	8,928	16-core SMP	285,696 GB DDR3-1600	Gemini (160 GB/sec)
JaguarPF	Cray XT5	2.6 GHz AMD Opteron 2435 (6-core)	18,688	2x6-core SMP	299,008 GB DDR2-800	SeaStar2+ (57.6GB/sec)
Jaguar	Cray XT4	2.1 GHz AMD Opteron 1354 (4-core)	7,832	4-core SMP	62,656 GB DDR2-800	SeaStar2 (45.6GB/sec)

The business results reported in this section 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, and the very small overlap in December 2011 beginning with the introduction of the Cray XK6.

Table 2.2. OLCF HPC System Production Dates, 2007 – Present

System	Type	Production Date	Performance End Date	Notes
JaguarPF	Cray XT5	August 19, 2008	July 28, 2009	151,000 AMD Opteron cores
JaguarPF	Cray XT5	September 25, 2009	October 9, 2011	224,256 AMD Opteron cores
JaguarPF	Cray XT5	October 10, 2011	October 16, 2011	162,240 AMD Opteron cores
JaguarPF	Cray XT5	October 17, 2011	December 11, 2011	117,120 AMD Opteron cores
JaguarPF	Cray XK6	December 12, 2011	January 04, 2012	142,848 AMD Opteron cores
JaguarPF	Cray XK6	January 05, 2012	NULL	299,008 AMD Opteron cores

The production date used for computing statistics is either the initial production date or the date of the last upgrade to the computational resource. Both new and upgraded systems have first-year business result targets in the year after they enter production. 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 HPSS Archive System, and the external Lustre File Systems (reference Tables 2.3-2.5).

Table 2.3. OLCF Business Results Summary for HPC Systems

	Measurement	2010 Target	2010 Actual	2011 Target	2011 Actual
Cray XT5 (JaguarPF)	Scheduled Availability	85.0%	94.1%	95.0%	96.37%
	Overall Availability	80.0%	89.2%	90.0%	92.88%
	MTTI (hours)	NAM	45.2	NAM	60.38
	MTTF (hours)	NAM	59.5	NAM	79.66
	Total Usage	NAM	81.64%	NAM	87.11%
	Core Hours Used	NAM	1,435,391,110	NAM	1,428,874,052
	Core Hours Available	NAM	1,758,135,000	NAM	1,640,290,505
	Capability Usage				
	INCITE Projects	NAM	36.9%	NAM	47.8%
	All Projects	35.0%	39.0%	40.0%	54.0%
Cray XT4 (Jaguar)	Scheduled Availability	95.0%	97.1%	95.0%	97.58%
	Overall Availability	90.0%	94.9%	90.0%	97.09%
	MTTI (hours)	NAM ³	95.8 hours	NAM	78.67 hours
	MTTF (hours)	NAM	134.0 hours	NAM	87.80 hours
	Total Usage	NAM	87.61%	NAM	90.73%
	Core Hours Used	NAM	228,266,799	NAM	39,079,672
	Core Hours Available	NAM	260,548,188	NAM	43,070,274
	Capability Usage				
	INCITE Projects	NAM	NAM	NAM	39.1%
	All Projects	NAM	NAM	NAM	57.1%

Table 2.4. OLCF Business Results Summary for HPSS

	Measurement	2010 Target	2010 Actual	2011 Target	2011 Actual
HPSS	Scheduled Availability	95.0%	99.6%	95.0%	99.81%
	Overall Availability	90.0%	98.6%	90.0%	98.65%
	MTTI (hours)	NAM	291.8	NAM	224.73
	MTTF (hours)	NAM	501.3	NAM	628.03

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

	Measurement	2010 Target	2010 Actual ⁴	2011 Target	2011 Actual
Widow 1	Scheduled Availability	95.0%	99.7%	95.0%	99.26%
	Overall Availability	90.0%	99.0%	90.0%	97.95%
	MTTI (hours)	NAM	481.6	NAM	536.27
	MTTF (hours)	NAM	623.8	NAM	785.84
Widow 2	Scheduled Availability	NIP ⁵	NIP	95.0%	99.93%
	Overall Availability	NIP	NIP	90.0%	99.34%
	MTTI (hours)	NIP	NIP	NAM	966.92
	MTTF (hours)	NIP	NIP	NAM	1750.78
Widow 3	Scheduled Availability	NIP	NIP	95.0%	99.95%
	Overall Availability	NIP	NIP	90.0%	99.36%
	MTTI (hours)	NIP	NIP	NAM	967.10
	MTTF (hours)	NIP	NIP	NAM	1751.09

³ NAM – Not a metric. No defined metric nor target exists for this system. Data provided as reference only.⁴ In 2010, there was a single very large file system (Widow). In 2011, this file system was segregated into three distinct segments as part of a upgrade that improved metadata performance.⁵ NIP – Not in Production. This system was not available as a production resource.

2.3 Resource Availability

2.3.1 Scheduled Availability

2011 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 rare scenario.

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

As shown in Table 2.6, the OLCF has exceeded the scheduled availability targets for the facility's computational resources for both 2010 and 2011.

Table 2.6. OLCF Business Results Summary: Scheduled Availability

	System	2010 Target	2010 Actual	2011 Target	2011 Actual
Scheduled Availability	Cray XT4	95.0%	97.1%	95.0%	97.58%
	Cray XT5	85.0%	94.1%	95.0%	96.37%
	HPSS	95.0%	99.6%	95.0%	99.81%
	Widow 1	95.0%	99.7%	95.0%	99.26%
	Widow 2	NIP	NIP	95.0%	99.93%
	Widow 3	NIP	NIP	95.0%	99.95%

Assessing Impacts to Scheduled Availability

The operational posture for both the Cray XT4 and XT5 contains a regularly scheduled weekly preventative maintenance (PM) period⁶. 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.

There were 27 scheduled outages in the 52-week period ending December 31, 2011.

⁶ PM for Cray XT4 terminated February 28, 2011.

2.3.2 Overall Availability

2011 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 both 2010 and 2011.

Table 2.7. OLCF Business Results Summary: Overall Availability

	System	2010 Target	2010 Actual	2011 Target	2011 Actual
Overall Availability	Cray XT4	90.0%	94.9%	90.0%	97.09%
	Cray XT5	80.0%	89.2%	90.0%	92.88%
	HPSS	90.0%	98.6%	90.0%	98.65%
	Widow 1	90.0%	99.0%	90.0%	97.95%
	Widow 2	NIP	NIP	90.0%	99.34%
	Widow 3	NIP	NIP	90.0%	99.36%

Increasing System Availability

The SA and OA targets are predicated on many factors, including the large physical scale of the system, the aggregate calculation of the failure rates of many disparate components, the architecture of the system and its resiliency to interrupt or failure due to component or software failure.

As part of existing HPC operations, ORNL and Cray continually assess the hardware component failure rates in the XT5 system against both the expected component failure rates using original equipment manufacturer and their own qualification data and against the failure rates of the same components at other Cray installations. During this reporting period, ORNL and Cray identified higher than expected failure rates for voltage regulator modules (VRM). On the Cray compute blade, each VRM is a step-down DC to DC converter that provides the associated 6-core AMD Opteron (Istanbul) the appropriate supply voltage of +1.3V from the higher voltage (nominally +12V, with $\pm 5\%$ variance) supplied to the compute blade.

VRM failures are typically associated with compute nodes powering down, heartbeat faults and link-inactive failures. These affect the SeaStar interconnect fabric, and can produce a condition that causes an unscheduled system interrupt. Cray and ORNL investigated multiple engineering solutions to this event and identified and implemented a solution related to a change to the VRM input voltage that has significantly reduced the failure rate of the VRM. The initial implementation of this engineering change was started in mid-June 2011. The dramatic reduction in VRM failures is shown in Figure 2.1.

Figure 2.1. Eliminating VRM failures increases system stability.

2.3.3 Mean time to Interrupt

2011 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	2010 Target	2010 Actual	2011 Target	2011 Actual
MTI (hours)	Cray XT4	NAM	95.8	NAM	78.67
	Cray XT5	NAM	45.2	NAM	60.38
	HPSS	NAM	291.8	NAM	224.73
	Widow 1 ⁷	NAM	481.6	NAM	536.27
	Widow 2	NIP	NIP	NAM	966.92
	Widow 3	NIP	NIP	NAM	967.10

2.3.4 Mean Time to Failure

2011 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

⁷ Due to the extremely long uptime of the Widow files systems, the formula for MTTF 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.

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

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

	System	2010 Target	2010 Actual	2011 Target	2011 Actual
MTF (hours)	Cray XT4	NAM	134.0	NAM	87.80
	Cray XT5	NAM	59.5	NAM	79.66
	HPSS	NAM	501.3	NAM	628.03
	Widow 1	NAM	623.8	NAM	785.84
	Widow 2	NIP	NIP	NAM	1750.78
	Widow 3	NIP	NIP	NAM	1751.09

2.4 Resource Utilization

2011 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 XT4 for the period January 1 – December 30, 2011, 39,079,672 core-hours were delivered from a scheduled maximum of 43,070,274 core-hours. This resulted in total system utilization for the Cray XT4 of 90.73%.

For the Cray XT5 for the period January 1 – December 30, 2011, 1,428,874,052 core-hours were delivered from a scheduled maximum of 1,640,290,505 core-hours. This resulted in total system utilization for the Cray XT5 of 87.11%.

Revising the Definition for Resource Utilization Measurements

The concept of core-hours remains applicable for 2011. There is also a direct translation for 2011 (and earlier years, for historical purposes) from core-hours to node hours, as the Cray XT4 and XT5 each were homogenous systems composed of multi-core processors. The use of node-hours, not core-hours, is necessary beyond 2011 due to the introduction of the Cray XK6. The current job scheduler for the OLCF compute resources is Adaptive Computing's Moab. Moab only allocates resources down to the granularity of a single node, regardless of the composition of that node. With the heterogeneous node types available on the Cray XK6, the use of just Opteron-core-hours will become insufficient to adequately describe resource utilization. The OLCF will work with the Program Manager to revise the definition for resource utilization measurement units beginning in 2012. These measurements will include segregation by node type as applicable.

2.4.1 Total System Utilization

2011 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. 2011 OLCF System Utilization

Program	Measurement Period	CPU Hours Allocated	CPU Hours Consumed	CPU Hours Available	% of Allocation Consumed
INCITE	CY2011	930,000,000	995,214,895		107.01%
ALCC	CY2011	393,730,910	207,434,655		52.68%
Director's Discretionary	CY2011	138,560,686	142,880,328		103.12%
Other	CY2011	215,415,138	122,423,846		56.83%
Total		1,682,706,734	1,467,953,724	1,683,360,779	87.24%
Cray XT4			39,079,672	43,070,274	90.73%
Cray XT5			1,428,874,052	1,640,290,505	87.11%

The OLCF tracks the consumption of core-hours by job. This allows the OLCF to track the consumption of core-hours, by Program, by Project, by User, and by System with high fidelity. Figure 2.2 describes the utilization by week and by Program for all of 2011. No adjustment is made to exclude any user group, including staff and vendors.

Figure 2.2. 2011 XT5 Resource Utilization – Core Hours by Program.

Assessing Total System Utilization

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

Note that non-renewed INCITE projects from 2010 were allowed by OLCF policy to continue running at low priority early in the first quarter of 2011 so that those projects could complete while 2011 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

2011 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 cores allocated to a particular job on the OLCF computing resources. In addition, this threshold is more stringent in the second and subsequent years of production. Leadership-class jobs must use more than 20% of the available cores of the largest system to qualify. Under this definition, the Cray XT4, available in the first two months of 2011, was too small to support capability computing.

The capability metric is defined by the number of CPU hours that are delivered by leadership-class jobs. For the initial year of production (2010), the metric stipulated that no less than 35% of the delivered CPU hours would reflect leadership-class jobs. For the second year of production (2011), the metric stipulates that no less than 40% of the delivered CPU hours reflect leadership-class jobs.

The OLCF Resource Utilization Council (RUC) used queue policy on the Cray XT5 in 2010 and 2011 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 XT5	20%	35% of delivered hours	30%	30% of delivered hours

The OLCF continues to meet – and 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 Joule metrics, where staff actively engage with code developers to promote application performance.

Table 2.12. OLCF Leadership Usage on Jaguar XT5 and XT4

	Leadership Usage	CY 2010 Target	CY 2010 Actual	CY 2011 Target	CY 2011 Actual
Cray XT5	INCITE	NAM	36.9%	NAM	47.8%
	Total	35.0%	39.0%	40.0%	54.0%
Cray XT4	INCITE	NAM	41.0%	NAM	39.1%
	Total	35.0%	34.7%	40.0%	57.1%

The average consumption of hours by leadership-class jobs was well above the CY 2011 Target of 40% at 54%. 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.

Strategic Results

HIGH PERFORMANCE COMPUTING FACILITY
2011 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2012

3. STRATEGIC RESULTS

CHARGE QUESTION 3: Is 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?”

OLCF RESPONSE: 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.

2011 Operational Assessment Guidance

In this section the Facility should report:

- *Science Output (publications);*
- *Scientific Accomplishments (true accomplishments, not just milestones); and*
- *Allocation of Facility Director’s Reserve Computer Time*

To be sure, the projects and user programs operating within the Oak Ridge Leadership Computing Facility (OLCF) advance the mission of the U.S. Department of Energy (DOE) 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 U.S. leadership in clean energy technologies.
- Goal 2: Maintain a vibrant U.S. effort in science and engineering as a cornerstone of our economic prosperity with clear leadership in strategic areas.

We will begin this discussion by focusing on Goal 2, within which DOE has a clear focus to be the leader in computational sciences and high-performance computing. Today this means that DOE will continue to advance the frontiers of energy-efficient computing and supercomputing to enable greater computational capacity with lower energy needs. OLCF embraces this goal and is focused upon it through the on time, on scope, and on budget delivery of Phase I of the OLCF-3 project (Reference Section 2). More specifically, OLCF has achieved and will continue to achieve DOE’s targeted

outcome to “Continue to develop and deploy HPC hardware and software systems through exascale platforms”.

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.

The OLCF currently follows the recommendation in the 2007 report of the ASCAC 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 (Joule metric). 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. The facility also collects quarterly reports from users, in which they are asked to provide updates on accomplishments and other activities, such as presentations given describing results of work under the allocation.

Table 3.1. List of OLCF Publications

	2011
Number of publications reported to the OLCF by users or identified by the OLCF	670
Number of refereed publications reportable within OAR guidance	300

3.2 Scientific Accomplishments

OLCF advances DOE’s science and engineering enterprise through robust partnerships with our users. Here we identify key accomplishments that advance the state-of-the-art in basic science and are advancing DOE’s science programs toward their targeted outcomes.

3.2.1 Subatomic Physics

DOE is the primary government sponsor of research in particle and nuclear physics. These fields advance knowledge at the extreme scales of energy and space to reveal the basic building blocks of our natural world. And these fields underlie technologies, e.g., based on nuclear isotopes, with demonstrated broad impacts on society and the economy. OLCF accomplishments described here represent insight into the fundamental nuclear forces and the structure of nuclei.

Accomplishment in Nuclear Physics: Nuclear Structure of C14 Requires Three-Body Forces

As reported in the May 20, 2011 edition of the journal *Physical Review Letters* by Maris et al., (*Phys. Rev. Lett.* **106**, 202502 (2011)), predicting nuclear structure and lifetimes using *ab initio* nuclear theory requires accounting for the complex nuclear interactions known as the three-body force. This accomplishment advances DOE toward its *targeted outcome in Subatomic Physics* to

- “test the theory of nuclear forces, and produce exotic nuclei of relevance in astrophysical processes”.

Nuclear theory to this point in time has assumed that the two-body force is sufficient to explain the structure and decay channels of a nucleus. In other words, the half-life or decay path of an unstable nucleus has been understood through the combined interactions of pairs of protons and neutrons. However, this project determined that the two-body force is not enough to describe the microscopic origins of the beta decay of ^{14}C to ^{14}N . Therefore, researchers must also tackle the far more

difficult challenge of calculating combinations of three nucleons at a time (three protons, three neutrons, or two of one and one of the other). This approach yields results that are both different from and more accurate than those of the two-body force, and are sufficient to describe the observed beta-decay phenomena. Two factors complicate the choice of approaches. First, two-body interactions do accurately describe some nuclei. Second, accurate calculations including three-body forces are very difficult and demand state-of-the-art supercomputers. Jaguar gave the team the computing muscle it needed to analyze the carbon-14 nucleus using the three-body force.

Carbon-14, with six protons and eight neutrons, is the isotope behind carbon dating, allowing researchers to determine the age of plant- or animal-based relics going back as far as 60,000 years. It was an ideal choice for this project because studies using only two-body forces dramatically underestimate the isotope's half-life, which is around 5,700 years.

"With Jaguar we are able to do ab initio calculations, using three-body forces, of the half-life for carbon-14," noted team member and OLCF computational physicist Hai Ah Nam said. "It's an observable that is sensitive to the three-body force. This is the first time that we've demonstrated at this large scale how the three-body force contributes."

This work is an accomplishment of the INCITE project entitled, "Nuclear Structure and Nuclear Reactions", led by PI James Vary, Iowa State University. This project received an allocation in 2011 at OLCF of 28 M hours. However, the project utilized 50 M hours, under the OLCF policy of allowing oversubscribed INCITE projects to continue to run at reduced priority. Hai Ah Nam contributed very directly by performing the code configuration on the Cray XT5 and all initial runs, as well as contributing the the analysis and communication of the results. In addition, OLCF uniquely enabled this achievement through Jaguar's very large memory. "Jaguar is the only system in the world with the capability to store the 240 terabytes of memory required for this calculation," Nam noted. "This is a huge, memory-intensive calculation."

3.2.2 Chemical and Materials Research

At the atomic and molecular scales, the DOE pursues world-class research in fundamental properties of materials and chemistry that explores the origins of macroscopic behaviors and their fundamental connections to atomic, molecular, and electronic structures. At its core is the quest for the deterministic design and discovery of new materials and chemical assemblies with novel structure, functions, and properties.

Accomplishment in Biochemistry: Researchers Show How Proteins Help DNA Replicate Past a Damaged Site

As reported in the October 17, 2011 issue of *Proceedings of the National Academy of Sciences*, a multi-institutional research team led by Ivaylo Ivanov of Georgia State University has employed the Jaguar supercomputer at Oak Ridge National Laboratory (ORNL) and x-rays a billion times brighter than the sun, produced at Lawrence Berkeley National Laboratory (LBNL), to illuminate how DNA replication continues past a damaged site so a lesion can be repaired later. The combination of computation and experiment reveals conformations that ubiquitin (Ub), a small protein that binds and orients DNA-editing enzymes, can assume when it associates with a molecular "tool belt" called proliferating cell nuclear antigen (PCNA). The combination jointly advanced DOE's targeted science outcomes to

- “explore the use of synchrotron light sources”, and
- “continue to develop and deploy high-performance computing hardware and software systems...”.

“The tool belt model is a longstanding model in the PCNA field, although it has not been conclusively proven,” said Ivanov. “There are three binding sites on PCNA, which is a ring made of three identical subunits,” Ivanov explained. “The model proposes that each site can bind a different DNA-editing enzyme.” One such enzyme is DNA polymerase, which catalyzes the formation of a new DNA strand from an existing template. It also plays a major role in repair by associating with the tool belt and iteratively adding one of four bases to a damaged strand. For the tool-belt model to be feasible, replicative and translesion polymerases need distinct binding sites. “Multiple polymerases—say, a high-fidelity replicase and a low-fidelity lesion bypass polymerase—could bind PCNA without steric clashes, forming a complex,” Ivanov said. “The model envisions that depending on the DNA context encountered by PCNA as it moves along DNA, it would swap the two polymerases and use the right tool from this set.”

The researchers used computing to winnow the number of places where Ub binds to PCNA from thousands of prospects to three. “Each stage of the refinement involved a different method with different timescale,” Ivanov explained. An application called Tethered Brownian Dynamics (TBD) explored the complex on the longest accessible timescale. It treats Ub and PCNA as rigid bodies as it quickly explores possible orientations in terms of electrostatic and geometric complementarity. With the electrostatic potential precomputed on a grid, interactions are approximate. RosettaDock, in contrast, uses a more sophisticated energy function, making it more accurate than TBD, and allows for side chain flexibility. It investigates protein–protein docking but keeps the backbone of both Ub and PCNA fixed to limit the expense of the conformational search. At the last stage, the researchers used the NAMD code to model the molecular dynamics of the entire system, including the solvent, in full flexibility. Clustering analysis bridged different stages in the multiscale modeling protocol.

“The job of the simulations is to provide low-energy, structurally compatible models that can then be vetted using the experimental data,” Ivanov said. Using LBNL’s Advanced Light Source (ALS), experimentalists obtained structures of Ub-PCNA complexes in solution using a technique called small-angle X-ray scattering (SAXS). Combining multiscale computational modeling and SAXS showed that Ub didn’t have to reside just at the site on PCNA identified by the crystal structure; it could exist at two other sites that were computationally derived. The researchers called the positions “flexible,” “back,” or “side.” While the latter two indicate positions relative to PCNA, the flexible position is generated by high-temperature molecular dynamics, so the Ub becomes somewhat removed from the PCNA surface.

When PCNA traveling down a double strand of DNA bases encounters a stalled replication fork—a place where replication has stopped due to DNA damage—its first response is to attach a Ub molecule. Enzymes called translesion polymerases then attach to PCNA via Ub and replicate bases through the DNA lesion. That process is called translesion synthesis (TLS). Finding new docking sites on the PCNA and equilibrium among the three Ub positions provides unexpected insight into previously unexplained biological observations. “The computationally identified positions explain the influence of mutations identified previously in genetic screens,” Ivanov said. “These are known to interfere with translesion synthesis but not with normal DNA replication.”

Like a construction worker swaps a hammer with a screwdriver on a tool belt to make a needed instrument more accessible, Ub can swap enzymes bound to PCNA when different ones are needed for TLS. “The dynamic range of positions offers an explanation of how the TLS polymerase could initially bind to the back side of PCNA and then transition to the polymerase binding face of PCNA,” Ivanov said. The researchers received an INCITE allocation of 4 million processor hours in 2010 and 2011, which they ran on ORNL’s Jaguar supercomputer.

Accomplishment in Materials Modeling: Boron Nitride Nanoribbons as Graphene Substrate

As reported in the July 7th, 2011 edition of the journal *Nano Letters*, by Lopez-Bezanilla et al., (*Nano Letters*, **11**(8), 3267 (2011)), predicting that boron-nitride monolayers are an ideal dielectric substrate material for future nanoelectronic devices constructed with graphene as the active layer. This accomplishment advances DOE toward its *targeted outcome in Chemical and Materials Research* to

- “[d]evelop and explore a broad spectrum of new materials that have novel properties [...] or otherwise contribute to the advancement of energy technologies by 2020”.

Graphene, which is carbon in the form of freestanding 1-atom-thick sheets, is a natural for next-generation computer chips, communications equipment, and solar energy devices. To live up to its potential, however, graphene needs an appropriate substrate. On its own, its edges wrinkle, tear, or roll up. The silicon dioxide substrate used for today’s microchips is not a good partner for graphene; it creates vibrations that slow the electrons, and its surface is too bumpy. An ideal substrate would not physically interfere with the graphene. “The substrate has to be a dielectric material, a material that is insulating and that can be polarized by an applied electrical field,” says Lopez-Bezanilla, a research associate in ORNL’s Computing and Computational Sciences Directorate. He is funded by the Petascale Initiatives program of the Department of Energy’s Office of Advanced Scientific Computing Research.

Any substrate will affect the electrons in the adjacent graphene layer, but boron nitride interferes less than silicon dioxide. Also, boron nitride resists chemical change and is unaffected by high temperatures. “Boron nitride is a covalent material with atoms tightly bonded to each other, but it also presents a strong ionic behavior,” explains Lopez-Bezanilla, making it a great insulator and poor conductor. Boron nitride, like graphene, can be formed as monolayers and nanoribbons with their atoms arranged in a hexagonal lattice (imagine chicken wire). And graphene and boron-nitride have well matched unit cells.

The team used two systems for the simulation. Jaguar ran the Vienna Ab-initio Simulation Package (VASP) and the Oak Ridge National Laboratory Institutional Cluster ran the SIESTA code. Both are density functional theory (DFT)-based codes.

This work is an accomplishment of the INCITE project entitled, “Petascale Modeling of Chemical Catalysis and Interfaces”, led by PI Bobby Sumpter, Oak Ridge National Laboratory. This project received an allocation in 2011 at OLCF of 75 M hours. The project utilized 76.8 M hours. The lead author of this paper, Alejandro Lopez-Bezanilla’s is funded by ASCR’s postdoc program (i.e., Petascale Initiatives Program). The achievement leverages the strong, sustained partnership between OLCF and the Center for Nanophase Materials Sciences (CNMS) and Computer Science and Mathematics Divisions (CSMD). In particular Paul Kent (CSMD/CNMS) has made improvements to VASP to increase scalability. However, scaling VASP to leadership computing performance remains a challenge.

Figure 3.1. Atomic-detailed model of plant components lignin and cellulose. The leadership-class molecular dynamics simulation investigated lignin precipitation on the cellulose fibrils, a process that poses a significant obstacle to economically-viable bioethanol production.

3.2.3 Bioenergy

As DOE continues to develop biotechnology solutions for energy, with particular emphasis on cost-effective technologies for next-generation production of biofuels, leading scientists come to OLCF to access our unique computational capabilities. The OLCF accomplishment described here represents a cutting-edge approach to application of world-class science facilities to advance research into cellulosic biofuels.

Accomplishment in Biofuels: Self-Similar Multiscale Structure of Lignin Revealed by Neutron Scattering and Molecular Dynamics Simulation

As reported in the April 14th, 2011 edition of the *Journal of Applied Physics*, by Petridis et al., (*J. Appl. Phys* **109**, 07A942 (2011)), a team led by Jeremy Smith from the University of Tennessee and ORNL has taken a substantial step in the quest for cheaper biofuels by revealing the surface structure of lignin clumps down to 1 angstrom (equal to a 10 billionth of a meter). The team employed two of ORNL’s signature strength – simulation on Jaguar and neutron scattering – to reach its conclusion that the surface of these clumps is rough and folded, even magnified to the scale of individual molecules. These results are important because lignin is a major impediment to the production of cellulosic ethanol, preventing enzymes from breaking down cellulose molecules into the sugars that will eventually be fermented. This accomplishment advances DOE toward its *targeted outcome in Bioenergy* to

- “[a]pply systems biology approaches by 2015 to create viable biofuels processes...”.

Lignin itself is a very large, very complex molecule made up of hydrogen, oxygen, and carbon (Figure 3.1). In the wild its ability to protect cellulose from attack helps hardy plants such as switchgrass live in a wide range of environments. When these plants are used in biofuels, however, lignin is so effective that even expensive

pretreatments fail to neutralize it. “Nature has evolved a very sophisticated mechanism to protect plants against enzymatic attack,” explained team member Loukas Petridis, a computational physicist at ORNL, “so it is not easy to make the fuels. What we’re trying to do is understand the physical basis of biomass recalcitrance—resistance of the plants against enzymatic degradation.”

The team used neutron scattering with ORNL’s High Flux Isotope Reactor to resolve the lignin structure from 1,000 down to 10 angstroms. A molecular dynamics application called NAMD used Jaguar to resolve the structure from 100 angstroms down to 1. The overlap from 10 to 100 angstroms allowed the team to validate results between methods. The two methods—neutrons and supercomputing—confirmed that the surface of lignin aggregates is highly folded, with a surface fractal dimension of about 2.6. The surface fractal dimension is a measure of the roughness or irregularity of a surface and ranges from 2 (very smooth) to 3 (very folded). The value of 2.6, indicates a highly folded, rough surface. This roughness explains enzymes captured by the lignin, a process that limits viable biofuel production.

Smith’s project is the first to apply both molecular dynamics supercomputer simulations and neutron scattering to the structure of biomass, and the two methods reinforce one another very well. On the one hand, neutron scattering could not reveal the structure at the smallest scales. On the other hand, simulation could not cover the full range of scales even on Jaguar, the United States’ most powerful supercomputer. “When you look at the combination of neutrons and simulation, first of all it has not been done on lignins before. The combination of techniques gives you a multiscale picture of lignin. Neutron scattering can probe length scales, for example, from 10 angstroms all the way up to 1,000 angstroms. “On the other hand, molecular dynamics simulations can go to smaller length scales—from 1 angstrom or even sub angstrom all the way to 10 or even 100 angstroms. This is why we have been able to study the structure of the lignin droplets over various length scales. Not only was the finding new, but these techniques were used for the first time to study lignocellulose.” The performance bottleneck is the calculation of long-range electrostatic forces. Employing reaction field electrostatics yields strong improvement in parallel efficiency, enabling the largest-ever biological MD simulation.

This work is an accomplishment of the INCITE project entitled, “Cellulosic Ethanol: Simulation of Multicomponent Biomass Systems”, led by PI Jeremy Smith, University of Tennessee and ORNL. This project received an allocation in 2011 at OLCF of 30 M hours. The project utilized 25 M hours. . Mike Matheson performed visualization of simulation results.

3.3 Accomplishments in Energy Systems Research

Goal 1 for DOE is to “Catalyze the timely, material, and efficient transformation of the nation’s energy system and secure U.S. leadership in clean energy technologies. The OLCF has powerful computational capabilities that are being applied to meet this goal. In the following, we present accomplishments from two projects within our industrial partnerships program that demonstrate the impact we are having on DOE’s targeted outcomes.

3.3.1 Energy Systems Simulation

Accomplishment in Energy Systems Simulation: Turbomachinery simulations show potential for efficient CO₂ compression for cheaper carbon sequestration

As reported in the Proceedings of the 4th European Conference for Aerospace Sciences (EUCASS), July 4-8, 2011 in Saint Petersburg, Russia, a team lead by Allan

Grosvenor of Ramgen Power Systems, Bellevue WA, performed 3D aerodynamics simulations of transonic and supersonic flow, where shock wave, boundary layer interaction plays a dominate role in understanding separating flow phenomena. Ramgen is applying the principles and techniques advanced in these studies, originally developed for supersonic aircraft technology, to produce new designs for turbo compressors for CO₂ with much higher efficiency than today's state-of-the-art compressor technologies. This accomplishment advances DOE toward multiple *targeted outcomes for strategic Goal 1, e.g.,*

- *“Reduce upfront risk and cost associated with geologic technologies, including carbon sequestration and geothermal energy systems”, and*
- *“Bring at least 5 commercial-scale carbon capture and storage (CCS) demonstrations online by 2016”.*

One of the most pressing scientific challenges facing the US and many other countries is the removal of greenhouse gases from the atmosphere. The Department of Energy (DOE) is currently sponsoring large-scale demonstration projects to prove the viability of capturing carbon and storing it underground (also called “sequestration”). But it turns out that a principle barrier to widespread application of the carbon capture and sequestration (CCS) process is the cost of large-scale compression of the CO₂ for storage, a process that represents approximately 33 percent of the total cost of CCS. Ramgen Power Systems, a small, Seattle-based engineering firm, is developing shock wave compression turbo machinery as a means to meet DOE goals for reducing carbon capture and sequestration costs, with a complementary goal to design gas turbines with dramatically lower costs and higher electricity-generation efficiency. This activity was brought to the OLCF from the highest levels of the Department of Energy.

Mike Matheson of OLCF Scientific Computing Group collaborated and supported Ramgen, including the process of taking their full-turbo-design workflow to a new paradigm on the Jaguar supercomputer. Mike worked with Ramgen personnel to port their application to Jaguar and gather the needed performance metrics. Together, they created a plan to improve performance and scalability, resulting in a 50-fold improvement in code scalability via more efficient memory utilization and a 10-fold improvement in I/O. The Cray XK6 upgrade late in 2011 doubled the node, significantly helping this project. This led to a revolutionary end-to-end workflow strategy that fully exploits Jaguar's leadership computing capabilities, allowing the intelligent use of ensembles to efficiently explore design parameter space using over 120,000 Jaguar cores. Workflows previously requiring months of work are now completed within 8 hours.

The end result is that Ramgen significantly advanced its shock wave based compression aerodynamic design process and revealed designs that exhibit valuable new aerodynamic characteristics. Based on recent results, Ramgen is already machining titanium to build a prototype turbocompressor. This was an outstanding effort that reflects well on the unique capabilities of our leadership computing facilities and on the truly unique intellectual capabilities of our computational science staff.

“Jaguar makes it possible to solve aerodynamics design problems that in the past would have taken prohibitively long periods of time,” said Ramgen's CEO Doug Jewett. “It simply would not have been possible without Jaguar and the assistance of Mike Matheson and the OLCF.”

Ramgen received an allocation on Jaguar, first through the Director's Discretionary Program, then later through the 2011 ASCR Leadership Computing Challenge (ALCC) project entitled, "High-Resolution Design-Cycle Analysis, Supporting CO2 Compression, Technology Development", led by PI Allen Grosvenor, Ramgen Power System. This project received an ALCC allocation in 2011 at OLCF of 6 M hours. The project utilized 5 M hours in CY 2011.

Accomplishment in Energy Systems Simulation: Comparison of Unsteady Flow to Traditional Assumptions of Steady Flow in Turbomachinery Helps Improve Efficiency and Reduce Noise.

As reported in the SC11 Masterworks Symposium (SC11), November 12-18, 2011 in Seattle, WA, a team from General Electric (GE) is looking to Jaguar to simulate the design of next-generation turbomachinery. Using their flagship code known as Tacoma, GE engineers ran their largest ever computational fluid dynamics simulation on Jaguar, an achievement that helps pave the way for the design and production of next-generation turbines and uniquely positions GE in the international turbomachinery marketplace. This accomplishment advances DOE toward its *targeted outcome for strategic Goal 1, e.g.,*

- *"Validate high-fidelity simulations of internal combustion engines, fission and conventional power plants in commercial by 2015, thereby integrating HPC simulation into the industrial energy sector".*

General Electric (GE) has been building turbomachines for nearly a century and is currently a major producer for the electric power generation and aircraft engine industries. Recently, however, GE took its turbomachinery R&D to the fast lane with the help of one of the fastest computers in the world.

"It's a very competitive business," said Principal Engineer Graham Holmes of General Electric (GE) Global Research. "If you could achieve a one percent increase in efficiency for a turbomachine, the market is yours." That one percent fuel-burning advantage would, over time, add up to enormous energy and cost savings for GE's customers and provide GE with a "business critical" advantage, according to Holmes. Through ORNL's HPC Industrial Partnerships Program, GE recently harnessed Jaguar's power to study in greater detail than ever before the unsteady fluid flows in turbomachines. Understanding these flows is essential in order to achieve greater efficiency, and for GE to gain an edge in an intensely competitive global marketplace.

The basic physics of turbomachinery operations have been well understood for years—jet engines and gas turbines go back to the mid-20th century. Essentially, turbomachines feature alternating rows of stationary and moving blades either expanding or compressing gas. The design process has evolved from experimentation and highly simplified analytical models to increasingly sophisticated simulations, carried out on increasingly powerful computers. Engineers typically shape blades, run a combination of simulations and experiments, tweak the design, and repeat; an expensive path to production by any measure. In order to make simulations practical within a reasonable wall clock time, turbomachinery designers have traditionally assumed that the velocity of air around and across the blades remains steady in the reference frame of the blades. However, designers have always understood that this airflow is unsteady and that the flow has to be unsteady for a turbomachine to work, i.e., perform work and transfer momentum from stator to rotor, and/or visa versa. But the assumption that the flow, as seen by each blade row, can be approximated as steady has proven to be remarkably powerful. The designs of all the most efficient

turbomachines, such as the turbines that drive the large diameter fans in modern jet engines, have been created using this paradigm.

Any future efficiency improvements will likely depend on understanding the unsteady nature of the fluid flow. For example, in jet engines, low-pressure turbines drive increasingly larger diameter fans, rotating at a lower RPM in order to increase efficiency. But running low pressure turbines at lower speeds presents a severe technical challenge—the lower the relative velocity between the blade rows, the harder it is to extract energy to drive the fan, requiring more rows of blades. Unfortunately, as more rows of blades are added, the turbine becomes heavier, and therefore less efficient. Understanding the unsteady flows should allow designers to make needed adjustments without adding to the weight of the turbine, resulting in overall greater fuel efficiency. Unsteady flow analysis is also essential in understanding other phenomenon like blade flutter, or the blade vibration induced by the fluid flow. Blade flutter can be catastrophic in an aircraft engine if it results in the damage of one or more blades and turbine failure.

In order to simulate unsteady flow, GE used their flagship code for turbomachinery, TACOMA. When paired with Jaguar, GE researchers ran their largest computational fluid dynamics calculation to date and were able to investigate for the first time the unsteady flows in turbomachinery. Simulations were ramped up from 3-D to 4-D, and researchers were able to look at the time-resolved unsteady flows in the moving blades. The team was able to examine a turbine test rig and compare steady and unsteady flows. In the two analyses, the efficiency remained the same, which is “an extremely valuable piece of information,” said Holmes. Furthermore, the team found plenty of interesting unsteady phenomena occurring throughout the device. For instance, the interactions between the blade and the hub created unsteady secondary flows, which behave differently than those witnessed in the steady analysis.

Overall, said Holmes and Moore, GE’s simulations on Jaguar advanced the company’s R&D in the turbomachinery arena and are providing it with a distinct competitive advantage as the company pours over the simulation results. The team believes that GE and its competitors will move further into unsteady flow analysis to achieve the final point in efficiency, a move that will require substantial HPC resources. In fact, largely as a result of these calculations, GE recently purchased its own Cray system, a move that significantly ramps up its in-house HPC capability.

GE received an allocation on Jaguar through the Director’s Discretionary Program, project entitled, “Unsteady Performance Predictions for Low Pressure Turbines”, led by PI Brandon Moore, GE Global Research. This project received a Director’s Discretionary allocation in 2010 at OLCF of 2 M hours. The project carried over time through this year and utilized 552 K hours in CY 2011.

In concluding this section on accomplishments, we summarize the targeted outcomes in science and energy systems research within DOE’s 2011 Strategic Plan advanced by the six significant accomplishments selected for inclusion in this present Operational Assessment Report.

OLCF Advances DOE’s Targeted Outcomes in Science and Energy Systems Research

Targeted Outcomes	OLCF Achievement
“Continue to develop and deploy HPC hardware and software systems through exascale platforms.”	Delivered Phase I of OLCF-3 Project on time, on scope, and on budget.
“Test the theory of nuclear forces, and produce exotic nuclei of relevance in astrophysical processes”	Nuclear structure and decay of ^{14}C requires three-body forces.

“Develop and explore a broad spectrum of new materials that have novel properties, and otherwise contribute to the advancement of energy technologies by 2020.”	Boron-nitride (BN) monolayers are ideal dielectric substrate materials for future, graphene-based nanodevices.
“Apply systems biology approaches by 2015 to create viable biofuels processes.”	Lignin molecular structure is revealed by neutron scattering and molecular dynamics simulation. Surface roughness explains enzymes being captured by the lignin, a process that limits viable biofuel production.
“Reduce upfront risk and cost associated with geologic technologies, including carbon sequestration”, and “Bring at least 5 commercial-scale carbon capture and storage (CCS) demonstrations online by 2016.”	Invent and implement a new paradigm for exploring shock compression turbo machinery and design of turbo compressors to reduce the cost of CCS and meet DOE’s aggressive goals.
“Validate high-fidelity simulation of internal combustion engines, fission, and conventional power plants in commercial use by 2015, thereby integrating high-performance computer simulations into the industrial energy sector.”	Comparisons of unsteady flow to traditional steady-flow assumptions in turbo machinery help improve efficiency and reduce noise. General Electric, Global Research, on the basis of the success of this project, purchased their own Cray supercomputer for in house use.

3.4 Allocation of Facility Director’s Reserve

2011 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 Director’s Discretionary (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 through the Advanced Scientific Computing Research Leadership Computing Challenge (ALCC) program.

3.4.1 Director’s Discretionary Program

The goals of the Director’s Discretionary (DD) program are threefold: development of strategic partnerships, leadership computing preparation, and application performance development and measurement. These goals are aligned with particular strategic goals for the OLCF, namely the expansion of the steady-state 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. These may be entirely new areas with respect to HPC, or areas in need of nurturing. Example candidate projects are those associated with the ORNL Laboratory Directed Research and Development Program, programmatic science areas (bioenergy, nanoscience, climate, energy storage, engineering science), and key academic partnerships (e.g., that with the ORNL Joint Institute for Computational Sciences). Included in this broad category is the Industrial 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. In addition, infrastructure software such as frameworks, libraries, and application tools, and support research areas for next-generation OSs, performance tools, and debugging environments are often developed in DD projects.

The Resource Utilization Council (RUC) makes the final decision on DD applications, using written reviews from subject matter experts. The actual DD project lifetime is specified upon award: most allocations are for less than 1 year. The typical size of DD awards is roughly 1M core-hours, but can range from tens of thousands of hours to 4 million hours or more.

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%

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). The complete list of current Director’s Discretionary projects is at the end of this section.

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

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

Annual DD allocations are typically less than the available hours. 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 utilized with 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.

3.4.2 Industrial HPC Partnerships Program

The Industrial HPC Partnerships Program is gaining traction and attracting both large and small firms. They are applying for time on Jaguar through every pathway available; INCITE, ALCC and the Director's Discretionary Program. (Table 3.4 lists projects active in CY 2011.) And although the program is only three years old, it is helping expand and accelerate U.S. industrial use of HPC for national competitive gain. For example, United Technologies Research Center began using Jaguar through smaller DD allocations and in 2011 was able to compete successfully for a larger ALCC allocation to tackle more complex problems. And largely as a result of the new insights GE gained from using Jaguar, the firm purchased its own Cray system, substantially increasing its in-house computing capabilities. This is a significant testament to how the Industrial HPC Partnerships Program is helping companies build an internal return-on-investment case for greater use of high performance computing.

Table 3.4. Industry Projects at the OLCF

Corporate Partner	Program	Description
Procter & Gamble	INCITE	Coarse Grained Molecular Dynamics Studies of Vesicle Formation and Fusion
Boeing	ALCC	Reliable Predication of Performance of High Lift Systems of Commercial Air
Ramgen	ALCC	High resolution design-cycle computational fluid dynamics analysis supporting CO2 compression technology development
GE Global Research	ALCC	High fidelity simulations of gas turbine combustors for low emissions engines
GE Global Research	ALCC	Non-icing Surfaces for Cold Climate Wind Turbines
United Technologies Research Center	ALCC	Large -eddy simulation for turbomachinery - advancing state-of-the-art
BMI/Smart Truck Corporation	DD	Class 8 long-haul truck optimization for greater fuel efficiency
GE Global Research	DD	Unsteady Performance Predictions for Low Pressure Turbines
United Technologies Research Center	DD	Nanostructured catalyst for water-gas shift and biomass reforming hydrogen production
United Technologies Research Center/Pratt & Whitney	DD	Multiphase injection for jet engine combustors
GE Global Research	DD	Investigation of Newtonian and non-Newtonian Air-Blast Atomization Using OpenFoam
United Technologies Research Center	DD	Surface Tension Predictions for fire-fighting foams
Ford	DD	Large Scale Engine Bay Package Optimization

Many of the industry projects in 2011 complemented DOE's strategic focus on addressing the nation's energy challenges. The cost and availability of energy, coupled with heightened environmental concerns, are causing companies to reexamine the design of products from large jet engines and industrial turbines to automotive engines. Their customers and the country are demanding products that have lower energy requirements and reduced environmental impact. However, the complexity of these design and analysis problems, coupled with the need for nearer term results, often requires access to computing capabilities that are far more advanced than those available in corporate computing centers. The OLCF is helping to address this gap by providing access to leadership systems and experts not available within the private sector.

For example, GE and United Technologies Research Center (UTRC) are both using Jaguar to tackle different problems related to jet engine efficiency. The impact of even a small change is enormous. A 1% reduction in specific fuel consumption can save \$20B over the life of a fleet of airplanes (20,000 engines $\times$ 20-year life).

Access to Jaguar allowed GE for the first time to study unsteady flows in the blade rows of turbomachines, such as the large diameter fans used in modern jet engines. Unsteady simulations are orders of magnitude more complex than simulations of steady flows, and GE was not able to attempt this on its in-house systems. GE engineers also ran their largest ever computational fluid dynamics simulation on Jaguar.

UTRC is using Jaguar to better understand the air-fuel interaction in combustors, a critical component of aircraft engines. They are validating first principles methods against experimental measurements, a first in this field given the complexity of the problem. Better understanding of the air-fuel interaction will enable UTRC to develop more efficient combustors that will reduce the emissions, lower the noise, and enhance the fuel efficiency of aircraft engines.

Access to Jaguar and OLCF experts is helping industry accelerate time-to-insight and time-to-solution for important energy-related problems with national impact. As industry delivers more energy efficient products, ORNL and DOE are delivering an additional return on the nation's investment in the OLCF.

3.5 Management of INCITE Projects

INCITE awards represent 60% of the allocable time on the OLCF's production system. Once INCITE allocations are awarded, the OLCF center takes on the responsibilities of ensuring projects are enabled and granting project members access, supporting users on INCITE projects by answering questions and resolving problems, tracking usage and capability metrics, and capturing and recording achievements.

The OLCF places a strong emphasis on ensuring that INCITE projects are enabled and ready for use at the beginning of the allocation period. In order to enable a project, certain requirements have to be satisfied like user agreements, PI agreements, Export Control rulings, identify proofing, etc. The OLCF has investigated several barriers that keep the projects from getting access on day one and have worked over the last few years to improve these processes. One of the barriers identified was the time it takes to do an Export Control ruling. After working with the ORNL Export Control office, it was determined that renewing projects could undergo a less intensive review thereby freeing up time to work on the rulings for the new projects. After working to improve these processes, the percentage of projects with access by the beginning of the allocation period improved from 68% in 2010 to 95% in 2012.

Usage and capability metrics are monitored through various reports and are reviewed routinely by center staff and the Resource Utilization Council (RUC). Additionally, data is reviewed from quarterly reports and survey to ensure that queue policies and other system parameters are optimized for capability computing.

Another important process is monitoring scientific accomplishments. The accomplishments are monitored through quarterly reports and regular interactions between the PI and center staff. The PIs are required to report quarterly and provide feedback on project milestones, accomplishments, and their productivity information

including publications. The information is then used in highlights for DOE and other media groups, reports to DOE, fact sheets and presentations, and in special reports.

There are three distinct ways that Scientific Computing liaisons can and do work to ensure INCITE project successes. The first role is that of trouble shooter and problem solver for specific, short-term challenges identified within the project. The second involves making major improvements in the code that require significant rewriting or refactoring of sections of code. The third, and most intimate relationship with projects, is the case in which the liaison is an indispensable collaborator -- possibly the lead PI or lead code architect -- for a project. In all cases, much of the credit for OLCF's scientific success in timely and effective project execution resides with the liaisons in the Scientific Computing Group.

The role of the Director of Science is comprehensive and pervasive in managing INCITE projects. Indeed, the responsibilities of the Director of Science begin before awards are made, as the DoS is a member of the INCITE Program Awards Committee. Throughout the extent of each project, the DoS is the person at the OLCF who is responsible for collecting and reporting the scientific results and achievements of the OLCF as a whole. The Director of Science also develops direct relationships with INCITE PI's and prospective PI's, ensuring that each project has the best possible circumstances and collaborations for success. The Director of Science serves as chair of the RUC, the primary mechanism for OLCF policy establishment and management. Enabling INCITE project success is one of the primary missions of the RUC, and the specific attention given to individual project needs while ensuring access for all users is the heart of its work.

2011 Director's Discretionary Allocations

PI	Affiliation	2010 Allocation	Carryover to 2011	New 2011 Allocation	Total 2011 Allocation	2011 Usage	Project Name
Mike Henderson	BMI Corporation	4,000,000	2,695,917		2,695,917	266,618	Smart Truck Optimization
Pablo Carrica	University of Iowa	750,000	20,167		20,167	0	Large-scale Computations of Wind Turbines using CFDShip-Iowa Including Fluid-Structure Interaction
Branden Moore	GE Global Research	2,000,000	172,836		172,836	640	Unsteady Performance Predictions for Low Pressure Turbines
Thomas Gielda	Caitin Inc.			500,000	500,000	69,106	Parallel Computing performance Optimization for Complex Multiphase Flows in Strong Thermodynamic Non-equilibrium
Alexander Akkerman	Ford Motor Company			1,000,000	1,000,000	230	Large Scale Engine Bay Package Optimization
Rainald Lohner	George Mason University			1,000,000	1,000,000	183,875	Highly Detailed Simulations of Blasts on Offshore Platforms
Paul Ricker	UIUC	3,150,000	2,000,000		2,000,000	2,335,283	Testing Active Galaxies as a Magnetic Field Source in Clusters of Galaxies
Salman Habib	LANL	1,000,000	999,735		999,735	0	Dark Universe
Patrick Fragile	ORAU	1,000,000	1,000,000		1,000,000	18,002	Radiation Transport in Numerical Simulations of Black-Hole Accretion Disks
Paul Sutter	University of Illinois			5,000,000	5,000,000	1,795,343	Exploring the origins of galaxy cluster magnetic fields
Stephen Nesbitt	UIUC	165,000	115,797		115,797	47,019	Dynamically Downscaling the North American Monsoon Using the Weather Research and Forecasting Model with the Climate Extension (CWRf)
Jason Hill	University of Minnesota	900,000	900,000		900,000	1,613,618	Air Pollution Impacts of Conventional and Alternative Fuels
Pratul Agarwal	ORNL	4,000,000	0		0	3,883,851	High Throughput Computational Screening Approach for Systems Medicine

2011 Director's Discretionary Allocations

PI	Affiliation	2010 Allocation	Carryover to 2011	New 2011 Allocation	Total 2011 Allocation	2011 Usage	Project Name
Lei Shi	Cornell University	1,000,000	999,980		999,980	1,102,929	Transport Mechanism of Neurotransmitter: Sodium Symporter
Jerome Baudry	ORNL			6,000,000	6,000,000	2,859,545	High Performance Computing for Rational Drug Discovery and Design
Xiaolin Cheng	ORNL			500,000	500,000	6,286	Scalable bio-electrostatic calculation on emerging computer architectures
Thomas Miller	California Institute of Technology	2,000,000	10,104		10,104	0	Proton Coupled Electron Transfer Dynamics in Complex Systems
Jean-Luc Bredas	Georgia Institute of Technology	1,000,000	1,000,000		1,000,000	475,035	Electronic and Geometric Structure of Inorganic/Organic and Organic/Organic interfaces Relevant in Organic Electronics
Amra Peles	United Technologies Research Center	100,000	7,979		7,979	8,286	Nanostructured Catalyst for WGS and Biomass Reforming Hydrogen Production
Ilian Todorov	STFC Daresbury Lab	500,000	440,888		440,888	298,166	An Investigation of the Channel-Opening Movements of the Nicotinic Acetylcholine Receptor
Erik Deumens	University of Florida	1,000,000	777,712	7,222,288	8,000,000	6,709,434	EOM-CC calculations on diamond nano crystals
Zhengyu Liu	University of Wisconsin Madison			2,000,000	2,000,000	2,517,586	Assessing Transient Global Climate Response using the NCAR-CCSM3: Climate Sensitivity and Abrupt Climate Change
Kate Evans	ORNL	5,000,000	0		0	317,687	Decadal Prediction of the Earth System after Major Volcanic Eruptions
Gil Compo	University of Colorado	3,000,000	2,769,235		2,769,235	463,861	Developing a High Resolution Reanalysis Dataset for Climate Applications (1850 to present)
Leslie Hart	NOAA-ESRL	50,000	50,000		50,000	0	NOAA Benchmark Portability
Moetasim Ashfaq	UT-Knoxville	1,000,000	993,364		993,364	26,822	Quantification of Uncertainties in Projections of Future Climate Change and Impact Assessments
David Erickson	ORNL	500,000	172,260		172,260	24,497	WRF Downscaling

2011 Director's Discretionary Allocations

PI	Affiliation	2010 Allocation	Carryover to 2011	New 2011 Allocation	Total 2011 Allocation	2011 Usage	Project Name
Atul Jain	University of Illinois			30,000	30,000	1,553	Land Cover and Land Use Change and its Effects on Carbon Dynamics in Monsoon Asian Region
James Joseph Hack	ORNL			15,000,000	15,000,000	4,900,556	Ultra High Resolution Global Climate Simulation to Explore and Quantify Predictive Skill for Climate Means, Variability and Extremes
Aytekin Gel	ALPEMI Consulting			600,000	600,000	689,737	Mitigation of CO2 Environmental Impact Using a Multiscale Modeling Approach
Don Lucas	LLNL			100,000	100,000	1,662	Uncertainty Quantification of Climate Sensitivity
Katherine Evans	ORNL			1,000,000	1,000,000	282,812	A Scalable, Efficient, and Accurate Community Ice Sheet Model (SEACISM)
Dale I Pullin	California Institute of Technology	500,000	194,776		194,776	1,062,723	Direct Numerical Simulation of the Mach Reflection Phenomenon and Diffusive Mixing in Gaseous Detonations
Alexei Khokhlov	University of Chicago	600,000	600,000		600,000	0	First-principles Petascale Simulations for Predicting DDT in H2-O2 Mixtures
Marios Soteriou	United Technologies Research Center			2,500,000	2,500,000	1,422,257	Multiphase Injection
Gregory Laskowski	GE Global Research	1,000,000	890,854		890,854	818,233	Investigation of Newtonian and non-Newtonian Air-Blast Atomization Using OpenFoam
Suresh Menon	Georgia Institute of Technology			1,000,000	1,000,000	75,826	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	1,000,000	32,671	Next Generation Turbulent Reactive Flow Simulation
Samuel Paolucci	University of Notre Dame			1,000,000	1,000,000	0	Reactive flows with detailed chemistry using an adaptive multiscale wavelet method

2011 Director's Discretionary Allocations

PI	Affiliation	2010 Allocation	Carryover to 2011	New 2011 Allocation	Total 2011 Allocation	2011 Usage	Project Name
David Bowler	University College London	2,650,000	2,321,114		2,321,114	0	Modeling of Large-Scale Nanostructures using Linear-Scaling DFT
Tommaso Roscilde	Ecole Normale Supérieure de Lyon	800,000	0		0	12,163,594	Emulating the Physics of Disordered Bosons with Quantum Magnets
Thomas Maier	ORNL			10,000,000	10,000,000	5,511,931	Predictive simulations of cuprate superconductors
Gabriel Kotliar	Rutgers University			1,000,000	1,000,000	236,001	Calculation of Strongly Correlated Systems Using GW+DMFT Method
Sean Ahern	ORNL	8,000,000	1,516,488		1,516,488	14,022	SciDAC 2 Visualization Center and Institute
Kalyan Perumalla	ORNL	2,000,000	1,999,980		1,999,980	1,955,081	An Evolutionary Approach to Porting Applications to Petascale Platforms
George I-Pan Fann	ORNL	1,000,000	0		0	451,274	Prototype Advanced Algorithms on Petascale Computes for IAA II
Stephen Poole	ORNL	300,000	0		0	1,867,803	FASTOS Community Allocation
Zizhong Chen	Colorado School of Mines	1,000,000	0		0	977,347	Fault Tolerant Linear Algebra Algorithms and Software for Extreme Scale Computing
Robert Patton	ORNL	1,000,000	934,680		934,680	0	High Performance Text Mining
Phil Colella	LBNL	2,500,000	228,877		228,877	152,906	Applied Partial Differential Equations Center. APDEC.
Kalyan Kumaran	ANL	1,000,000	1,000,000		1,000,000	1	Performance Measurements Using ALCF Benchmarks
Sean Ahern	ORNL			1,000,000	1,000,000	0	Large-Scale Data Analysis and Visualization
Terry Jones	ORNL			3,000,000	3,000,000	2,637,455	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	250,000	1,347	Scalable Graph Decomposition and Algorithms to Support the Analysis of Petascale Data

2011 Director's Discretionary Allocations

PI	Affiliation	2010 Allocation	Carryover to 2011	New 2011 Allocation	Total 2011 Allocation	2011 Usage	Project Name
Marc Snir	UIUC			100,000	100,000	0	Damaris
Terry Jones	ORNL			3,000,000	3,000,000	7,350,378	Extending Vampir IO for OLCF-3 Class Systems
Lee Berry	ORNL	3,000,000	80,635		80,635	128,026	Wave-Particle Intercations in Fusion Plasmas
Kai Germaschewski	ORNL			200,000	200,000	7,172	Load balancing particle-in-cell simulations
Thomas Jordan	University of Southern California			2,000,000	2,000,000	4,804	Deterministic Simulations of Large Regional Earthquakes at Frequencies up to 4Hz
John Dutton	Prescient Weather	100,000	100,000		100,000	0	CFS Reanalysis Extension
Omar Ghattas	University of Texas Austin	1,000,000	150,618		150,618	0	Forward and Inverse Modeling of Solid Earth Dynamics Problems on Petascale Computers
Cristiana Stan	Center for Ocean-Land-Atmosphere Studies			500,000	500,000	305,300	Simulations of Antropigenic Climate Change Effect Using a Multi-Modeling Framework
George Karniadakis	Brown University	1,500,000	1,276,488		1,276,488	1,559,654	NektarG-INCITE
Stephen Poole	ORNL	1,000,000	1,000,000		1,000,000	0	Gov-IP
John Cobb	ORNL	50,000	50,000		50,000	0	Neutron Scattering Science Exploratory Projects
James Chelikowsky	University of Texas Austin	500,000	406,510		406,510	295,241	Simulating the Emergence of Crystallinity: Quantm Modeling of Liquids
Rong Tian	Institute of Computing Technology, Chinese Academia of Sciences			900,000	900,000	2,594,605	Petascale simulation of fracture process
Shok Srinivasan	Florida State University			300,000	300,000	0	Accelerating Quantum Monte Carlo on Massively Parallel Computing Platforms
Bruce Harmon	AMES Lab			1,000,000	1,000,000	592,896	Beyond Rare Earth Magnets (BREM)
Predrag Krstic	UT-Knoxville			1,500,000	1,500,000	299,672	Science of the Plasma-Material Interface at Extreme Conditions

2011 Director's Discretionary Allocations

PI	Affiliation	2010 Allocation	Carryover to 2011	New 2011 Allocation	Total 2011 Allocation	2011 Usage	Project Name
Christopher Lynberg	Centers for Disease Control and Prevention	100,000	100,000		100,000	0	CSC Scientific Computing Architecture
Barry Schneider	National Science Foundation	2,000,000	18,574		18,574	0	Time-Dependent Interactions of Short Intense Laser Pulses and Charged Particles with Atoms and Molecules
Kenneth Smith	United Technologies Research Center	100,000	94,333		94,333	29,948	Surface Tension Predictions for Fire-Fighting Foams
Patrick Joseph Burns	Colorado State University	200,000	200,000		200,000	0	Parallel Lagged Fibonacci Random Number Generation
Shaikh Ahmed	Southern Illinois University Carbondale	1	0		0	9,241,935	Multimillion-Atom Modeling of Harsh-Environment Nanodevices
Bhagawan Sahu	University of Texas Austin	1,000,000	990,876		990,876	0	Gap Engineering in Trilayer Graphene Nanoflakes
William Martin	University of Michigan			1,000,000	1,000,000	0	Development of a Full-Core HTR Benchmark using MCNP5 and RELAP5-ATHENA
John Turner	ORNL			15,000,000	15,000,000	4,582,927	Fundamental studies of multiphase flows and corrosion mechanisms in nuclear engineering applications
Christopher Taylor	LANL	200,000	89,501		89,501	0	Fundamental Properties of the Stability of Exposed and Oxygen-covered Tc-Zr Alloy Surfaces from Density Functional Theory
Emilian Popov	ORNL	200,000	188,718		188,718	215,138	Testing STARCCM+ on Jaguar for Computing Large Scale CFD Problems
Dinesh Kaushik	ANL	2,000,000	2,000,000		2,000,000	0	Scalable Simulation of Neutron Transport in Fast Reactor Cores
Srdjan Simunovic	ORNL	100,000	14,493		14,493	22,163	Development of a Global Advanced Nuclear Fuel Rod Model

2011 Director's Discretionary Allocations

PI	Affiliation	2010 Allocation	Carryover to 2011	New 2011 Allocation	Total 2011 Allocation	2011 Usage	Project Name
Calvin Johnson	San Diego State University			500,000	500,000	0	Large-scale configuration-interaction nuclear shell-model code with factorization algorithms
Gary Grest	SNL	1,000,000	1,000,000		1,000,000	989,892	Assembly of Nanoparticles at Liquid/Vapor Interface
Masako Yamada	GE Global Research			100,000	100,000	220,478	Engineered icephobic surfaces
Xiao Cheng	University of Nebraska Lincoln			1,000,000	1,000,000	1,004,064	Exploration of Structural and Catalytic Properties of Gold Clusters
Bobby Sumpter	ORNL			4,000,000	4,000,000	4,548,189	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
Brian J Albright	LANL	1,000,000	2,000,000		2,000,000	5,407,844	Kinetic Simulations of Laser Driven Particle Acceleration
Nikolai Pogorelov	University of Alabama Huntsville	1,000,000	480,051		480,051	177,856	Modeling Heliospheric Phenomena with an Adaptive, MHD-Boltzmann Code and Observational Boundary Conditions
Homayoun Karimabadi	University of California San Diego	3,000,000	3,000,000		3,000,000	38,043,675	Enabling Breakthrough Kinetic Simulations of the Earths Magnetosphere through Petascale Computing
James Nutaro	ORNL	500,000	500,000		500,000	0	Qualitative System Identification for Massive Data Sets: Knowledge Discovery from Observations of Biological Systems
Ramesh Balakrishnan	ANL			100,000	100,000	0	The Performance of Turbulence Codes on Massively Parallel Computing Platforms with Multicore Processor Architectures
Oleg Zikanov	University of Michigan	400,000	396,401		396,401	101,228	Effect of Liquid-Phase Turbulence on Microstructure Growth During Solidification

2011 Director's Discretionary Allocations

PI	Affiliation	2010 Allocation	Carryover to 2011	New 2011 Allocation	Total 2011 Allocation	2011 Usage	Project Name
George Vahala	College of William and Mary	2,500,000	461,737		461,737	1,341,179	Lattice Algorithms for Quantum and Classical Turbulence
Michael Matheson	ORNL	500,000	1,084,560		1,084,560	1,059,237	Exploration of High Resolution Design-Cycle CFD Analysis
Praveen Ramaprabhu	University of North Carolina			862,160	862,160	4,082	Simulations of turbulent mixing driven by strong shockwaves
David Rector	PNNL			400,000	400,000	363,237	Solid-liquid tank mixing using the implicit lattice kinetics method
Pui-kuen Yeung	Georgia Institute of Technology			3,000,000	3,000,000	2,108,993	Frontiers of Computational Turbulence
Misun Min	ANL			900,000	900,000	3	Codes for High Order Methods

Total 138,560,686 142,880,327

Innovation

HIGH PERFORMANCE COMPUTING FACILITY 2011 OPERATIONAL ASSESSMENT OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2012

4. INNOVATION

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

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

2011 Operational Assessment Guidance

The Facility should highlight innovations that have improved its operations, focusing especially on best practices:

- *that have been adopted from other Facilities,*
- *that the Facility has developed and/or recommended to other Facilities, and*
- *that other Facilities have adopted.*

2011 Approved OLCF Metrics – Innovation

Innovation Metric 1: The OLCF will report on new technologies that we have developed and best practices we have implemented and shared.

The OLCF has carried out numerous activities ranging from working with users to update their applications to maximize their effective use of anticipated systems, to technology innovations that streamline workflow, to tools development.

Innovation Metric 2: The OLCF will report on technologies we have developed that have been adopted by other centers or industry.

The OLCF has developed a number of technical innovations that have been adopted by other centers and industry. Our work on exploiting hierarchical parallelism within applications to better map to next-generation architectures is being adopted by the communities who developed these applications. To this end, the OLCF established the Center for Accelerated Application Readiness (CAAR). A guiding principle of this effort has been to directly integrate these capabilities into the canonical source tree of each application thereby easing longer-term maintenance of the application and portability of these enhancements. The OLCF's work in topology aware I/O, specifically our topology aware Lustre network routing capabilities have been incorporated into the canonical Lustre source tree and the knowledge required to make use of these capabilities have been disseminated through a number of publications and presentations by OLCF staff. Our work on the Common Communication Interface (CCI) is a collaborative development effort

conducted in concert with other laboratories (SNL, INRIA) and industry (Cisco, Myricom). The OLCF has funded and managed contract development of scalable and heterogeneous debugging features that have been incorporated into the Allinea DDT debugging tool. To improve code portability and ease porting to advanced architectures the OLCF has funded and managed contract development of accelerator enhancements in the CAPS HMPP compiler, a commercially available product. The OLCF has funded and managed contract development of scalable performance analysis for heterogeneous systems in the widely used Vampir tool set allowing these capabilities to be utilized by HPC centers around the world. To ensure a long-term sustainable Lustre ecosystem, the OLCF has lead the way in establishing Open Scalable File Systems, a non-profit mutual benefit corporation. Through direct engagement with other HPC centers, vendor partners, and application development teams, the OLCF is ensuring that ASCR investments that culminate in technical innovations have broad impact to the entire HPC ecosystem.

Innovation is the heart of HPC. Innovation not just in the science enabled by the computing power inherent in high-performance computers, but in HPC itself. The increasing complexity of the world we live in is making innovation a matter of careful, long-range planning.⁸ OLCF activities this past year reflect this, with staff members across the organization contributing to planning for the next generation of HPC. Judging by the results, the OLCF will be more than ready to take advantage of the technological breakthroughs looming with the advent of such leading edge technologies as multithreaded parallelism, general purpose GPUs, and multicore-aware software. The following pages describe some of these exciting new developments, pioneered and led by OLCF staff.

4.1 Application Readiness

Preparing the Scientific Community for Titan – Innovative Application Readiness

In 2012 the OLCF will deploy a large-scale, hybrid- multicore node-based system known as Titan for use as a major compute resource for DOE SC. The nodes on this system will have an industry standard x86-64 architecture processor paired with a GPU-based application accelerator. The resulting node will provide a peak performance of more than 1 teraflop.

The new hybrid node architecture will require application teams to modify their codes to expose more parallelism. Given the marked difference in node architecture, substantial effort will be needed to bring scientific applications to the point of effective use of the new platform. The primary challenges involved in marshaling the GPUs are threefold:

- recognition and exploitation of hierarchical parallelism by scientific applications, including distributed memory parallelism via message passing interface (MPI), symmetric multiprocessing (SMP)-like parallelism via threads (OpenMP or pthreads), and vector parallelism via the GPU programming;
- development of effective programming tools to facilitate this (often) substantial rewrite of the application codes; and
- deployment of useful performance and debugging tools to speed this refactoring.

⁸Dosi, G., “Technological paradigms and technological trajectories,” *Research Policy*, 11 (1982), pp. 157–162.

To lead the way, the OLCF established the Center for Accelerated Application Readiness (CAAR), whose members include application teams, vendor partners, and tool developers. CAAR is charged with preparing six representative applications for Titan. The six applications, selected from among 50 of the most productive applications running on Jaguar, were chosen because they represent much of the range of demands that will be placed on Titan from a variety of scientific domains. application and Software development leadership

Each of the CAAR teams is led by an OLCF staff member from the Scientific Computing Group. The teams also include representatives from the individual code development groups, engineers from OLCF vendor partners Cray and NVIDIA, and, in some cases, other OLCF and ORNL staff members. The SciComp CAAR team leaders are responsible for coordinating the work of their teams and have shared responsibility with the code owners in formulating the science targets for OLCF-3. One of the most important responsibilities of the CAAR team leads is to ensure that changes made to facilitate the port to OLCF-3 are retained in the production trunk of each code. This vital step helps assure portable performance, as changes made that increase data locality and expose hierarchical parallelism prove useful even on non-hybrid architectures.

In 2011 the CAAR team has made significant progress. LAMMPS, a molecular dynamics application is currently 2x to 5x faster with GPU acceleration versus using 16 cores per node on the XK6 with existing LAMMPS algorithms. A new parallel linear scaling electrostatic solver using the Multilevel Summation Method (MSM) algorithm is approximately 95% complete. LSMS, a materials science application, has been modified to include a new main block inversion kernel (zblock_lu) that is ~30x faster on the GPU compared to a single Opteron core. This kernel represents over 95% of CPU runtime of LSMS. A new multithreaded version of LSMS is currently under development incorporating this GPU optimized kernel. S3D, a highly used combustion application is now running on the XK6 with OpenACC directives. The application has been refactored to either run on a multi-core architecture using OpenMP or on a XK6 with an NVIDIA accelerator. S3D will be ready to run a new chemical species as soon as the chemical mechanism is delivered. Denovo, a radiation transport application whose 3-d sweep algorithm represents 80-99% of the total runtime has been modified to offload much of the 3-d sweep to the GPU resulting in a ~40x speedup compared to a single Opteron core. This speedup is equivalent to a 7x speedup on a Fermi GPU relative to a single Istanbul 6-core processor. While these results are quite impressive, it is important to note that the next-generation NVIDIA GPU known as Kepler will substantially improved performance over the current generation Fermi GPU.

The totality of these CAAR code port experiences, like much of the work the SciComp liaisons produce in support of production work on Jaguar, is being transmitted to broader community through several means, including dissemination of best practices at OLCF workshops and conference proceedings and the availability of production software packages and libraries (e.g., the Multi-level Summation Method kernel from the CAAR code LAMMPS will be made available as a library to other MD practitioners). The CAAR experiences and lessons-learned will lead to the most complete and sustainable set of practices available for hybrid multicore computing for the near future.

4.2 Application Support

Vendor Partnership to Promote Effective Programming Environments – Innovative Application Support

Two years ago Cray Inc. and Oak Ridge National Laboratory determined that a heterogeneous architecture was the most power efficient path to Exascale computing. Once this decision was made, the two teams identified numerous challenges facing the effective utilization of the new heterogeneous system. In particular the programming approach was discussed and a direction was formulated to deliver compiler and tools to allow the user community to employ portable comment line directives for moving existing application to the new hybrid architecture.

Today Cray Inc. is delivering a GPU programming environment that consists of tools to assist the code developers in the refactoring of their applications for the heterogeneous system. The elements of the programming environment and the design of the components was achieved through numerous discussions between Cray developers and DOE Office of Science researchers and members of the OLCF Scientific Computing Group. In working with the six code teams comprising the CAAR, Cray representatives were able to identify what tools would help in the porting of codes to the new system. Without this close interaction between Cray Inc. and ORNL the resultant programming environment could have been ineffective; however, we are seeing today that the tools and compiler capabilities are exactly what the code developers need to address the challenges in moving their applications to the new system.

High-productivity Hybrid-programming Through Compiler Innovation – Innovative Application Support

Applications of interest to OLCF-3 are written in C/C++ and Fortran 77/90, with MPI; OpenMP; and, in some cases, DSLs. To improve user code porting and development productivity, the OLCF-3 project will support the use of high-level languages with accelerator directives. The Center is exploring the use of OpenACC, Cray, PGI, and HMPP accelerator directives and has initial performance assessments on kernels written in C and Fortran, which requires minor modification to the original source code and can be retargeted to different platforms. As part of this effort, Cray and CAPS (HMPP vendor) have each agreed to support the OpenACC set of directives, providing users with a set of compiler directives supported by all OLCF-3 compiler vendors. In addition, the Application Performance Tools group is working with CAPS enterprise (www.caps-enterprise.com) to ensure that the directive support provides the full range of application needs, investigating support needed beyond the Opencast specification.

Since the last review, CAPS has delivered a document co-authored with Allinea that describes a design for providing the debugging hooks needed by debuggers such as DDT. This enhancement would enable users to debug their original code as it was written, rather than debugging the HMPP-transformed code. In addition, HMPP version 3.0 was released, which is the production version of the compiler that includes enhancements that were developed as part of the earlier work for OLCF-3 such as memory placement and coordination between a device and host.

Scalable Debugging for Hybrid Systems – Innovative Application Support

In order to develop applications on a massive, hybrid, GPU-based cluster system, users will require a scalable, hybrid, platform-aware debugger. The OLCF is working closely with Allinea, who have been tasked with providing this essential component of the programming environment (PE), to enable their debugger, Distributed Debugging Tool (DDT), to meet OLCF-3's demanding needs. Current efforts have

resulted in a debugger that can scale to over 200,000 cores as well as handle the debugging of GPU kernels.

The collaboration with Allinea is allowing the OLCF to address the problem of scalable debugging. In the current period of performance, Allinea has delivered to ORNL a first version of visualization support for large-scale analysis. This feature is designed to help users look for anomalies in the data, to help hone in on the source of the errors. This feature is very new, and has not been exercised much yet, however, as the vendor has communicated to us “I thought you would also be interested to know, as we talk to others about the visualization work, they are extremely supportive and interested in the work we are doing with you. (More than I expected). Last week we conducted training sessions with Cray at NCSA (we are selected for Bluewaters machine – thanks to ORNL) and ANL. At both locations, we found great interest in our Visualization work. You are definitely investing in areas of very strong need.”

Work has also continued to enhance DDT’s GPU support, adding support for multi-warp stepping and initial evaluation of the requirements for Cray compiler GPU debugging support.

Scalable Performance Analysis for Hybrid Systems – Innovative Application Support

The Vampir (Visualization and Analysis of MPI Resources) tool set is used for performance analysis in OLCF-3. We are working together with Vampir’s vendor, the Technische Universität Dresden, to make this tool set ready for the targeted OLCF-3 system. Vampir uses program tracing to record a detailed list of events during the execution of an application. Using a set of compiler wrappers for C, C++, and FORTRAN, the application can be built with specific instrumentations.

VampirTrace enables instrumentation of the parallel paradigms MPI and OpenMP/Threads, as well as generic recording of function invocations through compiler or manual instrumentation. Vampir further provides a postmortem visualization of the program execution based on the recorded trace. This visualization features a rich set of displays to help understand the fine-grained behavior of the application. The Vampir visualization framework is provided through a parallel server and GUI application, allowing the processing of very large traces. The entire tool chain is tailored for scalable parallel analysis. To match the scale of the target OLCF-3 system, additional improvements have been and are being incorporated in Vampir. Specific optimizations to the communications layer of the VampirServer now enable the use of more than 10,000 analysis processes. Multiple improvements target the handling of an increasing amount of trace data from hundreds of thousands of processes. Pattern matching-based compression will improve recording capabilities, while filtering and highlighting of irregularities will support the evaluation of large-scale traces.

Over the period of performance, support for NVIDIA’s GPU hardware counters, via the CUPTI interface, integration with the HMPP compiler, reduction in the size of collected traces, timeline filtering, and comparison of multiple runs has been added to this tool suite.

4.3 Outreach

Empowering American Industry through High Performance Computing – Innovative Outreach

About 60 software experts gathered in Chicago on March 31, 2011, for the first Workshop for Independent Software Developers and Industry Partners, sponsored by the DOE Advanced Scientific Computing Research office. Jointly organized by

Lawrence Berkeley and Oak Ridge National Laboratories, this workshop introduced independent software vendors (ISVs) and industrial software developers to software resources that can help ease the private sector's transition to multicore computer systems. These tools, libraries, and applications were developed through DOE's Scientific Discovery through Advanced Computing (Sciacc) program to enable DOE's own critical codes to run in a multicore environment.

The cost and difficulty of scalable parallelizing legacy codes (codes written for nonoperational or outdated operating systems or computer technologies) often are prohibitive to independent software vendors, particularly if they are small businesses. They also hamper many firms that, for proprietary and competitiveness reasons, maintain their own code in addition to using commercial options. The problem is becoming acute as desktop workstations and small clusters are rapidly being designed and built using multicore processors.

The 1-day workshop was an important contribution to addressing these hurdles. It gave participants an overview of the SciDAC program and more than 60 SciDAC-developed software packages and outlined the process to obtain them, often at no cost. In addition, DOE explained its role in providing research grants through the U.S. Small Business Administration's Small Business Innovation and Research (SBIR) grant program. This program ensures that the nation's small, high-tech, innovative businesses are a significant part of the federal government's research and development efforts. Workshop participants then provided feedback on private sector software development requirements that could help DOE shape future SBIR research topics and jumpstart areas for collaboration.

"SciDAC has spent a decade developing world class software to ensure DOE can operate successfully in a multicore environment," explained David Skinner, workshop cochair and director of the SciDAC Outreach Center at Lawrence Berkeley. "The private sector software developers who participated now have direct links to key developers who can provide expertise in developing software for multicore systems and help guide integration of SciDAC software into commercial applications. We hope to extend these links to those who could not attend."

The workshop's participants represented 49 organizations, including small and large ISVs, companies with internal software development capabilities, academic institutions, other national laboratories, and HPC system vendors.

"This event launched a new opportunity to leverage DOE's investment in SciDAC for an additional return on investment for the country," said fellow chair Suzy Tichenor, director for the HPC Industrial Partnerships Program at Oak Ridge. "Most of the ISVs and companies that attended had never heard of the SciDAC program. Now they are aware of SciDAC's valuable software resources and how to access them."

4.4 Systems

Breaking Bottlenecks in Parallel I/O – Innovative Systems

In 2011, the user-achievable bandwidth on Spider was more than doubled. This was accomplished without purchasing any additional hardware by carefully considered configuration changes.

Spider is a "routed" file system, which means that it uses I/O nodes on the Jaguar system to move information between two physically incompatible interconnect topologies; in this case, the Cray SeaStar network on Jaguar and the 20 Gbps InfiniBand on Spider. Because Spider offers aggregate bandwidth far in excess of the single-link speeds of either interconnect, avoiding congestion is fundamental to

achieving efficient I/O. Unfortunately, simple configurations of Lustre at large scale inherently induce congestion in the InfiniBand fabric. By default, Lustre disperses traffic to all routers in a round-robin fashion. This causes traffic to be injected into the InfiniBand fabric's fat-tree topology in nonoptimal locations, which in turn causes oversubscription and congestion on internal links of the fabric. Significant performance degradation due to this issue has been measured. Additionally, this dispersal of traffic to the routers prevents using locality information to optimize application I/O performance, as it is impossible to know which router will service each request.

The OLCF has completely eliminated congestion inside the InfiniBand fabric by pairing routers with individual Spider servers. This one-to-one mapping keeps traffic inside the crossbar switch and prevents it from traversing the internal links of the fat-tree. In addition, traffic for a given server takes a more direct route within the torus. This configuration change improved demonstrated read bandwidth by 101% and gave a 93% improvement for write bandwidth for applications without regard to their locality. For tests in which the I/O targets were chosen based upon location in the torus, the new routing configuration allows improvements of up to 115% for reads and 137% for writes.

This information was shared with the larger user community during the 2011 Cray User Group meeting and is available as an ORNL technical report via <http://info.ornl.gov/sites/publications/Files/Pub30140.pdf>.

Intuitive Data Portal for Collaborative Climate Science – Innovative Systems

The Earth System Grid Federation (ESGF) is a large-scale, multi-institution, interdisciplinary project to provide climate scientists worldwide, as well as climate impact policy makers, a web-based platform to publish, disseminate, compare, and analyze ever increasing amounts of climate-related data. ORNL is a key contributor to the ESGF project with development and data publication efforts funded by the DOE Office of Science - Biological and Environmental Research. While BER funds the development and software maintenance of ESGF at ORNL, the OLCF has assisted in the architecture and deployment of the system infrastructure required to provide climate scientists with access to the high-value datasets resident within the OLCF. The ESG data storage is publicly accessible and leverages HPSS infrastructure without compromising OLCF data.

As a result of this work, the ORNL-ESG system hosted within the OLCF provides access to a number of high use, high value data sets, including the following.

- Climate Modeling Best Estimate atmospheric, cloud, and radiation quantities showcase data sets from the Atmospheric Radiation Measurement Program
- Carbon-Land Model Intercomparison Project data set
- Ameriflux (part of the FLUXNET global network of towers making continuous measurements of CO₂, water vapor, and radiation via eddy covariance in terrestrial ecosystems) and Fossil Fuel (gridded fossil-fuel CO₂ emission estimates) data from the Carbon Dioxide Information Analysis Center data set
- The Ultra High Resolution Global Climate Simulation project

Achieving High-Performance Communication While Preserving Portability – Innovative Systems

The sheer size of the OLCF imposes scalability issues for everything from storage to debugging tools. In addition to Jaguar, the OLCF includes many different types of hardware including multiple types of network infrastructures. Each network

provides at least two application programming interfaces (APIs): BSD sockets; and the network's native interface, which provides better performance through direct access to the network hardware. Jaguar, for example, provides Portals while the storage system uses Verbs. Cray's next generation of hardware replaces SeaStar with Gemini and replaces Portals with GNI.

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).

The Technology Integration Group (TechInt) is 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 jointly developed by ORNL, SNL, University of Tennessee, Myricom, and Cisco. CCI is designed for portability, scalability, and performance.

TechInt continues to refine the API and has completed support for UDP, Cray Portals, and InfiniBand Verbs. We are working on support for GNI and expect to complete it in early 2012. Performance results are encouraging with InfiniBand and GNI latencies as low as 1.4 u-sec and bandwidth in-line with that of the native interface.

Real-time Monitoring of Simulations through an Integrated Dashboard – Innovative Systems

Computational scientists have a new weapon at their disposal. On February 1, 2011, the electronic Simulation Monitoring (eSiMon) Dashboard, version 1.0, was released to the public, allowing scientists to monitor and analyze their simulations in real time. Developed by the Scientific Computing and Imaging Institute at the University of Utah, North Carolina State University, and ORNL, this “window” into running simulations shows results almost as they occur, displaying data just a minute or two behind the simulations themselves. Ultimately, the Dashboard allows scientists to concentrate on the science being simulated rather than having to learn HPC intricacies, an increasingly complex area as leadership systems continue to break the petaflop barrier. This work was funded through collaboration between DOE/ASCR and DOE/FES.

4.5 Leadership

Empowering a Sustainable Lustre Ecosystem through OpenSFS – Innovative Leadership

The Lustre parallel file system is the most used parallel file system technology in HPC, with use on more than 70 of the top 100 HPC systems and all of the top 5 systems in the November 2010 Top500 list. In 2010 the OLCF teamed with Cray, DDN, and LLNL to form Open Scalable File Systems, Inc. (OpenSFS), a nonprofit mutual benefit corporation for development of high-end open-source file system technologies, with a focus on the Lustre parallel file system.

OpenSFS is now embarking on the development of next-generation features within the Lustre file system, allowing the OLCF to meet its current and future HPC requirements. The OpenSFS model allows the OLCF to leverage others' investment in the Lustre file system while preserving its ability to oversee collaborative development efforts. Having released a request for proposals for Lustre development in April 2011, OpenSFS awarded a contract to Whamcloud to develop a number of new features in the Lustre file system. These features include single server metadata performance acceleration, horizontally scalable metadata services, and an online file system checker. These features were identified as the highest priority items through a broad canvassing of the Lustre community as part of the OpenSFS

Technical Working Groups requirements analysis process. Development of these features is currently underway and they are scheduled for release in Lustre 2.3 in 2012 and Lustre 2.4 in 2013.

In addition to leading the way in the continued evolution of Lustre, OpenSFS has partnered with Whamcloud to maintain current Lustre releases and provide an open and transparent community development model. This model allows the entire Lustre community to contribute to the Lustre codebase while maintaining rigorous quality control of the canonical Lustre source tree. A number of services are now provided to the Lustre community through OpenSFS to make this a reality, including Lustre gatekeeping, patch inspection, and testing.

OpenSFS is now recognized as the leader in the Lustre community. In 2011 OpenSFS has grown to include representatives from Cray, Data Direct Networks, Indiana University, Lawrence Livermore National Laboratory, Netapp, Oak Ridge National Laboratory, RAID Incorporated, Sandia National Laboratories, SGI, Terascale, Whamcloud, and Xyratex. Xyratex joined OpenSFS at the promoter level and now has a seat on the OpenSFS board. Indiana University was nominated and confirmed as the community board representative in 2011. A number of other organizations are currently under the due-diligence process of joining and we expect even broader participation in 2012. In total, OpenSFS manages an annual operational budget of over \$2.7 million dollars making it the largest single investor in open source parallel file system technologies.

The OLCF's leadership role in OpenSFS has resulted in a single Lustre community represented by OpenSFS and the European Open File System consortium (EOFS). This collaboration, the first of its kind in the HPC world, was announced at the first Lustre User Group Meeting (organized by the OLCF) and ratified through a memorandum of understanding between OpenSFS and EOFS signed at this year's International Supercomputing Conference (June 19–23, 2011, Hamburg, Germany). OLCF leadership fostered this collaborative approach to continued Lustre development and thus ensured the future of the Lustre file system.

4.6 Energy Management

Effects of CRU Top Hats on Air Flow – Innovative Energy Management

The ORNL Computer Facilities Manager and Facilities & Operations continue to evaluate various methods for improving the airflow within the data center, especially in high-density areas, and in constrained-supply areas. The target goals include increasing the capacity or effectiveness of an air handler, providing greater control over the air-distribution process, and providing more optimal inlet air temperatures to high-density air-cooled equipment.

In July 2011, ORNL installed air handler top hats on two 30-ton units. These top hats are simple ducting extensions that pull return air from a higher stratification in the data center. With the top hats installed, ORNL measured an increased return air temperature of 6 degrees Fahrenheit. According to the ASHRAE psychrometric chart for mechanical cooling performance, a rise from 70°F to 75°F at 50% Relative Humidity is equivalent to a 45% increase in cooling capacity at identical motor kW. Given the relatively low material cost for the top hats, and the high performance increase, ORNL is extending the installation of these top hats to the remaining air handlers in the Computational Sciences Building.

The results of this experiment are shown in Figure 4.1. Two CRUs, labeled CRU 39 and CRU 40 were sampled before and after top hat installation. These two units reside in a very dense air-cooled equipment area that has traditionally demonstrated

mechanical challenges with both control of inlet temperatures, and control of exhaust heat. The summary of the impact of the top hats on the CRU on the return-air temperatures is shown in Table 4.1.

Also in July 2011, a pilot program was initiated to implement a hot air containment solution on several cabinets within the air cooled section of the data center. This pilot consisted of fan assisted chimneys placed on pairs of cabinets that were open to each other on one side. While there were some positive effects measured with the implementation, they were not sufficient to warrant full deployment throughout the data center, and were removed in December 2011. Implementation of a cold aisle containment system on the same cabinets was designed for implementation in January 2012.

Figure 4.1. The Effect of Top Hats on CRU Efficiency

Table 4.1. The Positive Impact on CRU Return-air Temperatures with Top Hats

	CRU 39			CRU 40		
Degrees Fahrenheit	71.0	76.9	6.0	81.7	87.0	5.3
Configuration	Without top hat	With top hat	Temp. increase (measured, average)	Without top hat	With top hat	Temp. increase (measured, average)

A number of activities continue, including studies on the effectiveness of hot/cold air separation techniques; use of water-side economizers; addition of VFDs to Central Energy Plant chillers and chilled-water pumps; cool-roof technologies; new controls for chilled-water delivery that optimize cooling load, environmental conditions, and available equipment; increasing the delivered chilled-water temperature; chilled-water storage; and load shedding.

Risk Management

HIGH PERFORMANCE COMPUTING FACILITY
2011 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2012

5. RISK MANAGEMENT

CHARGE QUESTION 5: Is the Facility effectively managing risk?

OLCF RESPONSE: The OLCF has a very successful history of anticipating, analyzing and rating, and retiring risk for both project-based 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, are retired, re-characterized, or mitigated. The major risks currently being tracked are listed and described below. Any mitigation(s) planned for or implemented are included in the descriptions. The OLCF has only one "high" operational risk: that the funding picture for 2013 is uncertain. To address this risk, the OLCF will continue to work with our DOE sponsors to understand their projections and adjust our plans accordingly.

5.1 Risk Management

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. Each Project Execution Plan refers to the main RMP but may incorporate some tailoring specific to that 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 focus their attention on the high and moderate risks and any low risks within the impact dates associated with the risk. Trigger conditions are stated in the Risk Response narrative section of the registry when appropriate and 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 their risks for which they are responsible and for bringing appropriate attention to the management of those risks whatever their rating level.

The OLCF sends aggregated risk reports monthly to the DOE program office. 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).

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 WBS. However, for operations, costs of accepted risks and residual risks are estimated by expert opinion and are accommodated as much as possible in management reserves. This reserve is re-evaluated continually throughout the year.

The following are the major risks in the OLCF Operations Risk Register.

5.2 Major Risks Tracked in the Current Year

Risk	Register ID	Rating	Notes
Insufficient funding to meet DOE commitments (FY2013)	979	High	Uncertainty is a concern
Insufficient funding to meet DOE commitments (FY2012)	974	High reduced to Medium	Funding appears to be sufficient for FY2012. Will be reduced to low when funding arrives.
Insufficient funding to meet DOE commitments (FY2011)	404	High reduced to Medium	Retired: Funding was sufficient for FY2011
Lack of infrastructure for an exascale system	948	Medium	Long lead time required to resolve. Remains a concern
Lustre metadata performance continues to impact applications	721	Medium	Mitigation includes participation in OpenSFS. Remains a concern
Scientists decline to port to heterogeneous architecture	361	High reduced to Medium	Mitigation includes in house experience and training development. Remains a concern
Programming environment tools may be insufficient	906	High reduced to Medium	Mitigation includes subcontracts with tool vendors. Remains a concern
New architecture requires new programming model	912	High reduced to Medium	Mitigation includes portable programming model. Remains a concern
Supply chain issues e.g. flooding in Thailand may impact planned hard disk acquisitions	973/975	High / Medium	#973 retired and replaced by #975 with lowered risk.
System upgrade from XT5 to XK6 takes too long causing users to seek alternative computing resources	915	Medium	Retired: System returned to operations.

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

Risk Owner Buddy Bland
Probability Medium
Impact *Cost:* High *Schedule:* Medium *Scope/Tech:* High
Rating High
Status Accepting the risk

Annual budgets are set with guidance from 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.

The FY13 budget is being formulated now and substantial reductions from current projections are being discussed. We will maintain close contact with Federal Project Director and ASCR Program Office to understand the changing funding projections so alternative plans can be made in sufficient time. Where possible, structure contracts to accommodate flexible payment terms.

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

Risk Owner	Buddy Bland		
Probability	Low		
Impact	Cost: Low	Schedule: Low	Scope/Tech: High
Rating	Medium		
Status	Accepting the risk		

Annual budgets are set with guidance from 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.

At this time it appears that funding will be sufficient for FY2012. We will reduce the rating to Low once funding is in hand. Maintain close contact with Federal Project Director and ASCR Program Office to understand the changing funding projections so alternative plans can be made in sufficient time. Where possible, structure contracts to accommodate flexible payment terms.

5.2.3 ID# 404 – Insufficient Funding to Meet DOE Commitments (FY2011)

Risk Owner	Buddy Bland		
Probability	Low		
Impact	Cost: High	Schedule: Low	Scope/Tech: High
Rating	Medium		
Status	Retired		

Annual budgets are set with guidance from 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. Funding was sufficient in FY2011.

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

Risk Owner	Jeff Nichols		
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 risk if the systems are significantly larger or use more power than projected.

ORNL has a plan to house the exascale system in building 5600 by moving other systems out of the building. However, the much preferred approach would be to build a new building that is designed for exascale from the beginning. OMB has rejected third party financing as a method of building such a facility so this will need a congressional line item.

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

Risk Owner	Galen Shipman		
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. 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.

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. Contract development through the OLCF with Whamcloud is accelerating the deployment of Lustre 2 on Jaguar which has demonstrated improved performance, confirmed during dedicated Lustre test shots on Jaguar. Some improvement has been realized in the Lustre 2.2 version which is going to be released by Whamcloud over the next month or so (it is in code freeze now). We will be deploying Lustre 2.2 along with our storage system upgrade in September. The Lustre upgrade will include Lustre metadata enhancements to for both interactive workloads "ls -l" as well as simulation workloads which often create large numbers of files – tens of thousands or more – in a single I/O epoch.

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 in order to minimize the impact to the overall user population. Multiple file systems have been deployed reducing load on the metadata server. Even with all of this mitigation effort, progress toward a solution is slower than expected and the risk remains a concern.

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

Risk Owner	Jack Wells		
Probability	Medium		
Impact	Cost: Low	Schedule: Medium	Scope/Tech: Low
Rating	Medium		
Status	Mitigating the risk		

Some users may determine that it is too much effort to port their code to the new heterogeneous architecture. This risk and risks ID# 906 and 912 are all related to the introduction of the new heterogeneous architecture by the OLCF-3 project.

Original risk evaluation rated this as High. Mitigation with outreach, training, and the availability of libraries and development tools will ameliorate some of the resistance. Current trends in publication venues imply that many development teams are exploring architectures with accelerators which are contrary to this risk. The marked improvement of compiler directive technology from Cray, CAPS, and PGI (including the OpenACC standardization) has removed significant technical barriers for computational scientists in porting their codes to GPUs.

5.2.7 ID# 906 – Programming Environment Tools May Be Insufficient

Risk Owner	Rich Graham		
Probability	Medium		
Impact	Cost: Medium	Schedule: Low	Scope/Tech: Medium
Rating	Medium		
Status	Mitigating the risk		

As was pointed out by the CD-1 Lehman Review panel, the OLCF-3 system will not be perceived as successful if programming the system requires that the users use a very different programming method that would not be compatible with other large system such as Jaguar, and the new BG/Q system at ANL. We have developed a strategy to prevent this problem by using compilers, debuggers, and performance measurement tools that are compatible with other systems for the programming environment of OLCF-3. The risk is that these tools will not be sufficiently effective in generating acceptable code with reasonable performance on the applications.

Original risk evaluation rated this as High. For mitigation, we have moved the development of the compilers and tools into the OLCF-3 project as an initial risk mitigation. We will monitor the progress of the tools developers, and check out early versions of the tools on new Fermi processors in Jaguar, and on other GPU enabled systems to ensure the compatibility with existing programming models. We are developing portable programming models (through our vendor partners) such as the CAPPs source-to-source compiler and OpenMP directives for accelerators.

5.2.8 ID# 912 – New Architecture Requires New Programming Model

Risk Owner	Bronson Messer		
Probability	Medium		
Impact	Cost: Low	Schedule: Medium	Scope/Tech: Low
Rating	Medium		
Status	Mitigating the risk		

The OLCF-3 system has a new computer architecture, using both traditional x86 CPUs and GPUs to achieve unprecedented performance and energy efficiency. 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 have developed 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.

Original risk evaluation rated this as High. 960 Fermi cards were integrated 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 is an operational risk, not a project risk, it is important to work with users early to begin porting to the system so that the machine will be judged as successful by delivering breakthrough science.

5.2.9 ID# 975 – Supply chain issues i.e. flooding in Thailand may impact disk drive acquisitions

Risk Owner Galen Shipman
Probability Medium
Impact *Cost:* Medium *Schedule:* Medium *Scope/Tech:* Medium
Rating Medium
Status Mitigating the risk

Planned advanced storage system deployment may be delayed beyond Q4 2012 with the continued impact of flooding in Thailand on the availability of hard disk drives.

Original risk (ID# 973) was rated as High. For mitigation, we maintained very close contact with suppliers to monitor the situation and to continually evaluate price projections. As a result of this action, we were able to develop a purchase plan that we hope to initiate in May 2012. We still monitor the situation, but we retired #973 replacing it with #975 which accepts the risk. Should the situation not improve or indicate improvement by May, prior to release of the storage system RFP, we will consider a reduced scope (less performance due to increased disk prices), continuing to hold the RFP until the market improves, and/or increasing the level of funding for the acquisition. Until issue is resolved, we can continue to use the existing Spider file system.

5.2.10 ID# 915 – System Upgrade from XT5 to XK6 Takes Too Long Causing Users to Seek Alternative Computing Resources

Risk Owner Ann Baker
Probability Medium
Impact *Cost:* Low *Schedule:* Low *Scope/Tech:* Medium
Rating Medium
Status Retired

With a new system upgrade of the size and complexity of XK6, there may be problems that delay completion of the acceptance tests, thus delaying user access. The risk was retired when the system was returned to operations.

5.3 Risk That Occurred during the current year and the effectiveness of their mitigations

Risk	Register ID	Rating	Notes
Supply chain issues i.e. flooding in Thailand may impact planned hard disk acquisitions	973*	High	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 973 retired and replaced with a new risk 975.
Lustre metadata performance continues to impact applications	721*	Medium	Working with other major Lustre stakeholders and working with apps to restructure codes. Progress slow so risk remains a concern.
Loss of key personnel	407	Low	Long term incapacity of a key group leader occurred in 2011 but effective OLCF career development programs and a maintenance of sufficient depth in remaining staff allowed uninterrupted progress.

*These risks have been described in Section 5.2 above.

5.4 Risks retired during the current year

Risk	Register ID	Rating	Notes
Insufficient funding to meet DOE commitments (FY2011)	404*	Medium	Retired: 12/8/2011 - Funding was sufficient for FY2011
System upgrade from XT5 to XK6 takes too long causing users to seek alternative computing resources	915*	Medium	Retired: 2/13/2012 The system was returned to operations.
Applications are not ready for new technologies	410	Medium	Retired: 8/6/2011 – Split into 912, 361, and 906 because of multiple aspects of the risk, e.g., user reluctance, programming model and programming environment tools readiness.
Supply chain issues i.e. flooding in Thailand may impact planned hard disk acquisitions	973*	High	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 973 retired and replaced with a new risk 975.
Differences between Lustre versions on Spider and the Cray systems may impede integration.	138	Low	Retired: 8/6/2011- Operational processes were developed to test and integrate new Lustre releases and stage upgrades to maintain compatibility of systems across the OLCF complex.
INCITE hour goals may not be met because of upgrade to Jaguar	931	Low	Retired 12/8/2011 - Risk did not occur.
Future disk technology may be different from expected	393	Low	Retired: 8/6/2011 - We have a very good understanding of what disk technologies will be available for our next procurement through careful market analysis.

*These risks have been described in Section 5.2 above.

5.5 Major New or re-characterized risks since last review

Risk	Register ID	Rating	Notes
Insufficient funding to meet DOE commitments (FY2013)	979*	High	Uncertainty is a concern
Lack of infrastructure for an exascale system	948*	Medium	Long lead time required to resolve. Remains a concern
Supply chain issues i.e. flooding in Thailand may impact planned hard disk acquisitions	975*	Medium	Purchase plan developed that includes possible scope or schedule adjustments if risk impacts price projections.
Applications are not ready for new technologies	410*	Medium	Retired: 8/6/2011 – Split into 912, 361, and 906 because of multiple aspects of the risk, e.g., user reluctance, programming model and programming environment tools readiness.

*These risks have been described in Section 5.2 above.

5.6 Major Risks for next year

Risk	Register ID	Rating	Notes
Insufficient funding to meet DOE commitments (FY2013)	979*	High	Uncertainty is a concern
Insufficient funding to meet DOE commitments (FY2012)	974*	High reduced to Medium	Funding appears to be sufficient for FY2012
Lack of infrastructure for an exascale system	948*	Medium	Long lead time required to resolve. Remains a concern
Lustre metadata performance continues to impact applications	721*	Medium	Remains a concern
Supply chain issues i.e. flooding in Thailand may impact planned hard disk acquisitions	975*	Medium	Remains a concern.

*These risks have been described in Section 5.2 above.

Summary of the Proposed Metric Values

HIGH PERFORMANCE COMPUTING FACILITY
2011 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2012

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 2011, and proposed metrics and targets for 2012 and 2013.

2011 Metric	2011 Actual	2012 Metric	2012 Target	2013 Target
<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 OLCF score on the user survey will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	OLCF exceeded the target in 2011 with a survey mean score for overall customer satisfaction of 4.2, "very satisfied." (Reference survey questions #11, #15, #20, #25, and #27).	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 user survey results will show improvement in at least ½ of questions that scored below satisfactory (3.5) in previous period.	No user responses in the 2011 period were below the 3.5 satisfaction level. OLCF met this target.	Improvement on results that scored below satisfactory in the previous period.	Results will show improvement in at least ½ of questions that scored below satisfactory (3.5) in the previous period.	Results will show improvement in at least ½ of questions that scored below satisfactory (3.5) in the previous period.
<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.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.
80% of OLCF user problems will be addressed within three working days, either resolving the problem or informing the user how the problem will be resolved.	Through December 31, 2011, 89.8% of queries were addressed within three working days. OLCF exceeded this target.	OLCF user problem resolution time period.	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.	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.
<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. The 2012 user survey incorporates additional specific questions about User Support.	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.
N/A	N/A	OLCF will provide a summary of training events including number of attendees.	At least 4 training events.	At least 4 training events.

2011 Metric	2011 Actual	2012 Metric	2012 Target	2013 Target
<i>Customer Metric 4: Communications with Key Stakeholders</i>				
N/A	N/A	OLCF survey results related to communication with our user community, if any.	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.
N/A	N/A	OLCF will provide representative communications with key stakeholders.	An example of at least one representative communication with users and one representative communication with DOE ASCR.	An example of at least one representative communication with users and one representative communication with DOE ASCR.
<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: 95%	Through December 31, 2011: Cray XT5 (96.4%); Cray XT4 (97.6%); HPSS (99.1%); Widow 1 (99.3%); Widow 2 (99.9%); Widow 3 (99.95%). OLCF exceeded each target in 2011.	Scheduled availability.	85% (lower in FY12 due to the compute system upgrades).	85% (lower in FY13 due to the compute system upgrades).
Overall Availability: 90%	Through December 31, 2011: Cray XT5 (92.9%); Cray XT4 (97.1%); HPSS (98.7%); Widow 1 (98.0%); Widow 2 (99.3%); Widow 3 (99.7%). OLCF exceeded each target in 2011.	Overall availability.	Jaguar: 80%; HPSS 90%; External File Systems 90%	Titan: 80%; HPSS 90%; External File Systems: existing, 90%, new, 80%
<i>Business Metric 2: Resource Utilization</i>				
OLCF will report on INCITE allocations and usage.	CY 2011 INCITE allocations of 930 million hours. Through December 31, 2011, INCITE usage was 995,214,895 core-hours, or 107.0% of the total allocation.	OLCF will report on INCITE allocations and usage.	Report only, no target.	Report only, no target.
<i>Business Metric 3: Capability Usage</i>				
At least 40% of the consumed core hours will be from jobs requesting 20% or more of the available cores.	Through December 31, 2011, the capability usage was 54%. OLCF exceeded the target in 2011.	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.	At least 30% of the consumed node hours will be from jobs requesting 20% or more of the available Opteron nodes.
N/A	N/A	OLCF will report GPU usage in CY 2013.	N/A	Report only, no target.

2011 Metric	2011 Actual	2012 Metric	2012 Target	2013 Target
<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 publications resulting from work done in whole or part on the OLCF systems.	Through December 31, 2011, 670 publications were reported to the OLCF by users or identified by the OLCF. 300 are reportable within OAR guidance.	The OLCF will report numbers of refereed publications resulting from work done in whole or part on the OLCF systems.	Report only, no target.	Report only, no target.
<i>Strategic Metric 2: Scientific Accomplishments</i>				
The OLCF will provide a written description of major accomplishments from the users over the previous year.	Reference Section 3.	The OLCF will provide a written description of major accomplishments from the users over the previous year.	Descriptions of at least 5 major accomplishments.	Descriptions of at least 5 major accomplishments.
<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.	Reference Section 3.	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.	Report only, no target.
<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 we have developed and best practices we have implemented and shared.	Reference Section 4.	The OLCF will report on new technologies that we have developed and best practices we have implemented and shared.	Report only, no target.	Report only, no target.
<i>Innovation Metric 2: Technology Transfer</i>				
The OLCF will report on technologies we have developed that have been adopted by other centers or industry.	Reference Section 4.	The OLCF will report on technologies we have developed that have been adopted by other centers or industry.	Report only, no target.	Report only, no target.
<i>Is the Facility effectively managing risk?</i>				
<i>Risk Management</i>				
The OLCF will provide a description of major operational risks.	Reference Section 5.	The OLCF will provide a description of major operational risks, including realized or retired risks.	Report only, no target.	Report only, no target.